The London School of Economics and Political Science

Essays on Growth and Fiscal Policy

Julio Brandao Roll

A thesis submitted to the Department of Economics of the London School of Economics and Political Science for the degree of Doctor of Philosophy

London, June 2025

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

Chapter 3 of this thesis is joint work with Maarten De Ridder, Simona Hannon, and Damjan Pfajfar. In this chapter, I contributed 25% of the work. The views in this chapter are those of the authors and do not necessarily reflect those of the Board of Governors of the Federal Reserve System, the Federal Reserve Bank of Cleveland, or their staff.

I declare that my thesis consists of approximately 55,500 words (87,000 words with appendices).

Acknowledgments

I am indebted to several people who played important roles and supported me throughout my Ph.D. and in writing this thesis, either directly or indirectly.

I am enormously grateful to my supervisors Xavier Jaravel and Maarten De Ridder for guiding me, both in writing this thesis and in becoming a better researcher. They have been extremely generous with their time and patience, and showed unwavering support. I thank Xavier specifically for his constant support in pushing me to do better at my research with direct, but gentle, feedback. His encouragement was vital, especially at times when I thought the research work was not good enough. The empirical skills I learned with him have helped me in all chapters of this thesis and will be invaluable in my future as a researcher. I thank Maarten specifically for fruitful discussions on both macroeconomic modeling and research writing. His generosity in helping me craft better and clearer sentences was a gift. I am also thankful that he saw the potential in me to contribute with the work in Chapter 3. Further, I thank Matthias Doepke for helping me see the "big picture" contribution of Chapter 1, which is my Job Market Paper. It showed me the importance of properly communicating the findings of my research and it taught me the skills to do it.

I am also grateful for being able to have helpful conversations with great members of the LSE faculty. In particular, I want to thank Philippe Aghion, Oriana Bandiera, Alessandro Gavazza, Jonathon Hazell, Camille Landais, Alan Manning, Guy Michaels, Ben Moll, Joana Naritomi, Jane Olmstead-Rumsey, Steve Pischke, Ricardo Reis, Mark Schankerman, Johannes Spinnewijn, and Daniel Sturm. It was a privilege to have access to such a vast source of support and inspiration. I also want to thank my co-authors, Simona Hannon and Damjan Pfajfar, for their support and kind words throughout this journey.

My colleagues at the LSE have also provided invaluable support. In particular, I thank Fraser Clark, Gabriela Deschamps Ochoa, Arnaud Dyèvre, Tiernan Evans, Andres Fajardo Ramirez, Lin Fan, Nilmini Herath, Yuxiao Hu, Ananya Kotia, Peter John Lambert, Gabriel Leite Mariante, Runhong Ma, Isaac Martinez Centeno, William Matcham, Roberto R. Maura-Rivero, Sidharth Moktan, Thomas Monk, Canishk Naik, William Parker, Derek Pillay, Rebecca Rose, Soroush Sabet, Yasaman Saeidi, Yannick Schindler, Patrick Schneider, Javad Shamsi, Ameek Singh, and Caterina Soto Vieira. A special mention goes to Ananya Kotia and Sidharth Moktan, who made the journey significantly more fun. Finally, I want to thank colleagues outside the LSE. In particular, Scott Behmer, whose conversations have kept me sane, and Felipe Bordini do Amaral, whose friendship since high school is both a treasure and a delight.

My deepest thanks go to a loving family and close friends who have helped me persevere on this academic path. In particular, to my parents, whose excitement for my achievements has been a valuable source of inspiration. To my brothers, my grandmother, Neuza, and my aunts and uncles for keeping the home I left in Brazil an ever-lasting constant in my life. To my cousins, who fate decided they would be in the UK while I was doing my Ph.D.. And to my close friends, Alvaro, Arthur, Caio, and Elcio, for helping me push the boulder up the hill, and to Felipe and Gustavo, whose long friendship is close to my heart and has defined a home to where I can always go back.

Lastly, I dedicate this thesis to two people. First, to my grandmother, Eneida. Although she never understood why I had left Brazil, she showed utmost delight in seeing me happy. Thank you for your unconditional love. Your memory is and will continue to be a constant blessing. I only wish you could see me graduate. Second, to my fiancée, Emma. I am the luckiest man alive to have met you, though fate would have it that such moment happened a few months prior to my departure to London. I will spend the rest of my life trying to show my gratitude for your unwavering support and love.

Abstract

This thesis consists of one essay on economic growth and human capital, and two essays on local fiscal policy.

In Chapter 1, I raise and test the hypothesis that the effect of human capital on economic growth depends crucially on the concentration of high-skilled labor across firms. First, I causally estimate that new colleges had a positive impact on local economic growth in municipalities with lower concentration of high-skilled labor, but a negative effect in municipalities with higher skill concentration. Second, I isolate the causal effect of changes in local high-skilled labor concentration on local growth. Third, I develop and estimate an endogenous growth model, which quantitatively matches the preceding results and which I use to assess policy counterfactuals.

In Chapter 2, I show evidence of novel heterogeneity in local fiscal multipliers. First, I present evidence from the UK of an average local multiplier of 1.69 and 1.71 for services and capital spending, respectively. There are, however, significant inter-council differences in multiplier estimates which are unrelated to variation in local MPCs. I rationalize my results with a model of heterogeneous labor and productivity shocks that impose a psychological toll on workers' cognitive load capacity. Results show potential gains from removing fiscal misallocation between councils.

In Chapter 3, we examine the short-run effects of education expenditures on local income and employment. We estimate fiscal multipliers using city-level exposure to the US Federal Pell Grant Program. An increase in grants by 1 percent of a city's income raises local income by 2.8% and local employment by 1.9% over the next two years. The higher multiplier is partly driven by Pell grants enabling students to take up student loans. Multipliers are also higher during recessions than in expansions, suggesting that Pell grants can be an effective tool for countercyclical policy.

Contents

List of Figures

List of Tables

Chapter 1

Human Capital and Growth: The Role of high-skilled labor Concentration

1.1 Introduction

Why do increases in human capital supply, particularly in developing countries, not always lead to higher growth rates? A longstanding macroeconomic literature associates improvements in human capital with higher economic growth, either through better labor productivity (Mankiw et al., 1992) or innovation (Romer, 1990). This has led to a push in the last decades by governments and international organizations for an accumulation of skills, especially in middle-income countries, under the assumption that low-growth countries lack human capital. The empirical evidence on this positive association, however, has produced mixed results, mainly due to data limitations and identification issues.[1] For instance, average TFP growth has been lackluster in middle-income countries since the mid-2000s even though high-skilled labor supply soared, as shown in Figure 1.1. This begs the question of whether there is something about skill demand, or lack thereof, that could explain the absence of higher economic growth from more skill supply.

To tackle this question, I propose a new channel that links skill supply and growth via skill concentration at large firms. This new channel works in two steps. First, an increase in skill supply raises skill concentration, defined here as the share of high-skilled people working at large firms over total local supply. This is because a larger supply of human capital benefits large firms more than small ones, helping the former grow even larger. Second, this increase in skill concentration lowers firms' incentives to innovate. As big firms grow larger, the additional profit from further improvements over their own products keeps declining due to lower incremental gains in market share, creating a discouragement effect (Arrow, 1962). Moreover, small firms find it increasingly harder to catch up to big ones as the latter grows larger, which lowers small firms' incen-

[1] C.f. Bils & Klenow (2000), Durlauf et al. (2005), Pritchett (2006).

Figure 1.1: *Evolution of college enrollment and TFP growth in middle-income countries*

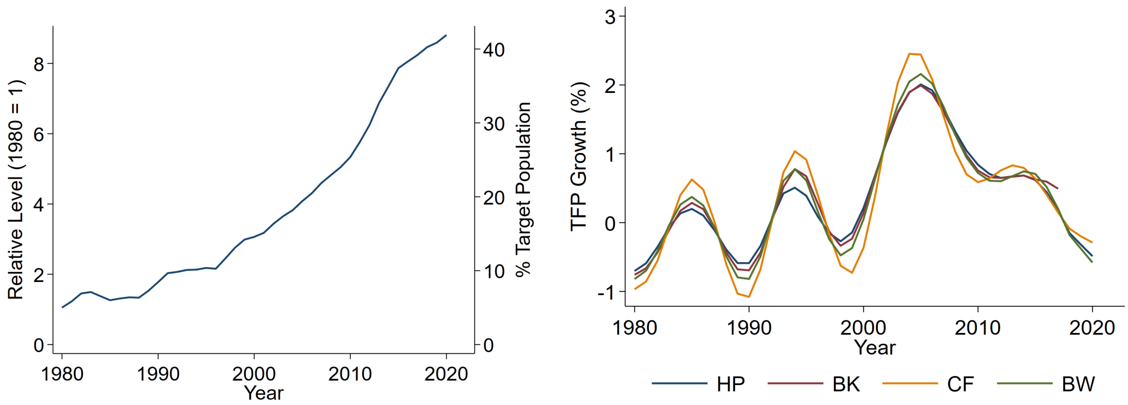

Note: College enrollment data is from the World Bank. Target population is the age group corresponding to college education. Middle-income countries are defined using the World Bank income group classification in 2024 which includes lower and upper middle-income countries (i.e. countries with gross national income per capita between \$1,136 and \$13,845 in 2022). TFP growth data is the trend from the Penn World Table (Feenstra et al., 2015) after using the Hodrick-Prescott (HP), Baxter-King (BK, Baxter & King, 1999), Christiano-Fitzgerald (CF, Christiano & Fitzgerald, 2003), and high-pass Butterworth (BW) filters. Series show 3-year moving averages weighted by population, which is done after filtering for TFP growth.

tives to innovate. These two steps, then, create an offsetting effect that can cancel out, and even overcome, the positive effect on growth that we would expect from a reduction in innovation costs as skill supply increases. This novel channel can, then, explain why successive past increases in skill supply, particularly in developing countries, did not lead to higher growth rates and can actually induce a growth slowdown if skill concentration is high enough.

To test this novel channel empirically, I start by showing that the relationship between high-skill supply and growth depends on local high-skill concentration. I do this using municipality-level data from Brazil on new college and university creation in a difference-in-differences design where I compare municipalities that received a new college with those that did not. The estimation of the causal effect relies on the assumption that the choice of where to open a new college is unrelated to local growth trends. I substantiate this identification assumption by showing evidence of no pre-trends on growth, employment flows, and proxies of college demand, and by showing robustness of results to changes in the control group. Results show that places where high-skill concentration was high before the arrival of a new college grow around 10% less in the long term than places with low skill concentration. This relative difference is due to an initial increase in local growth at municipalities with lower skill concentration, which subsides in the long term, and a decline in long-term growth at places with higher skill concentration, a counterintuitive result. This is evidence that local skill concentration plays a key role in moderating the link between skill supply and growth.

Results also show around a 6% decline in long-term growth from new colleges across all municipalities.

I then investigate the exact mechanism underlying this new channel by splitting it into two steps. In the first step, I exploit the same difference-in-differences design as the one above to pin down the causal effect of higher high-skill supply on local skill concentration. I show that college creation has led to a rise in high-skill concentration of almost 12% a decade after students started graduating, which implies that large firms are the ones hiring most of the new college graduates. This effect is quite significant as a back-of-the-envelope calculation shows that the increase in skill supply in Brazil can potentially explain almost half of the increase in skill concentration in the same period.

This surprising result can be understood through the following thought experiment: assume two firms of different sizes compete in the same market by innovating on their products with high-skilled labor, though only the market leader produces anything since it has the best product variety. Both firms face a strategic incentive to improve their products: the leader aims to make its product harder to copy, while the follower wants to catch up with the leader and become dominant in the market. However, if each innovation improves a firm's marginal cost, then when the leader innovates it is able to extract higher profits from its market share. This profit incentive is exclusive to the leader as the follower does not produce. Hence, when we make innovation less costly by raising skill supply, though both firms want to innovate at a higher rate, the leading firm wants to innovate relatively more because it has an extra incentive to do so. This implies the larger firm will increase its relative share of high-skill hiring, which raises skill concentration.

In the second step, I present novel evidence that local skill concentration has a non-monotonic relationship with GDP growth. To identify the causal link between both variables, we require random variation in local skill concentration. I achieve this through a shift-share instrumental variable (SSIV) design that leverages heterogeneous municipality exposure to national changes in the loan portfolio of the Brazilian Development Bank (BNDES). Importantly, identification relies on changes to loan amounts for different economic sectors being as-good-as-random, an assumption which I test through different falsification tests. As large (small) firms are as-good-as-randomly allocated loans, local skill concentration rises (falls) as firms use such loans to hire skilled labor. Results show that at low levels of concentration, increasing skilled labor at large firms boosts local growth rates. This trend, however, reverses at higher concentration levels when the relationship is negative. These results, which visually characterize an inverted-U relationship, are evidence that incentives for firms to grow depend crucially on skill concentration.

We can, then, join these two steps to fully understand the high-skill concentration channel. As human capital supply increased from new colleges, large firms benefited relatively more than small ones which raised local skill concentration incrementally. Since skill concentration can lead to either higher or lower growth depending on its prior level, the final piece is to understand how Brazil moved along the curve between growth and skill concentration as the latter increased. We can extrapolate the reduced-form estimates to the aggregate economy, abstracting from missing intercept issues, and show that the rise in skill concentration due to the increase in human capital can potentially explain a decline in long-term growth rates of around 18.3% of average growth between 1999 and 2010. Importantly, this is not the net effect of skill supply on growth but the partial effect through the skill concentration channel as we are assuming that the only effect of increasing skill supply is to raise skill concentration. This result is useful as it highlights that the magnitude of the skill concentration channel can be quite significant, particularly in counties where skill concentration was high. As the empirical estimate for the net effect of skill supply on growth is around -6%, we can conclude that the skill concentration channel can more than offset the positive effect of skill supply on growth. Hence, we can link the increase in human capital supply to a slowdown in growth in Brazil. This skill concentration channel is also likely to be applicable to other developing countries which saw a large boost to college enrollment without experiencing higher growth rates.

I rationalize these empirical findings in a model of step-by-step innovation with firm strategic interaction and high-skilled labor search. As in Aghion et al. (2001), two firms compete in a duopoly through a quality ladder where a leading firm can be a number of steps ahead from a lagging one. I then add two novel aspects to my model. First, I require both firms to search for high-skilled labor used in R&D as in the Diamond-Mortensen-Pissarides (DMP) framework (Diamond, 1982, Mortensen, 1982, Pissarides, 1985). Adding search frictions allows firms to shed labor when incentives to hire are low, which is important for the mechanism, and allows the model to capture empirical trends in both high-skill unemployment and skill premium. Second, I make innovation catch up, also interpreted as R&D imitation, a function of high-skilled labor.[2] Though not necessary for the skill concentration channel, this modification allows us to gauge how active R&D imitation changes with respect to higher skill supply, which I then link to innovation diffusion.

I then bring the model to data and show that it can reproduce the non-monotonic relationship between skill concentration and growth. I do so by linking both variables to the technological gap between leading and lagging firms. The intuition is the following. Start with both firms at the same step on the quality ladder. When

[2] An idea made explicit in Cohen & Levinthal (1989).

a firm innovates and moves ahead, R&D competition intensifies as the leading firm wants to defend its profit flow from the laggard's threat while the latter wants to catch up. As such, economic growth increases along with skill concentration as the leader has the additional profit incentive, which leads to relatively more hiring. As the technological gap keeps increasing, both firms face lower incentives to innovate. For the leader, the likelihood of the laggard ever catching up gets smaller and incremental profits from product improvements decline, reducing the marginal benefit of innovating. As for the laggard, a large gap implies a low likelihood of catching up, while any reduction of the gap results in a more intense competitive response from the leader. As such, growth declines though skill concentration increases as the laggard's incentives to invest in R&D fall quicker than the leader's since the latter always earns some incremental profit from improvements in marginal cost. The corollary, then, is that skill concentration goes up while growth increases at first to then decline, leading to a non-monotonic relationship between both variables.

I can, then, use the model to study the effect of an increase in human capital. I start by decomposing the effect of higher skill supply on growth into two opposing channels. On the one hand, more skill supply boosts R&D output through a reduction in hiring costs, which has a positive effect on growth. On the other hand, since the leader benefits more from a higher supply of skills, the average gap between leader and follower increases. This has a negative effect on economic growth due to the stronger disincentives to innovate. Hence, whether growth increases or decreases depends on the strength of each one of these two channels. I show that growth stops increasing at a large enough skill supply and that it even decreases in per-capita terms as the skill concentration channel gets stronger. I also show that the model is able to capture additional empirical trends in Brazil: a decline in the skill premium and an increase in high-skill unemployment, both due to the leading firm becoming unwilling to absorb the extra high-skill supply. I make an important contribution by showing a "double-whammy" effect on the skill premium: skilled wages go down relative to unskilled ones not only because skill supply increases, but also because demand for skilled labor declines due to the skill concentration channel.

Linking a growth slowdown to more high-skill supply is an important contribution for two reasons. First, it explains why education policy might not produce higher economic growth and may even lead to a growth slowdown, a surprising result.[3] What is key here is the role of skill concentration which can make education policy backfire as it ends up helping large firms grow even larger. Second, the possibility of a growth slowdown from more skill supply is also not expected

[3]From the World Bank: "Having a skilled workforce has been recognized as paramount to boosting competitiveness in an increasingly global and interdependent economic environment, fostering innovation and business creation and increasing productivity" (Roseth et al., 2016).

in several endogenous growth models.[4]

Finally, I assess a social planner's role in counteracting the skill concentration channel. I assume the planner is able to tax the local leading firm and use those funds to subsidize high-skilled labor at the lagging firm. Helping laggards "fight back" unlocks the expected growth boost from more high-skill supply (1.6% vs. 1.25%). Results, then, highlight the important role of firm interaction in the link between human capital and growth, particularly in the form of skill concentration. Once we take this into account, the relevant policy lever in places where skill concentration is high becomes not only raising human capital but also innovation by smaller firms, linking education and competition policies.

Related Literature: My work relates to different strands of literature. In introducing the high-skill concentration channel, I shed new light on the relationship between human capital and economic growth. Previous studies have focused on two important channels that deliver a positive association between skills and growth. In the first one, a more educated workforce is more productive (Becker, 1962, Lucas, 1988, Mankiw et al., 1992, Black & Lynch, 1996) due to the higher quality of human capital. In the second one, improvements to human capital boost innovation, either by pushing the technological frontier or through higher adoption rates (Romer, 1990, Aghion & Howitt, 1992, Benhabib & Spiegel, 2005, Toivanen & Vaananen, 2016). My contribution is to propose a new channel where increasing high-skill supply raises skill concentration, which lowers growth. While in practice all three channels happen in tandem, I show that the skill concentration channel is useful in explaining the observed growth slowdown in Brazil and possibly other developing countries where human capital soared. Empirically, I build on work showing the effects of colleges on local outcomes (Abramovsky et al., 2007, Toivanen & Vaananen, 2016, Azoulay et al., 2019, Valero & Van Reenen, 2019, N. Hausman, 2022, Nimier-David, 2023, Cox, 2024). Particularly, I follow Nimier-David (2023) in using college creation in an event study research design to identify the effect of a new college on both high-skill concentration and growth. I add to their results by showing heterogeneity with respect to the degree of local high-skill concentration on the effect of a new college on local growth, which I rationalize in a model of step-by-step innovation.

In terms of both mechanism and model, I build on the literature on endogenous growth, particularly on strategic interaction models (Aghion et al., 2001, Acemoglu & Akcigit, 2012, Liu et al., 2022), to explain my findings. I link my results to the non-monotonicity induced by the "escape-competition" effect, where a market leader invests heavily in innovation to be further ahead of the compe-

[4] I show a derivation based on Romer (1990) in Section A.1 in the Appendix. More generally, the positive link between human capital and growth is a shared feature of models that follow Nelson & Phelps (1966).

tition, and the "lazy-monopolist" effect, where the leader stops investing when it is too far ahead. My contribution lies in adding high-skilled labor demand and search to the step-by-step model which not only introduces the role of skill concentration but also extends results to the skill premium and high-skill unemployment. This paper also extends two previous results. First, I can get lower economic growth, similar to Liu et al. (2022), without requiring low interest rates which did not happen in Brazil (and other developing countries) on the same scale as in the US after the Great Recession. Second, I offer a potential mechanism for the observation made in Akcigit & Ates (2023) that there has been less knowledge diffusion in the US. Although ideas can be understood as public goods, turning them into productivity requires internal capabilities and skills (Cohen & Levinthal, 1989). By incorporating labor-dependent catching up in my model, I show how skill concentration lowers active R&D imitation.

I also contribute to the literature on the rise of firm concentration. This rise, documented for developed countries, has been attributed to different reasons, including less antitrust policy (Döttling et al., 2017), technological change (Autor et al., 2020, Olmstead-Rumsey, 2022, De Ridder, 2024) and diffusion (Akcigit & Ates, 2023), and lower business dynamism (De Loecker et al., 2021). My contribution is in identifying a new channel through which skill concentration increases as large firms benefit the most from more human capital. This channel is particularly useful in the context of developing countries which have experienced a large increase in skill supply (c.f. Figure 1.1). My paper is more closely related to Olmstead-Rumsey (2022) and De Ridder (2024) in using an endogenous growth model to propose a novel mechanism that links the rise in skill concentration to a decline in growth. This paper, however, focuses on the role of human capital.

Finally, this paper makes an important contribution to the recent literature examining the effects of firm concentration in the labor market. Previous papers have linked firm labor market concentration to a reduction in wages (Dix-Carneiro & Kovak, 2015, Azar et al., 2022, Felix, 2022, Schubert et al., 2024), which we also see in the aggregate in Brazil with high-skilled workers. However, I provide causal estimates that the relationship between local skill concentration and the skill premium is non-monotonic, which I rationalize theoretically. This shows how conclusions may differ for high-skill wages. This paper is more closely related to Akcigit & Goldschlag (2023), which shows empirical evidence of inventor concentration at large firms in the US, and Manera (2022), which uses defensive R&D to explain inventor concentration at the sector level. Similar to my results, both studies show that leading firms face lower incentives to implement new ideas once inventors are hired. However, while their results are based on leapfrogging models of innovation, I show how firm interaction and labor market search are crucial to understanding the non-monotonic empirical trends with respect to economic growth.

The remainder of the paper is organized as follows. Section 2.2 describes the data. Section 3.3 explains the empirical strategy and shows estimation results. Section 2.5 rationalizes results with a step-by-step model of innovation with high-skilled labor search. Section 1.5 uses the model to analyze counterfactuals and policy. Finally, Section 3.5 concludes.

1.2 Data

My main data source is the Annual Social Information Report (RAIS) which has annual, non-identifiable socio-economic data on employer-employee links from 1999 to 2017 in Brazil. This includes data on employer, employee, and job characteristics. Employers are required to send their employee data to the Ministry of Labor, which oversees RAIS, and face fines if they do not. As such, the database represents almost the entire labor force under formal employment, which is my main focus since I am interested in high-skilled labor. I do, however, exclude workers in the armed forces, police, firefighting, and politicians from the data, as these are not usually associated with a firm. Importantly, RAIS data identifies an establishment as an employer. This will be relevant when discussing the mechanism behind my results as establishments, in being a smaller constituent of a larger firm, face more intra-municipal competition. As most firms consist of a single establishment, I use "firm" and "establishment" interchangeably throughout the paper.

It is important to specify a few definitions regarding workers and employers used in the empirical estimation. I define high-skilled workers as those who have at least some undergraduate education, though they might not have completed their degree. This group corresponds to around 17.8% of workers in my sample. This broad definition of high skill allows me to capture workers who have the capacity for productivity-enhancing activities regardless of their current occupation.[5] Particularly, it is not uncommon for a worker with a degree in an innovation-related field (e.g. engineering) to be hired in a non-innovative occupation (e.g. financial analyst). This worker, however, could still produce innovation if employed to do so. To show robustness of results, I also use a narrower definition of high skill which I label "high critical-thinking workers." Specifically, starting with the set of aforementioned high-skilled workers we narrow it down to those who are also employed in occupations requiring innovation-prone skills. Importantly, given the Brazilian context I consider throughout the paper the idea that innovation includes not only frontier R&D but also technology adoption and

[5]This is backed by Harrigan et al. (2023) who shows how broadly defined "techies," i.e. technically trained workers, are important for innovation and technology adoption (vs. the narrow "scientists").

incremental improvements. Data on skills by occupation comes from O*NET.[6] As for employers, I classify those with 500 employees or more as being large. If a municipality does not have an establishment that matches this criterion, I consider as large those with 250 employees or more. If there are still no large local employers, I label as large those with 100 employees or more. Large employers, then, correspond to around 14.4% of all employer-employee links.[7]

I also use the RAIS Estabelecimentos dataset which contains similar data to the employer-employee dataset collapsed at the employer level. This allows me to calculate the municipal-level Herfindahl-Hirschman index (HHI) for total employment. This dataset, however, does not have variables separated by different types of employees such as high or low skill. As such, I am limited to calculating a firm-level, HHI-style measure of concentration for total employment, which we can then compare with trends in high-skill concentration.

High-skill concentration in Brazil has increased significantly since the late 1990's. I show this in Figure A1 which plots the evolution of high-skill, non-high-skill, and total employment concentration, calculated as the share of workers at large firms.[8] We notice that only high-skilled labor saw an important increase in concentration at large firms as non-high-skill concentration went up to then mostly decline. To quantify the relative increase in high-skill concentration, I also show in Figure A1 the evolution of the ratio between high- and non-high-skill concentration which increased by almost 50% relative to its value in 1999. This observation is robust to two other measures of skill concentration. First, we observe a similar rise in skill concentration with a HHI-based measure calculated using firm-size bins. Second, there is also little change in the overall trend when we compare this bin-level HHI with a firm-level HHI, a comparison that we can only make for total employment due to data limitations. This provides evidence that the observed increase in high-skill concentration would have remained the same had we been able to calculate a firm-level, high-skill HHI. I show these different skill-concentration measures in Figure A2. These stylized facts indicate that the rise in high-skill concentration could potentially play an important role in explaining empirical trends in Brazil since the late 1990's.

Apart from RAIS, I use data from the General Registry of Employed and Unemployed Workers (CAGED) to calculate the municipal-level net change in total number of workers. While there is no continuous panel data on municipal unemployment, CAGED captures movements in the local unemployment rate as government requires firms to report any hiring and firing of formal workers.

[6]Formally, I define a high critical-thinking worker as those with at least some college education who are employed in occupations at the top skill quartile for one of the following: Math, Science, Critical Thinking, Active Learning, and Complex Problem Solving.

[7]Throughout the paper, I refer to firms not classified as large ones as either small or non-large firms.

[8]All A# Figures and A.# Tables can be found in the Appendix.

However, data points before 2020 do not systematically include temporary workers as firms were not required to report those. As my main concern is high-skilled workers who are generally hired for full-time positions, this does not seem to be a problem. For the municipality-year pairs where data is missing, I input zeros which should be understood as no changes in the number of full-time formal workers.

I also use minimum wage data from the Institute of Applied Economic Research (IPEA) to calculate nominal wages as RAIS provides wage data in units of the national minimum wage. I deflate wages using inflation data from the Brazilian Institute of Geography and Statistics (IBGE). Municipal-level population and GDP are also obtained from IBGE along with data on the share of informal workers which is available in 2000 through the census. Finally, state-level data on electricity consumption, which I will use as a proxy for capital investment, comes from the Energy Research Office, a government-affiliated company. GDP and other firm-related variables are deflated using the GDP deflator.

For my difference-in-differences estimation, I use college-level data from the National Institute of Educational Studies and Research (INEP), which also has data on college quality, between 1999 and 2019. Although the data is non-identified for most years, we can identify any changes to the total number of colleges within a municipality. As for course quality, I rely on two national-level assessments called ENADE (National Student Performance Exam) and the CPC (Preliminary Course Score). ENADE is a test created in 2004 that most college students have to take to graduate which measures both general-level knowledge and content that is specific to degree fields ("broad" and "specific," respectively).[9] Those taking the test must be at the end of their courses. The CPC is a composite indicator of quality, available since 2007, which takes into account the ENADE grade, teaching staff quality, student feedback, and an indicator of learning value added.

To construct my SSIV, I use loan-level data from the BNDES, which also includes information on borrower characteristics. This dataset covers the period between 2002 and 2017. The development bank played an increasingly relevant role in the Brazilian economy throughout my sample, representing around 20% of total bank loans in 2015 according to the Brazilian Central Bank. The BNDES is mainly funded through taxes and the Treasury, and can offer loans to most firms that meet the criteria of its different loan products. Most of its loans, designed to support national development and social causes, have below-market interest rates. Loan eligibility criteria depend on firm size, sector, and the purpose of the loan. Its loan offer is heavily influenced by the Executive Branch of the national government, who can determine funding changes and pick the bank's CEO.

Finally, I use IBGE data from annual sector surveys for the sector-level balance

[9] Though most colleges apply the test, it is only mandatory for private or federal universities.

tests regarding my shift-share strategy. These are run for the following economic sectors: manufacturing, construction, retail and wholesale, and most services. Although data coverage varies over time and between surveys, I am able to compile supply-side data on revenues, value added, intermediate inputs, wages, and number of production workers.

I report summary statistics in Table B.1 in the Appendix.

1.3 Empirical Results

In this section, I first show reduced-form evidence using a difference-in-differences design that an increase in local high-skill supply in Brazil has led to a relative decline in GDP growth in municipalities where skill concentration was high. Using the same empirical design, I also show causal evidence that the increase in human capital led to higher high-skill concentration at large firms. Finally, I present causal estimates for the relationship between local high-skill concentration and GDP growth using an SSIV.

1.3.1 Difference-in-Differences Design: Increase in High-Skill Supply

We first look at how an increase in human capital led to lower GDP growth in Brazilian municipalities where skill concentration was high relative to places where it was low.

To identify the effect of high-skill supply on growth, I leverage data on college and university creation in Brazil between 1999 and 2019. The college-education sector has seen a boom since the late 1990's as a result of government policy. In particular, the 1996 reform which made it easier for institutions to set up programs, the Higher Education Student Financing Fund (FIES) from 1999 which offers subsidized loans to low-income students, and the 2004 College For All program (ProUni) which mainly offers grants to low-income students from public schools. Figure A3 shows the strong increase in both the number of colleges and the population share of college graduates since the 1990's, particularly from private institutions. Importantly, by comparing the flat trends in the share of graduates before 1996 with the steady expansion afterward, it is clear that college supply was being restrained by the legal framework in Brazil before the reform. Moreover, as the growing trend has yet to stop, we can infer that supply has yet to catch up with student demand. We can, then, exploit this substantial expansion of colleges to assess the effects of increasing high-skill supply on growth by comparing municipalities that received a new college in the period ("treated") to those that did not ("control").

Before doing so, it is important to assess how these two groups differ. Both public entities and the private sector might well pick municipalities for new colleges based on specific characteristics that correlate with local economic growth. In particular, we can conceive that for-profit colleges, which constitute around half of the private sector in the 2010's, choose municipalities where student demand is high. This could potentially threaten the identification assumption of the difference-in-differences which relies on the choice of municipality and timing of opening a new college to be as-good-as-random with respect to local growth. To assess this threat, I report in Table A.2 the summary statistics for both treated and control groups on different demographic and economic observables. Importantly, Table A.2 shows that groups do differ on some observables: places picked for new colleges are more populous, have a lower share of minimum wage workers, and a higher share of workers in the Other Services sector relative to those in the Public Sector. However, groups look similar regarding the share of workers in all other sectors and educational profile. Though relevant, these differences between the two groups do not constitute an impediment *per se* to using untreated localities as a control group so long as the choice of where to place a new college was not correlated with local growth trends.

For the case of Brazil after the 1996 reform, the choice of municipality for a new college is *ex-ante* likely to be as-good-as-random with respect to local growth. This is because college supply was suddenly unleashed, implying that it faced significant excess demand, especially after the government launched the subsidy programs in 1999 and 2004. This implies that demand was not the differential factor in choosing one municipality over another. Instead, marginal factors such as the support of a local politician or ties with the local economy became the determining elements behind the choice of where to build a new college.[10] Finally, note that setting up a new college takes years, mostly because the process needs to be approved by the federal government, a step that can take up to three years. This makes targeting particular trends in local observables unreliable.

We can check the data for evidence that our assumption of municipality choice being as-good-as-random holds. I do so in several ways. First, I check for pre-trends to see whether new colleges target economic growth or employment flows. Second, I assess pre-trends on proxies for local college demand and competition to check whether local demand for education did not play a significant role in the choice for a municipality because it was high everywhere. Third, I build a placebo group of municipalities by matching treated and control groups on population level, share earning minimum wage or lower, share who only completed the 5^{th} grade, unemployment rate, and illiteracy rate, all in 2000. I show in Table A.2 the summary statistics of both the placebo and the matched treated groups.

[10] Anecdotal evidence includes a mayor donating land, a representative using funds to support the construction of a campus, and educators with local ties taking advantage of the 1996 reform to open a college.

Although both groups still differ on population size, all other observables are a close match which allows me to test whether differences in observables can explain results. Fourth, instead of comparing places that received a new college with places that did not, I compare the former with places that received a college in the last year of my sample. With the exception of the year when they are treated, these last-treated municipalities provide a valid comparison group with the treated subsample (Sun & Abraham, 2021). Finally, I also compare treated municipalities with those that, at a given time, have not yet received a new college, though will get one in future years. This not-yet-treated group also provides a valid comparison group (Callaway & Sant'Anna, 2021).[11]

Having defined the empirical strategy, I identify the effect of an increase in skill supply on local growth by running the following specification for municipality i at time t:

$$Y_{i,t} = \sum_{\substack{k=-7 \\ k\neq -1}}^{17} \mathbb{1}_{\{D_{i,t}=k\}} \left[\beta_{1,k} \mathbb{1}_{\{HSConc_{i,init} \leq p\}} + \beta_{2,k} \mathbb{1}_{\{HSConc_{i,init} > p\}} \right] + \alpha_i + \delta_t + \epsilon_{i,t} \tag{1.1}$$

where $Y_{i,t}$ is the logarithm of GDP per capita, $D_{i,t}$ is a binary treatment which is equal to one if a college or university were created at municipality i at time t, $HSConc_{i,init}$ is the initial high-skill concentration,[12] p is is the percentile threshold that defines both low and high concentration municipalities, and δ_t and α_i are time and municipality fixed-effects, respectively.[13] Skill concentration is defined as the sum of high-skilled workers in large firms divided by the total number of local high-skilled workers. Importantly, Equation 1.1 allows us to capture whether skill concentration plays a role in the effect of human capital on local growth by comparing $\beta_{1,k}$ and $\beta_{2,k}$. Notice also that, in using the logarithm of local GDP, estimates can be interpreted approximately as the difference in long-term growth between treated and control municipalities.

We can identify the set of $\beta_{1,k}$ and $\beta_{2,k}$ from the assumption of parallel trends. As aforementioned, the intuition behind identification is that the decision and timing of opening a new college are unrelated to local growth trends. We can assess evidence supporting this assumption by looking at pre-trends between

[11] C.f. Section A.3 in the Appendix for results on the falsification tests regarding the identification assumption. Table A.3 also shows results of regressing the year of arrival of the first new college on first-period observables. This balance test assesses whether the timing of college creation correlates with observables, which might bias the comparison between earlier and later-treated places. Only one variable (out of 20) shows a significant correlation, though a one standard deviation increase in the share of minimum wage workers is, on average, associated with an increase of less than 2 years in the year of treatment.

[12] For untreated municipalities, the initial skill concentration is the average concentration in the first two periods for which I have data. For treated units, I use the average concentration in the three years prior to treatment. Results are robust to varying this time window.

[13] While some municipalities report receiving new colleges multiple times, I consider treatment timing to be the year when a municipality reports receiving its first new college.

Figure 1.2: *Difference-in-differences estimates of the effect of college creation on local growth at municipalities with high and low high-skill concentration*

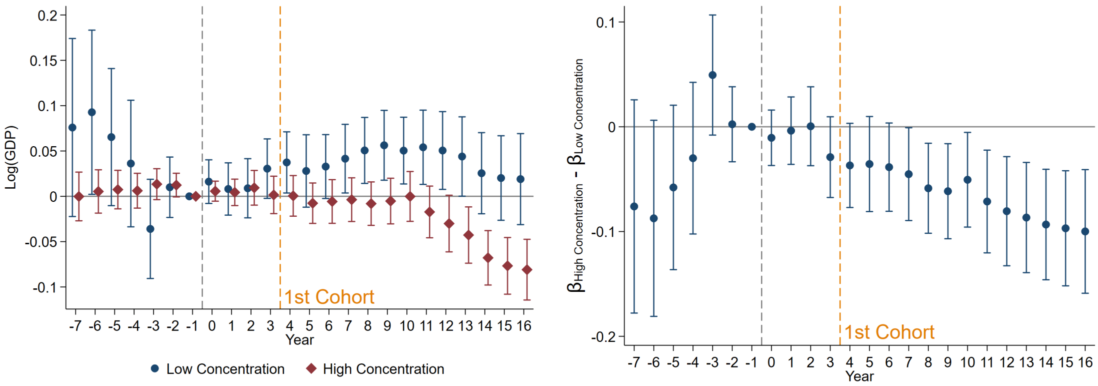

Note: $Log(GDP)$ is the log of local real GDP per capita. High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 14^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. Year is relative to the arrival of a new college and dashed orange line represents the period when the first student cohort is expected to graduate. Vertical bars represent the 90% confidence interval.

treated and control groups. I estimate Equation 1.1 and plot the set of $\beta_{1,k}$ and $\beta_{2,k}$ in Figure 1.2 along with the difference $\beta_{2,k} - \beta_{1,k}$ between municipalities with high and low skill concentration. We notice three important points that provide initial support for our identification strategy. First, results show evidence of no pre-trends for both groups of treated municipalities, which is what we would expect if the choice of where to open a new college is unrelated to local growth trends[14] Second, point-estimates are close to zero in the first three years of treatment, in line with the fact that it takes around four years for the first student cohort to graduate, implying that we should not expect significant effects in the early years after treatment. Third, the difference between low and high skill concentration intensifies in time. This is expected as further cohorts add to the local supply of human capital, as shown in Figure A4.

Results show heterogeneity in the effect of human capital on local growth. After the creation of a new college, we observe a positive and significant effect on local growth in municipalities where high-skill concentration was low before treatment. This effect, however, is not significant in the long term. For municipalities where skill concentration was high, we see no significant results for several years after the first cohort graduates though, surprisingly, coefficients turn negative (and significant) in the long term as growth declines. This highlights how an increase in human capital can result in a decline in growth depending on the local level of skill concentration. The difference between both sets of coefficients

[14]P-values for the joint test of significance: 0.14 (low concentration) and 0.72 (high concentration).

is significant (and negative) at 10% significance level from five years after the first cohort graduates onward. The estimated relative decline in long-term growth in places with elevated high-skill concentration is around 10%, or a 0.9% average relative decline in yearly growth. While the long-term effect of a new college on growth looks remarkable, it is important to notice that in most places the increase in the local supply of skills is relatively quite significant. As I show in Figure A4, after around 10 years since the first cohort graduates the local supply of high-skilled workers increases, on average, by around 90% relative to the pre-treatment average.[15] Finally, I show in Figure A6 that the significant decline in long-term growth is also captured in a specification without the heterogeneity by the level of skill concentration as, across all municipalities, we observe a decline in growth of around 6%. This is evidence that results are not being driven by confounding variables related to the level of skill concentration.[16]

I leave to Section A.3 in the Appendix a series of robustness checks of our results. These include using different estimators, for instance using never-treated, last-treated, or not-yet-treated municipalities as our control group, and estimators robust to heterogeneous treatment effects or non-binary treatment. I also show robustness to restricting the sample to places that do not have colleges, either prior to treatment or for all periods in the control group case, and to changes in the p threshold. Finally, results are robust to using a matched sample on observables, to weighting the specification by the logarithm of local population, and to adding different local-level controls.

Hence, the relationship between human capital and local economic growth depends crucially on skill concentration. Similar increases in skill supply affect municipalities differently depending on whether skill concentration is high or low. Having established this result, we now investigate the underlying mechanism that explains this heterogeneity by shifting our focus to local high-skill concentration.

We start by noting the significant rise in local skill concentration in Brazil since 1999. As shown in Figure A7, high-skill concentration increased around 25 percentage points between 1999 and early 2010's. Naturally, this trend could have different causes. There is a rich literature, mostly on developed countries, linking different mechanisms to a rise in firm concentration, either measured in terms of revenue or total employment.[17] While some of the proposed explanations may apply to Brazil in the same period as my analysis and might explain a rise in

[15] I further show in Figure A5 that the effect heterogeneity cannot be explained by large differences in the evolution of local high-skilled labor between high and low skill concentration places. Figure A5 is also evidence that migration patterns after graduation do not seem to differ significantly.

[16] Places where skill concentration is high or low look similar on observables, as I show in Table A.4.

[17] C.f. Section 3.1.

high-skill concentration, I propose adding a new driver which is the increase in high-skill supply. The intuition behind this channel, which I formalize in Section 2.5, is that large firms are the ones which mostly benefit from the additional supply of high skill because they are able to extract higher markups.

To identify the effect of high-skill supply on skill concentration, I leverage the same empirical strategy as the one I used to pin down the effect of skill supply on GDP growth. That is, I run a specification that is similar to Equation 1.1. For municipality i at time t:

$$\Delta HSConc_{i,t} = \sum_{\substack{k=-7 \\ k\neq -1}}^{17} \beta_k \mathbb{1}_{\{D_{i,t}=k\}} + \alpha_i + \delta_t + \nu_{i,t} \tag{1.2}$$

where $\Delta HSConc_{i,t}$ is the cumulative growth rate of local high-skill concentration between the first period for which we have data and time t.

Similar to Equation 1.1, the identification of the parameters of interest β_k requires an assumption on municipality choice for a new college. Consistent estimation assumes that both control and treated municipalities would have experienced similar trends in skill concentration had there been no treatment. As aforementioned, these two groups do look different in some observables (Table A.2). However, we can apply a similar reasoning to the one used for GDP growth: any imbalance between groups is not a problem as long as the choice of where to build a new college is unrelated to local trends in high-skill concentration at large firms. While I provide evidence of the validity of this assumption, it makes *a priori* sense that it holds for the same reasons discussed in the specification using GDP. Along with checking for the presence of pre-trends, I also assess imbalance by running a robustness check using the matched sample of placebo and treated municipalities.

We can then proceed with estimating the set of β_k. I show results in Figure 1.3. As with our results on growth, Figure 1.3 highlights a few reassuring points. First, there is no evidence of pre-trends, which is in line with the assumption that both treated and control groups would have behaved similarly in the absence of treatment. Second, we observe a similar delay, relative to treatment period, in significant results as the time between starting college and graduation takes around four years. After this initial period, however, results are significant and show an increase in high-skill concentration in large firms due to college creation. The magnitude of the increase is also important as it represents a rise of around 12% in concentration a decade after the first students start graduating.

As with GDP growth, I leave to Section A.3 in the Appendix a series of robustness checks of results on skill concentration. These include using the matched placebo group as control and removing places that do not have colleges prior to treatment or across all periods for the control group. Results are also robust

Figure 1.3: *Difference-in-differences estimates of the effect of college creation on local high-skill concentration*

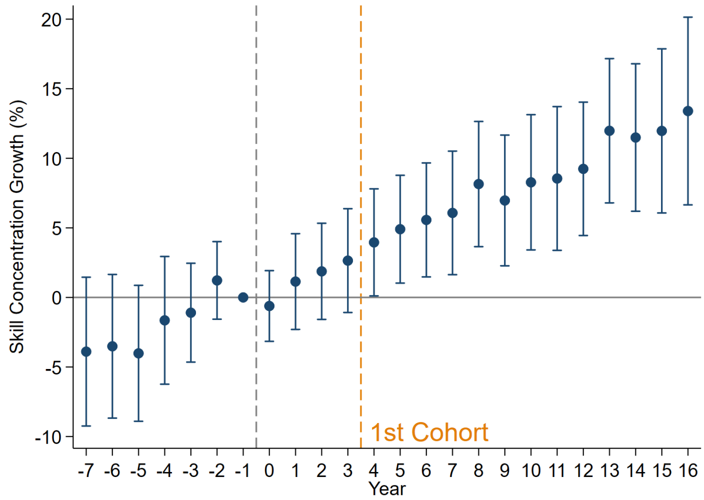

Note: High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Sample excludes observations with no workers at large firms. Year is relative to the arrival of a new college and dashed orange line represents the period when the first student cohort is expected to graduate. Vertical bars represent the 95% confidence interval.

to using estimators robust to heterogeneous treatment effects, with or without local-level controls, or non-binary treatment. Finally, results are unchanged if we add dummies for the arrival of a second or third new college, and to using a HHI-based measure of local high-skill concentration.

Evidence, then, points towards an important role of the steep increase in high-skill supply in Brazil in the rise in local skill concentration. To properly identify the effect of high-skill supply, we had to focus on college creation which limits how we can translate those results to the aggregate economy due to the missing intercept problem.[18] Nonetheless, we can proceed with a back-of-the-envelope calculation to gauge the magnitude of this high-skill supply channel on high-skill concentration. If our conclusions on college creation can be applied broadly to the rise of college graduates, whose numbers more than tripled between 2000 and 2010, the 12% increase in high-skill concentration could potentially explain almost half of the average national increase in skill concentration between 2000 and 2010.[19] Although a simplification, this calculation shows that the high-skill supply channel seems quite relevant in explaining the increase in local skill concentration.

[18]For instance, factor mobility can complicate translating the local estimates to the aggregate economy.

[19]To arrive at this conclusion, note that the 12% increase in local skill concentration is associated with a rise of around 2 percentage-points in the local share of high-skilled people (Figure A4). We then extend our result by assuming that the 4.5pp increase in the national share of high-skilled people in the same period (Figure A7) caused a similar growth in skill concentration as the one measured for new colleges. Finally, we compare this number to the around 24pp increase in aggregate skill concentration.

1.3.2 Shift-Share Design: From Skill Concentration to Growth

After showing the causal link between human capital supply and high-skill concentration, we now proceed with the second step of the skill concentration channel. That is, the relationship between local skill concentration and economic growth.

Results from Section 1.3.1 suggest the relationship between local skill concentration and growth is non-monotonic. The reason for this is the following. We saw that the increase in skill supply has caused both an increase in local skill concentration and a heterogeneous effect on local growth depending on the level of this concentration. If skill concentration plays a role in connecting human capital and growth, then its increase should lead to different effects on growth depending on whether its level is high or low.

We, then, proceed to test this hypothesis by looking at how high-skill concentration at large firms affects GDP growth rates at the municipal level. Our goal is to assess whether the relationship between these variables is non-monotonic. As I do not want to impose a functional form *a priori*, for municipality i at time t the main specification is the following:

$$\begin{aligned} y_{it} = {} & \beta_1 HSConc_{i,t-1}\mathbb{1}\{HSConc_{i,t-1} > p\} + \beta_2 HSConc_{i,t-1}\mathbb{1}\{HSConc_{i,t-1} \leq p\} + \\ & \gamma X_{i,t-2} + \epsilon_{it} \end{aligned} \tag{1.3}$$

where y_{it} is growth in real GDP per capita, $HSConc_{i,t-1}$ is high-skill concentration at large firms, p is a percentile threshold, and $X_{i,t-2}$ are controls which include time and municipality fixed-effects, and a constant for $\mathbb{1}\{HSConc_{i,t-1} > p\}$.[20] I use lagged skill concentration to account for the delay between the hiring decision and actual employee deployment. Effectively, Equation 1.3 estimates two slopes: one for municipalities where skill concentration is relatively low (β_1) and one for places where it is relatively high (β_2). We can then compare the signs of β_1 and β_2 for evidence of non-monotonicity.

As high-skill concentration may depend on other endogenous variables and be affected by local GDP growth, I address endogeneity concerns with an instrumental variable. A possible issue with estimating Equation 1.3 is that a municipality experiencing high growth could be seen by entrepreneurs as a good place to start (or expand) a company. This, in turn, may affect high-skill concentration at large firms, biasing my results. Municipality-specific confounders such as local productivity changes can also pose a threat to identification. I, then, propose a shift-share IV to address this endogeneity. The SSIV is constructed by leveraging heterogeneous exposure to public loans from the BNDES. As explained in Section 2.2, the BNDES loan portfolio, both in terms of size and client character-

[20] Controls are twice lagged to match the timing of the shocks in the SSIV.

istics, is heavily influenced at the national level. While local demand for public loans is affected by local supply and demand conditions, I assume exogeneity relative to Equation 1.3 of changes to the sector-level, national loan amount offered by the BNDES in any given year. This identification strategy consists of the "shift-approach" discussed in Borusyak et al. (2021). As such, I can use the heterogeneity in local-level exposure, measured by sector employment shares, to national changes in loan offer to estimate Equation 1.3 consistently.

Specifically, I instrument Equation 1.3 with the following SSIV:

$$B_{i,t-2} = \sum_n s_{in,t-3} g_{n,t-2} \tag{1.4}$$

where $g_{n,t-2}$ is the sector n shock ("shift") at time $t-2$, defined as the growth rate of the national loan amount, and $s_{in,t-3}$ is the exposure of each municipality i to sector n's shock at time $t-3$, measured as the local-level high-skill employment share in sector n. The SSIV is one-period lagged relative to the endogenous variable to account for the timing between loan issuance and actual spending, and I use the 2-digit Brazilian National Classification of Economic Activities (CNAE) to classify the $n = 1, ..., N$ sectors.

Following Borusyak et al. (2021), the validity condition for the shifts can be written as:

$$E\Big[\sum_t \sum_i B_{i,t-2}\epsilon_{it}\Big] = E\Big[\sum_t \sum_i \epsilon_{it} \sum_n s_{in,t-3} g_{n,t-2}\Big] = E\Big[\sum_t \sum_n \bar{\epsilon}_{n,t} s_{n,t-3} g_{n,t-2}\Big] = 0 \tag{1.5}$$

where $s_{n,t-3} = \sum_i s_{in,t-3}$ and $\bar{\epsilon}_{n,t} = \frac{\sum_i s_{in,t-3}\epsilon_{it}}{\sum_i s_{in,t-3}}$. Equation 1.5 shows how we can rewrite the SSIV orthogonality condition as a condition on the orthogonality of shocks $g_{n,t-2}$. Intuitively, SSIV validity assumes national shocks are uncorrelated with municipality-level confounders and do not systematically favor certain industries in a way that may bias results. We assess this point with falsification tests in Section A.5 in the Appendix.

Before the estimation, it is important to split the SSIV between large and small firms. Using $B_{i,t-2}$, calculated by bundling all firm sizes together, is problematic as loans to large and small firms affect skill concentration at large firms differently. It is reasonable to expect a national increase in loans to large firms to raise skill concentration as those firms increase hiring from the BNDES boost, while an increase in small-firm loans may have the opposite effect. As such, we should separate shocks to loan levels between large and non-large firms. Moreover, since the BNDES has different loan programs by firm size, there is enough variation to warrant the split. Effectively, I separate shocks $g_{n,t-2}$ and shares $s_{in,t-3}$ of small and large firms as if they were from different sectors, and calculate two shift-share instruments: $B_{i,t-2,large}$ and $B_{i,t-2,small}$.[21] I then instrument $HSConc_{i,t-1}$ in Equation 1.3 with both SSIVs, each interacted with $\mathbb{1}\{HSConc_{i,t-1} > p\}$, totaling four IVs.

One concern with using these instruments is that they may affect other firm-level inputs which could confound the effect of skill concentration on growth. For instance, an increase in BNDES loans to large firms increases not only high-skill hiring but also non-high-skilled labor and investments in capital, both of which could increase revenues and affect GDP growth. While capital takes longer to adjust than labor, changes in non-high-skill hiring are a potential issue. To deal with the latter, I add total non-high-skill hiring as a control in Equation 1.3, treating it as endogenous and instrumenting it with the additional IVs.[22] To further highlight the particular role of high-skill concentration, I also add one specification where I control for the municipal-level employment HHI measure of concentration, similarly instrumented with the available IVs.

Even though capital adjusts more slowly than labor, we may still worry about a violation of instrument validity from capital investing. As there is no municipality-level data on capital stock, I proxy investment with changes in electricity consumption which is available at the state level.[23] To get municipality-level variation, I multiply the per-firm, per-worker consumption with the local number of firms.[24] I then divide local consumption by local GDP and add the change in local electricity consumption as a control variable, which I treat as endogenous and instrument with the IV set, in Equation 1.3.[25]

I report 2SLS results for Equation 1.3 in Table 1.1. I set p to the 16^{th} percentile to maximize IV relevance.[26] Column (1) only includes fixed-effects while Column (2) adds local-level controls. Column (3) adds the 2000 local informality share interacted with year fixed-effects as a control. Columns (4)-(8) assess the potential bias from non-high-skill hiring (4), capital formation (5), total employment concentration (6), and both non-high-skill hiring and capital formation (7-8). I

[21]Shocks are winsorized at the 3^{th} and 97^{th} percentiles to avoid results being driven by extreme values.

[22]We may worry that controlling for non-high-skill hiring might introduce bias via a "bad control" problem. I report Monte Carlo simulations in Section A.2 in the Appendix that validate the identification strategy.

[23]An idea that goes back to Taylor (1967) and Moody (1974).

[24]I use electricity consumption of the manufacturing sector as it likely correlates more strongly with capital spending. Results are unchanged if I include service sector consumption as well.

[25]Table A.5 shows that the SSIV leads, as expected, to more hiring of high-skill and high critical-thinking workers. The effect on non-high-skill hiring is different whether loans are for large or small firms. Results also indicate that the SSIV using both small and large firm loans lowers the ratio of per-worker electricity consumption over GDP, evidence that capital formation is not happening at significant levels.

[26]Though relevant for the estimation, the choice of threshold p does not matter for the conclusion on non-monotonicity. If $f(X)$ is the true function that relates dependent and independent variables, we want to choose a value of p that is close to the point where $f'(X) = 0$, i.e. a local minimum/maximum. I show robustness to the choice of p in Section A.5 in the Appendix.

add a third SSIV in Columns (5)-(8) which is the same as the one in Equation 1.4 except that I do not separate shocks to large and small firms and I use total employment share (vs. high-skill shares). As joint IV relevance declines when we add all endogenous variables, I run in Column (8) the same specification as in Column (7) using the Limited Information Maximum Likelihood (LIML) estimator instead of 2SLS as the former has lower small sample bias due to weak instruments.[27] Joint F-statistics are above the usual weak-IV threshold in all specifications except (7) and (8) though the negligible change in estimates between 2SLS and LIML suggests a small bias. I assess this point further by reporting the effective F-statistics (Olea & Pflueger, 2013) for $HSConc_{i,t-1}\mathbb{1}\{HSConc_{i,t-1} > p\}$ and $HSConc_{i,t-1}\mathbb{1}\{HSConc_{i,t-1} \leq p\}$ separately, along with the respective critical values at significance level 5% and a 10% "worst-case" bias. Effective F-statistics are above the critical values in all specifications. Finally, we do not reject the null for the J-test of overidentification. This provides initial support for IV validity, a point which I analyze further for the SSIV in Section A.5 in the Appendix.

Results in Table 1.1 show a non-monotonic relationship between skill concentration and local GDP growth. In all specifications, an increase in skill concentration at large firms increases GDP growth in places where this concentration is low to begin with. This effect, however, reverses in places where skilled labor was already highly concentrated. This is in line with our initial hypothesis that skill concentration is the key channel between skill supply and growth given results in Section 1.3.1. Results are significant in all specifications and show that a growth slowdown can be induced by an accumulation of skilled labor at large firms. Controlling for non-high-skill hiring and our proxy for capital formation does not change coefficients significantly which indicates that potential biases from changes in other inputs are less of a concern here. Moreover, controlling for total employment concentration in Column (6) shows that results are specific to high-skill concentration.

We can see the non-monotonicity visually in Figure A8 where I plot the binned scatter plot between growth in local GDP per capita and the predicted value from the first stage of the 2SLS estimation (i.e. $\widehat{HSConc}_{i,t-1}$). As the threshold p is defined over high-skill concentration level and not over the first-stage predicted values, I also plot the binned scatter plot where I split the values of $\widehat{HSConc}_{i,t-1}$ between those where high-skill concentration ($HSConc_{i,t-1}$) is below or above the threshold p. We can clearly observe the non-monotonic shape, which visually constitutes an inverted-U.

I also assess the importance of sector-level correlation by calculating exposure-

[27] While LIML is known to be inconsistent under heteroskedasticity and many IVs, the bias is small when the number of IVs is < 10 (J. A. Hausman et al., 2012). All specifications include the local sum of shares $s_{i,t-2} = \sum_n s_{in,t-2}$ as recommended in Borusyak et al. (2021). Results using heteroskedastic LIML are unchanged.

Table 1.1: Effect of high-skill concentration in large firms on local GDP growth in places with high and low concentration

	GDP Growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}$=0 × $HS\ Conc._{t-1}$	0.973***	0.903**	0.972**	1.275***	0.889**	0.876**	1.291***	1.288***
	(0.286)	(0.287)	(0.301)	(0.360)	(0.284)	(0.292)	(0.380)	(0.378)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}$=1 × $HS\ Conc._{t-1}$	-0.489***	-0.489***	-0.499***	-0.570***	-0.489***	-0.481***	-0.573***	-0.572***
	(0.109)	(0.108)	(0.112)	(0.123)	(0.108)	(0.110)	(0.125)	(0.125)
N	74,090	74,090	74,090	74,090	74,090	74,090	74,090	74,090
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Informality			Yes					
Non-High-Skill				Yes			Yes	Yes
Capital Proxy					Yes		Yes	Yes
Employment HHI						Yes		
LIML								Yes
Joint F-statistic	32.4	32.6	28.9	13.0	14.6	25.7	6.0	6.0
J-test, p-value	0.12	0.13	0.17	0.94	0.18	0.15	0.99	0.99
OP F-statistic, $\mathbb{1}_{\{HS\ Conc._{t-1}>p\}} = 0$	38.0	38.2	35.2	33.8	37.2	36.1	30.7	30.7
OP Critical Value, $\mathbb{1}_{\{HS\ Conc._{t-1}>p\}} = 0$	16.7	16.8	16.5	15.5	16.5	16.5	14.8	18.2
OP F-statistic, $\mathbb{1}_{\{HS\ Conc._{t-1}>p\}} = 1$	41.8	41.4	37.7	48.9	39.5	38.9	47.3	47.3
OP Critical Value, $\mathbb{1}_{\{HS\ Conc._{t-1}>p\}} = 1$	13.4	13.4	12.7	5.2	12.1	12.3	4.9	12.9

High-skill concentration ($HS\ Conc$) is the local share of high-skilled workers at large firms over total local supply. Threshold p is set at the 16^{th} percentile. GDP Growth is real per-capita local GDP growth winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(4) use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as IVs, each interacted with $\mathbb{1}\{HSConc_{i,t-1} > p\}$. Columns (5)-(8) add an SSIV calculated using both small and large shocks together, and using total employment shares. All specifications control for the lagged local sum of shares $s_{i,t-2}$. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. Non-High-Skill refers to total non-high-skill hiring. Capital Proxy refers to the proxy variable calculated using the change in local electricity consumption. Employment HHI refers to the HHI measure of concentration calculated for total employment. J-test refers to the overidentification test in Hansen (1982). OP F-statistic and Critical Value refer, respectively, to the Olea-Pflueger effective F-statistic and the critical value for a 5% significance level and a 10% "worst-case" bias. Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

robust standard errors. One concern with the "shock-based" identification strategy is that localities with a similar sectoral composition (i.e. similar shares $s_{in,t-3}$) may have correlated errors in Equation 1.3 which are not taken into account when we use municipality-clustered standard errors. Adão et al. (2019) develop "exposure-robust" standard errors which can be extended to a case with an interacted endogenous variable and multiple SSIVs.[28] I calculate these robust standard errors for specifications in Columns (2) and (3) in Table 1.1. For Column (2), the exposure-robust errors are 0.356 (5% significance CI = [0.206,1.600]) and 0.161 (5% significance CI = [-0.803,-0.174]) for the bottom and top coefficients, respectively. As for Column (3), the exposure-robust errors are 0.307 (5% significance CI = [0.371,1.574]) and 0.154 (5% significance CI = [-0.801,-0.197]) for the bottom and top coefficients, respectively. Although robust standard errors are larger, coeffi-

[28]I follow one of the author's additional notes on extensions to cases with multiple regressors and IVs.

cients remain significant, as shown by the confidence intervals, and conclusions are unchanged.

Results, then, complement our findings on human capital supply in identifying skill concentration as the underlying channel that explains the heterogeneity in results on GDP growth. We started this analysis by showing evidence that a boost to local skill supply only led to higher economic growth in municipalities with low skill concentration, and only in the short term. In the long term, places where skill concentration was high see a relative decline in GDP growth. The channel that explains these findings can be summarized as follows. As high-skill supply increases, local large firms benefit relatively more than small firms, increasing skill concentration. The latter, however, has a non-monotonic relationship with growth depending on the level of skill concentration. Hence, an increase in high-skill concentration, due to higher skill supply, causes higher growth in places with low skill concentration and lower growth in places where such concentration is high.

As with the link between skill supply and concentration, we can gauge the economic importance of the link between skill concentration and growth. As shown in Figure A7, skill concentration increased from around 42% in 1999 to around 67% in the 2010s. We can then use our baseline estimates in Column (2) of Table 1.1 to assess the potential change in growth rates from the increase in skill concentration. As before, local estimates do not translate easily into the aggregate economy due to the missing intercept problem. Nonetheless, this exercise is useful to gauge whether the high-skill concentration channel seems relevant or not. To do so, we assume all municipality-year pairs undergo a 25 percentage point increase in their local high-skill concentration. We can then calculate the population-weighted average change in growth rates. Doing so implies a long-term decline of 1.07 percentage point, or almost half of per-capita real GDP growth in the period for Brazil.[29] Even if we only consider the increase in skill concentration from the increase in high-skill supply, I estimate a decline in long-term growth of around 18.3%.[30] Notice that we assume here that the only effect of an increase in skill supply is to raise skill concentration. As such, the 18.3% decline represents the partial effect of skill concentration in the link between human capital and growth, whereas the total net effect corresponds to the decline shown in Figure A6 of around 6%. Estimates show that the skill concentration channel can more than offset the positive effects of a larger skill supply on growth.

I leave to Section A.4 in the Appendix additional results on the causal relationship between skill concentration and the skill premium. Leveraging the same

[29] I assess the increase of 25 percentage-points over a period of 11 years. The national per-capita growth rate between 1999 and 2010 was 2.25%.

[30] To arrive at this number, I use the estimated increase in aggregate skill concentration from the aggregate change in high-skill supply using results from Section 1.3.1.

SSIV research design, I show that this relationship, as with GDP growth, also follows an inverted-U shape. This will be useful when we assess model validation in Section 2.5.

Finally, I show in Section A.5 in the Appendix a series of falsification tests of the SSIV research design along with robustness checks of our results. The former consists of several balance tests, both at the municipality- and at the sector-level, to check whether sectors that experienced large shocks look similar on observables to those that experienced small shocks. On the latter, I first show that the non-monotonic results are robust to using polynomial regressors and IVs instead of interacting both with $\mathbb{1}\{HSConc_{i,t-1} > p\}$. Second, results remain unchanged if we use a narrower definition of high-skilled workers which includes information on the type of skills required for different occupations. Third, results are robust to restricting the sample to the non-tradable sector only, which is reassuring as my skill concentration mechanism involves competition in local labor markets. Finally, I show robustness to several additional changes to the specification, including weighting the specification by the log of local population, lagging the SSIV exposure shares one additional period, and to changes in the threshold p.

1.4 Model

I now rationalize my findings from Section 3.3 in an endogenous growth model with high-skilled labor demand and search. I first describe the model's framework. I then show GMM estimation results and how they relate to the empirical findings in the previous section.

1.4.1 Model Framework

We start with a closed economy in continuous time and a unit continuum of markets j where two firms compete in a technology ladder in each market (similar to Aghion et al., 2001).[31] In each market j there is also a firm producing a non-innovative good, i.e. there are two goods in each market: one produced by the competing R&D firms and one produced by the non-innovative firm, referred to as i (or $-i$) and o respectively. At any moment in time an innovative firm i is located at step m on the technology ladder, the latter being shared among all markets though firms both within and between markets can be at different steps. Consumers have log-utility preferences, own firms, cannot save or borrow, and provide one unit of work of one out of two types: high or low skill. Intertemporal

[31] While step-by-step models of innovation are usually applied to developed economies, there is relatively less frontier R&D effort in Brazil. Nonetheless, these models can offer useful insights if we consider a broader definition of innovation which includes process innovation and adopting foreign technologies.

preferences and the budget constraint for the representative consumer in market j are as follows:

$$U = \int_0^\infty e^{-rt}\Big\{\nu \ln x_j(t) + (1-\nu)\ln x_{o,j}(t)\Big\}dt$$
$$\sum_{k\in\{i,-i,o\}} p_{k,j}(t)x_{k,j}(t) = w_{LS}(t)l_{j,LS}(t) + \Pi_j(t) \tag{1.6}$$
$$x_j(t) = x_{i,j}(t) + x_{-i,j}(t)$$

where $x_{k,j}(t)$ and $p_{k,j}(t)$ are the demand and price of firm k's product in market j, respectively, $w_{LS}(t)$ is low-skill wage, $l_{j,LS}(t)$ is low-skilled labor, $\Pi_j(t)$ are aggregate profits in market j which are also used to pay high-skill wages $w_{i,j,HS}(t)$, $\nu \in (0,1)$, and r is the discount rate.

Innovative firms i engage in both production and R&D. While innovation requires skilled workers, production uses low-skilled ones. The production function for firm i in market j is:

$$y_{i,j}(t) = \gamma_{i,j}(t)l_{i,j,LS}(t) \tag{1.7}$$

where $y_{i,j}(t)$ is output, $\gamma_{i,j}(t)$ is productivity, and $l_{i,j,LS}(t)$ is low-skilled labor. Productivity evolves according to the following law of motion:

$$\gamma_{i,j}(t+\Delta t) = \begin{cases} \gamma^{m+1}, \; if \; R\&D \; successful \\ \gamma^{m}, \; if \; R\&D \; fails \end{cases} \tag{1.8}$$

where $\gamma > 1$ is a constant. Equation 1.8 implies that each successful R&D effort moves the firm one step further on the technology ladder. The arrival rate of successful innovations happens at a Poisson rate $\eta_{i,j}(t)$ which is determined by the R&D production function:

$$\eta_{i,j}(t) = A_\lambda \lambda_{i,j}(t) + A_l l_{i,j,HS}(t)^\alpha \tag{1.9}$$

where $\alpha \in (0,1)$, A_λ, and A_l are constants, $\lambda_{i,j}(t)$ is R&D investment, and $l_{i,j,HS}(t)$ is high-skilled labor. We assume low-skilled labor supply is perfectly elastic and paid at an exogenous wage w_{LS}, while high-skilled labor supply is paid $w_{i,j,HS}$ which is determined through labor search.[32] The cost of investing $\lambda_{i,j}(t)$ is $C(\lambda_{i,j}(t)) = \rho\lambda_{i,j}(t)^2/2$, where ρ is a constant.

Innovative firms compete a la Bertrand.[33] Define the technological gap between two firms in a market as $s(t) = m_i(t) - m_{-i}(t)$. We shall call the firm

[32] I assess the assumption of perfect elasticity of low-skilled labor supply in Sections A.6 and A.7 in the Appendix. I also assess results under different labor market assumptions in Section A.7.

[33] Model results under an assumption of competition a la Cournot remain qualitatively similar.

that is ahead the "leader" (subscript L) and the one that is behind the "follower" or "laggard" (subscript F). As such, for $s > 0$ the leader takes the entire market and charges a price that is the marginal cost of its competitor. For $s = 0$, both firms split the market equally. Note that in this economy different markets j are characterized by a gap s which varies as firms innovate. Then, from log-utility:

$$x_{L,j}(t) = \frac{\nu D_j(t)}{p_{i,j}(t)} \,,\ x_{F,j}(t) = 0 \tag{1.10}$$

where $D_j(t) = D_s(t) = w_{LS}(t) l_{i,j,L}(t) + w_{o,j}(t) l_{o,j}(t) + \Pi_j(t)$ is aggregate demand.[34] It is straightforward to show that the optimal low-skilled labor demand for the leader when $s > 0$ is:

$$l_{i,j,LS}(t) = \frac{\nu D_j(t)}{\gamma^s w_{LS}(t)} \tag{1.11}$$

We can then write the static problem for the innovative firms. Consider a leading firm i that is s steps ahead of the laggard. Profits can be written as (I henceforth drop the time dependency to simplify the notation):

$$\pi_s = \max_{p_{i,j}} \left(p_{i,j} - \frac{w_{LS}}{\gamma^{m+s}} \right) x_{i,j} = \left(\frac{w_{LS}}{\gamma^m} - \frac{w_{LS}}{\gamma^{m+s}} \right) \frac{\nu D_s}{p_{i,j}} = (1 - \gamma^{-s}) \nu D_s \tag{1.12}$$

Given Bertrand competition, follower's profits are zero, i.e. $\pi_{-s} = 0$. When $s = 0$ the industry is "neck-and-neck" and both firms make no profits ($\pi_0 = 0$). Firms decide strategically how much to invest in R&D and how much high-skilled labor to hire given the technological gap s with their competitor. Conditional on s, profits are no longer time-dependent, nor do they depend on where each firm is on the technology ladder.

Regarding the non-innovative firm, it only engages in production via the same production function as in Equation 1.7. However, differently from the R&D firms it employs high-skilled labor in production, i.e. $w_o = w_{o,HS,s}$ and $l_o = l_{o,HS,s}$. This aspect of the model captures an important fact about the Brazilian economy which is that a significant share of high-skilled workers does not work in jobs that require a college degree (38% in 2018, Lameiras & Vasconcelos, 2018).[35] Since such employees still earn more than low-skilled workers, we assume that the non-innovative firm has to hire its workers through search. I show later on when estimating the model that adding a non-innovative firm that hires high-skilled labor does not affect the qualitative results regarding the innovative firms, though

[34] Since high-skill wages are paid out of profits, only $\Pi_j(t)$ shows up.

[35] Realistically, non-innovative firms may hire both high and low-skilled workers for production, potentially with different labor productivities. As allowing both types of labor would not change anything significantly in the model, we make the simplifying assumption that such firms only hire high-skilled labor.

it will prove important quantitatively to match labor market empirical moments. To guarantee the existence of a balanced growth path, we assume the productivity of the non-innovative firm γ_o grows at the same rate as the expected growth rate of γ_s.

high-skilled labor search works similarly to the DMP framework where high-skilled workers are either employed in R&D or searching for work while being unemployed. One important difference relative to the DMP framework is that I make an assumption, explained below, that removes the necessity of keeping track of a firm's current labor force. Let u_s be the unemployment rate when the gap between both innovative firms is s and v_s (v_{-s}, $v_{o,s}$) the vacancies posted by the leader (follower, non-innovative firm) such that $\overline{v}_s = v_s + v_{-s} + v_{o,s}$. Let $M(u_s, \overline{v}_s) = B u_s^{\varphi} \overline{v}_s^{1-\varphi}$, where φ is a constant, be the matching function and define $\theta_s \equiv \overline{v}_s / u_s$ as the labor market tightness. Then the worker flow rate from unemployment to employment is $M/u_s = B\theta_s^{1-\varphi}$ and the vacancy matching rate for a firm posting v_s vacancies is $v_s M/\overline{v}_s = v_s B \theta_s^{-\varphi}$. Let the cost for a firm of posting a vacancy be $C_{v,s} = \kappa v_s^2/2$, where κ is a constant.

We can now define the value functions for high-skilled workers and firms. Let W_s (U_s) be the value of employment (unemployment) for a worker. The value function of being a high-skilled worker is:

$$rW_s = w_{s,HS} + \delta(U_s - W_s) \tag{1.13}$$

where δ is an exogenous separation constant. Equation 1.13 is straightforward: while employed, a high-skilled worker receives wage $w_{s,HS}$ and faces a probability of being laid off.

Conversely, the value of unemployment is:

$$rU_s = b + \frac{v_s}{\overline{v}_s} B\theta_s^{1-\phi}(W_s - U_s) + \frac{v_{-s}}{\overline{v}_s} B\theta_s^{1-\phi}(W_{-s} - U_s) + \frac{v_{o,s}}{\overline{v}_s} B\theta_s^{1-\phi}(W_{o,s} - U_s) \tag{1.14}$$

where b is the value of the outside option. Equation 1.14 describes the value flow for an unemployed worker who can find a vacancy from any firm and move back to employment.

I then make an important change regarding innovation diffusion. In the original set-up (Aghion et al., 2001), the follower pays for an arrival rate of innovation of η_{-s} yet gets $\eta_{-s} + h$, where $h \geq 0$ is a constant that represents the relative easiness of catching up to the leader.[36] Instead, I consider the case where the diffusion parameter is a function of the high-skilled labor currently working at the laggard firm. As such, the follower gets $\eta_{-s} + (h_l l_{-s,H}^{\alpha} + h_c)$, where $h_c, h_l \geq 0$ are

[36] Although this term is referred to as an "imitation" parameter in Aghion et al. (2001), I interpret it here as a parameter that regulates the diffusion of ideas from the innovation frontier to the firm catching up.

constants. To simplify, I assume high-skilled workers at the follower firm work in internal and catch-up R&D at the same time.

The idea behind making innovation catch-up a function of high-skilled labor is two-fold. First, it brings the model closer to reality as firms have to develop internal capacity to absorb external knowledge even when the latter is a public good. Second, it strengthens the link between the R&D efforts of leader and laggard through the labor market. As the leader hires skilled labor, the labor market becomes tighter allowing the leader to indirectly hinder innovation catch-up. Though adding an active catch-up term does not change the model qualitatively, it will be important for results on knowledge diffusion.

We can, then, write the dynamic problem of the innovative firms as a function of s:

$$\begin{aligned}
rJ_s = & \max_{\lambda_s, l_{s,HS}} \Big\{ \pi_s - \rho\frac{\lambda_s^2}{2} - w_{s,HS}l_{s,HS} - \kappa\frac{v_s^2}{2} + [A_\lambda \lambda_{-s} + A_l l^\alpha_{-s,HS} + h_l l^\alpha_{-s,HS} \\
& + h_c](J_{s-1} - J_s) + [A_\lambda \lambda_s + A_l l^\alpha_{s,HS}](J_{s+1} - J_s) \Big\} \\
rJ_{-s} = & \max_{\lambda_{-s}, l_{-s,HS}} \Big\{ \pi_{-s} - \rho\frac{\lambda_{-s}^2}{2} - w_{-s,HS}l_{-s,HS} - \kappa\frac{v_{-s}^2}{2} + [A_\lambda \lambda_s + A_l l^\alpha_{s,HS}](J_{-s-1} - J_{-s}) \\
& + [A_\lambda \lambda_{-s} + A_l l^\alpha_{-s,HS} + h_l l^\alpha_{-s,HS} + h_c](J_{-s+1} - J_{-s}) \Big\} \\
rJ_0 = & \max_{\lambda_0, l_{0,HS}} \Big\{ \pi_0 - \rho\frac{\lambda_0^2}{2} - w_{0,HS}l_{0,HS} - \kappa\frac{v_0^2}{2} + [A_\lambda \lambda_{-0} + A_l l^\alpha_{-0,HS}](J_{-1} - J_0) + \\
& [A_\lambda \lambda_0 + A_l l^\alpha_{0,HS}](J_1 - J_0) \Big\}
\end{aligned} \tag{1.15}$$

where $(\lambda_{-0}, l_{-0,HS})$ refers to the competing firm at $s = 0$.

The dynamic problem in Equation 1.15 can be understood as follows. The leader (first two lines) receives a static flow of profits and pays for investment, the high-skill wage, and the cost of vacancies. At a rate $\eta_{-s} + h_l l^\alpha_{-s,HS} + h_c$ the follower is able to reduce the gap relative to the leader from s to $s-1$. Conversely, the leader is able to increase the gap by one at a rate η_s. The situation is analogous for laggard and neck-and-neck firms.

At this point, I make an important simplifying assumption. Both the firm's dynamic problem in Equation 1.15 and the high-skilled labor search, which yields the high-skill wage, need to be solved simultaneously as both require the firm's value function J_s for all s. Moreover, in the usual search framework labor is a state variable as we have to keep track of how many workers a firm has and solve the problem at every value of the gap s. To make things tractable, I split the firm's decision-making into two steps. First, the firm searches for high-skilled labor until it hires the optimum amount $l^*_{s,HS}$ for its current gap s. Then, it engages in R&D and finds out whether it was successful or not. This can be understood as a *collective hire bargaining* assumption: firms gather all workers they find and make a collective offer to hire all of them at once. This assumption can also be understood from a time-frame perspective: by the time a firm successfully innovates, it has already managed to hire the amount of labor it wants given s, i.e. labor adjusts quicker relative to the time between two innovation steps.[37] As a result, labor is no longer a state variable and we only need the value of labor demand at the steady state for each s.

This simplifying assumption, which effectively implies that firms achieve their desired level of labor demand before engaging in R&D, allows us to get an equation for high-skilled labor demand in steady state where $l_{s,HS}(t) = l_{s,HS}(t+1) = l^*_{s,HS}$:[38]

$$l^*_{s,HS} = (1-\delta)l^*_{s,HS} + v_s B\theta^{-\varphi}u_s \tag{1.16}$$

As for the non-innovative firm, it solves the following static problem every period:

$$\pi_{o,s} = \max_{l_{o,HS,s}} p_{o,s}\gamma_o l_{o,HS,s} - w_{o,HS,s}l_{o,HS,s} - \kappa\frac{v_{o,s}^2}{2} - c_{f,s} \tag{1.17}$$

where $c_{f,s}$ is a fixed cost which we add to make $\pi_{o,s} = 0, \ \forall s$ without loss of generality. This not only simplifies the wage equation later on but also highlights how results about the R&D firms will not depend on the non-innovative sector. Analogous to the R&D sector, demand for the non-innovative good is $y_{o,s} = (1-\nu)D_s/p_{o,s}$. We can then solve Equation 1.17 using Equation 1.16 to get labor demand at the non-innovative firm.

The final step is to solve the labor search problem. I define the net value of a match (i.e. the surplus) as $S_s \equiv W_s - U_s + J_s - V_s$, where V_s is the value function of the firm when it hires no labor, i.e. when collective hire bargaining has failed.[39] To solve the bargaining problem between firm and workers, I adopt the usual Nash bargaining solution. Let ξ be the weight for workers. We can, then, write the surplus as:

$$\xi S_s = W_s - U_s \ ; \ (1-\xi)S_s = J_s - V_s \tag{1.18}$$

Using Equation 1.18 along with the definitions of W_s and U_s in Equations 1.13 and 1.14, we arrive at the following expression for high-skill wage at the leading R&D firm:[40]

$$w_{s,HS} = b + \xi S_s(r+\delta) + \xi B\theta_s^{1-\varphi}\left[\frac{v_s}{\overline{v}_s}S_s + \frac{v_{-s}}{\overline{v}_s}S_{-s}\right] \tag{1.19}$$

[37] I, hence, assume the transitory effect on R&D effort from adjusting labor between $l_{s,HS}$ to $l^*_{s,HS}$ to be of second order.

[38] I derive the optimal investment choice and labor demand in Section A.6 in the Appendix.

[39] To get V_s, we have to solve a version of Equation 1.15 where collective hiring fails. For simplicity, I assume that firms do not invest in R&D when collective hiring fails (though they may do so if labor demand is zero) and pay the same vacancy costs as if hiring was successful.

[40] I provide the formal proof of Equation 1.19 in the Appendix. Notice from the $\pi_{o,s} = 0$ condition that the surplus for the non-innovative firm is zero.

Finally, we require the following labor market clearing conditions:

$$\begin{aligned} L_{HS} &= l_{s,HS} + l_{-s,HS} + l_{o,HS,s} + u_s L_{HS} \\ L_{LS} &= l_{s,LS} + l_{-s,LS} \end{aligned} \tag{1.20}$$

where L_{HS} (L_{LS}) is the total amount of high-skilled (low-skilled) labor available locally.

1.4.2 Model Estimation

We can now solve for the steady state. This requires us to pin down 15 parameters: $\{\xi, \varphi, \delta, \alpha, r, B, \gamma, b, \rho, A_l, A_\lambda, \kappa, h_l, h_c, \nu\}$. First, I set $L_{HS} = 1$ and the low-skill wage w_L to match the in-sample average which is R$1,1734.8 monthly.[41] I then pick $\xi = 0.45$ for the bargaining power parameter following Ulyssea (2010), which is close to the usual value in the literature (0.5). I set the elasticity with respect to unemployment φ in the matching function to 0.5 (Ulyssea, 2010, Dix-Carneiro et al., 2021). I calculate the average separation rate for high-skilled workers in sample and set $\delta = 0.084$. α is set to 0.438 in line with the estimate in Growiec et al. (2023) for a TFP production function. I calibrate r to the average nominal baseline interest rate (SELIC) deflated with the 12-month inflation expectation series for the period between 2000 and 2017. This gets us $r = 8\%$.

As for the matching function scaling parameter B, I calibrate it to the following unemployment flow equation which equates flows from and to unemployment:

$$\delta(L_{HS} - E[u_s]L_{HS}) = BE[\theta_s]^{1-\varphi}E[u_s]L_{HS} \tag{1.21}$$

where we then set $E[u_s] = 6.07\%$ and $E[\theta_s] = 0.48$ to arrive at $B = 1.88$.[42]

That leaves us with nine remaining parameters to estimate: $\{\gamma, b, \rho, A_l, A_\lambda, \kappa, h_l, h_c, \nu\}$. I do so via a GMM estimation using 10 empirical moments: average real GDP per capita growth rate, average municipality-level skill premium weighted by the local number of workers, average labor market tightness, average high-skill wage at non-large firms weighted by number of workers, average skill concentration, average firm profitability, R&D share of sales, average cost of hiring per job, average high-skill unemployment, and share of markets where skill concentration is below or equal to 50%. While there is no 1:1 mapping between parameters and moments, especially since moment fit depends on the distribution of sectors over gaps s, we can associate sets of parameters to their most closely related moments. The R&D investment cost parameter ρ directly influences the R&D investment-to-sales ratio. Similarly, we can pin down the vacancy cost scalar κ with the average cost of hiring. Firm profitability depends only on

[41] To get the annual wage, I multiply the monthly rate by 13 to take into account the mandatory end-of-the-year bonus which is equivalent to a month's payment.

[42] C.f. Section A.8 in the Appendix for details on data moments.

γ. ν influences labor market tightness and high-skill unemployment as the non-innovative sector hires most of the labor supply. These moments, along with the skill premium, are also influenced by the value of the outside option b and R&D labor productivity A_l. Finally, h_l and h_c help us pin down skill concentration, both on average and the sector distribution.

It remains to derive the expression of the growth rate in the model. Note that in steady state both leaders' and followers' productivities grow at the same rate g while the average gap s remains the same. As R&D follows a Poisson arrival, leader productivity improves, in expectation, by $\gamma\eta_s\Delta t$ while the follower's improves by $\gamma[\eta_{-s} + (h_l l^{\alpha}_{-s,HS} + h_c)]\Delta t$. Under such steady state, the inflow and outflow of firms between gap levels s have to balance out. Let μ_s be the share of sectors where the gap between leader and follower is s. Then:

$$\begin{aligned} 2\mu_0\eta_0 &= \eta_{-1} + h_l l^{\alpha}_{-1,HS} + h_c \\ \mu_s\eta_s &= \eta_{-(s+1)} + h_l l^{\alpha}_{-(s+1),HS} + h_c,\ s > 0 \end{aligned} \tag{1.22}$$

where $\sum_s \mu_s = 1$. As such, if we now consider a single sector, growth can be expressed as:

$$\begin{aligned} g_s &= ln(\gamma)2\eta_0,\ s = 0 \\ g_s &= ln(\gamma)\eta_s,\ s > 0 \end{aligned} \tag{1.23}$$

while aggregate growth is simply $g_{agg} = \sum g_s\mu_s$. I provide the formal proof of Equation 1.23 in Section A.6 in the Appendix.

It is worth explaining at this point how I calculate skill concentration in the model. Importantly, not all high-skilled workers are employed at innovative firms, a fact reflected in the data. Yet, for simplicity, we did not split the non-innovative sector between large and non-large firms. However, we will do so now to calculate high-skill concentration. We assume that the high-skilled labor in the non-innovative sector is split endogenously between large and small firms according to a Cournot profit split determined by the productivity levels of the innovative firms. Specifically, if two firms compete a la Cournot in the non-innovative sector, one with productivity γ^{m+s} and one with productivity γ^m, then it is straightforward to show that profits for both large and small firms can be written as:

$$\pi_{o,s} = \left(\frac{\gamma^s}{1+\gamma^s}\right)^2,\ \pi_{o,-s} = \left(\frac{\gamma^{-s}}{1+\gamma^{-s}}\right)^2 \tag{1.24}$$

We can, then, use Equation 1.24 to calculate high-skill concentration assuming that the large firm share of the non-innovative sector is the large firm profit share, i.e. $\pi_{o,s}/(\pi_{o,s} + \pi_{o,-s})$. I define skill concentration in the model as the ratio between employees at large firms, both innovative ($l_{s,HS}$) and non-innovative ($\pi_{o,share} l_{o,HS,s}$), over the total number of high-skilled workers ($l_{s,HS} + l_{-s,HS} +$

$l_{o,HS,s}$).[43] Note, however, that high-skill concentration is at its lowest at $s = 0$ since firms are competing neck-and-neck which implies a minimum model-generated level of 50%. This makes model fit difficult as in reality skill concentration can be below 50%. As a solution, I first calculate high-skill concentration as aforementioned (LC_1). I then calculate a second measure (LC_2) which takes the value of 1 whenever the follower does not hire ($l_{-s,HS} = 0$), is a linear function of s when $l_{-s,HS} > 0$, and at $s = 0$ we assume $LC_2 = 1/s_{min}$, where s_{min} is the lowest value of s where $l_{-s,HS} = 0$. This second measure is more in line with the fact that in reality competition through a quality ladder involves several firms and that at a neck-and-neck state the interaction between firms looks more like perfect competition. Finally, the model-generated skill concentration is the average between LC_1 and LC_2. I show further below that results remain unchanged using different approaches to calculating skill concentration. Importantly, this only affects model fit as it only affects how we calculate high-skill concentration.

I show GMM estimation results in Table 2.5. Overall, model fit is good as data and model-generated moments are close, especially for the growth rate and the skill premium. I further assess the model fit by checking the match relative to a non-targeted moment, i.e. the R&D worker share. Though the non-targeted fit is worse than the targeted ones, it is reassuring that the model-generated value is not far from the empirical moment.

Table 1.2: Model estimation and moment fit

Parameter	Value	Parameter	Value
γ	1.05	κ	1.34
b	0.62	h_l	1.90
ρ	3,084	h_c	0.31
A_l	2.23	ν	0.21
A_λ	29.6		
Moments		**Data**	**Model**
Growth Rate (%)		1.31	1.32
Skill Premium, Large Firms		2.76	2.77
Labor Market Tightness		0.48	0.48
High-Skill Wage, Non-Large Firms		0.58	0.54
High-Skill Concentration		0.59	0.59
Firm Profitability		0.20	0.21
R&D Investing-to-Sales Ratio (%)		0.19	0.21
Cost-per-Hire		0.12	0.11
High-Skill Unemployment		0.19	0.22
Share of High-Skill Concentration $\leq$ 50%		0.38	0.40
Non-Targeted Moment		**Data**	**Model**
R&D Worker Share (%)		0.91	0.70

We can then analyze the firm's choice. I show in Figure A9 the value function of both leader and follower, and the hiring and investment decisions as a function of s. Starting with the value functions, they are both monotonic: increasing for the

[43] I later show robustness of results if, instead, we ignore the non-innovative sector when calculating high-skill concentration.

Figure 1.4: *Left: Growth and high-skilled labor concentration; Right: Skill premium and high-skilled labor concentration, all as a function of the gap s*

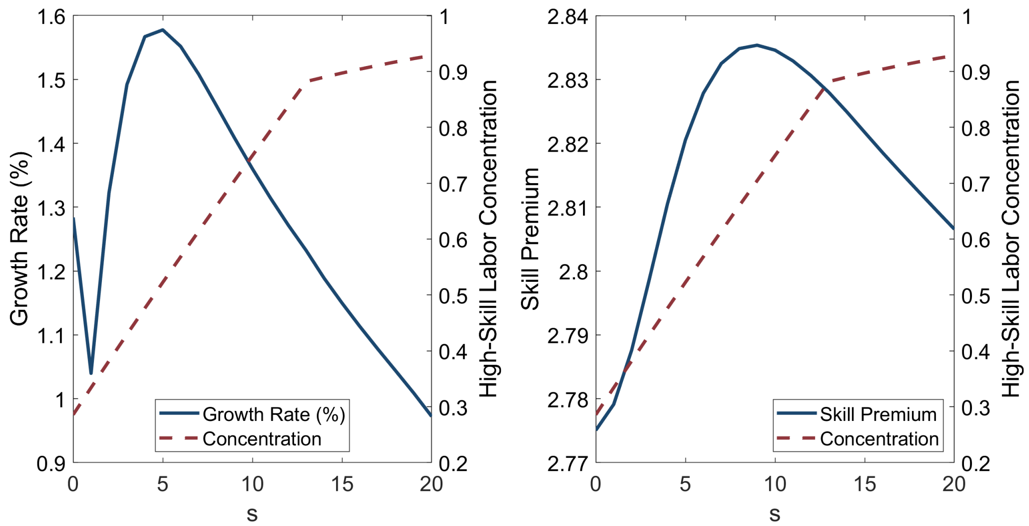

leader and decreasing for the follower. At a high enough s the follower's value function is essentially zero while the leader's value function changes concavity. This identifies the region of most intense innovation effort by the leader as it attempts to escape competition from the follower which can be seen in the peak in both R&D investment and high-skill hiring. This is followed by a reduction in R&D effort in the "lazy monopolist" region where R&D effort falters due to the discouragement effect as the gap is too large for any credible competitive threat. These results are expected in step-by-step models of innovation.

The novelty lies in what we gain by adding high-skilled labor to the model. This can be seen in Figure 1.4 where I show the growth rate, the skill premium,[44] and skill concentration as a function of s. Except for the neck-and-neck region, both plots show non-monotonic curves for the growth rate and the skill premium resembling an inverted-U shape while skill concentration increases. This captures the same patterns estimated in Sections 1.3.2 and A.4 in the Appendix: as skill concentration increases, at first both the local growth rate and skill premium increase. However, as skill concentration keeps increasing, the relationship inverts as growth and skill premium go down.[45] The changes in the growth rate are significant as it moves from around 1.6% at peak to a bottom near 1%. Notice that the reduced-form results capture municipality-level differences, where each local area is at a different gap s. This is why we are analyzing results as a function of s.

Here is the intuition behind the increase in skill concentration as the gap be-

[44]The skill premium in Figure 1.4 is calculated only for the innovative firms. While it is important to take the non-innovative sector into account when matching moments, here I want to highlight the firm interaction in the innovative sector. Results with the non-innovative firm have a similar inverted-U shape.

[45]Figure A10 shows the inverted-U relationships directly for both growth and the skill premium.

tween firms increases. Notice that this happens both in the region of intense competition and when firms shed labor, implying that the laggard lowers its labor demand more intensively than the leader. This stems from two complementary forces. First, the leader has a higher marginal incentive to innovate at $s \geq 1$ because it benefits from the decline in its marginal cost, increasing its profit flow relative to the follower. This higher incentive, however, exists even in the case of constant profits due to a second force: as firms move from $s = 0$ to $s = 1$, the leader starts to receive positive profits. At this stage, it has more incentive to innovate than the follower because a reduction in the gap removes such profits straightaway, while the follower has to innovate twice (i.e. from -1 to 0, then to 1) to experience the same change in profits. Hence, since at any s the follower is farther from the sudden change in profits than the leader, incentives to innovate are higher at the latter.

Figure A9 also highlights the importance of search frictions. We see this clearly if we start from a model with only R&D investment. As firms' incentives to innovate decrease once a leading firm is far ahead, they can adjust investment down accordingly. Once we add high-skilled labor but without labor market frictions, firms cannot shed labor as Equation 1.20 requires the labor market to clear, that is for both firms to jointly hire all available high-skilled labor for all levels of the gap s. This has an important effect on results as it imposes, rather mechanically, that R&D effort from skilled labor does not change with s as total hiring stays the same. Hence, allowing for unemployment is important as it allows firms to adjust labor in tandem with their incentives to innovate.[46]

The model also links rising skill concentration to a decline in active R&D catch-up. By making R&D imitation partly depend on high-skilled labor, we can assess the link between higher skill concentration and changes to innovation diffusion. I show this in Figure A11. We see that active catch-up declines as skills get concentrated at the leader. Importantly, this shows how skill concentration can lead to further disincentives for the laggard: the likelihood of catching up is not only small due to a high gap s, but it also affects the knowledge diffusion from the technological frontier. In other words, it is hard to catch up because the laggard does not have enough high-skilled labor. This relates to the observation in Akcigit & Ates (2023) that the growth slowdown in the US is associated with lower knowledge diffusion. Here, this happens tangibly through skilled labor.

I also assess the robustness to different model specifications. First, I show in Figure A12 the growth rate and the skill premium as a function of s in a version of the model where I remove the non-innovative firm. Although results are slightly different from those in Figure 1.4, curve shapes are similar. This shows that results are not being driven by the inclusion of the non-innovative firm as

[46]Naturally, any framework that is isomorphic to having unemployment would also lead to similar results. I provide further details on this in Section A.7 in the Appendix.

most of them depend on the strategic interaction between the innovative firms. Its inclusion, however, is important for the quantitative fit of the model, particularly with respect to labor market moments. As shown in Table A.6, the model without the non-innovative sector struggles to match the empirical labor market tightness, cost-per-hire, and unemployment, which also affects the other moment fits. I also show robustness of results to different values of the R&D production function labor elasticity α in Figure A13. Finally, results are robust to changes in the convexity of the R&D cost function, assumed as quadratic in the baseline. I show this in Figure A14.[47]

Results remain unchanged if we change how we calculate high-skill concentration. As aforementioned, calculating skill concentration in a duopoly so as to match the data is not straightforward. A different way of doing it from my baseline method is to not use the high-skilled labor employed at the non-innovative sector when calculating high-skill concentration. I show in Figure A16 that estimation results using this measure of skill concentration remain largely the same. Another alternative, shown in Figure A17, is to forego the adjustment using LC_2, i.e. letting high-skill concentration start at 50%. As expected, high-skill concentration levels at low values of s are excessively high. Nonetheless, we get similar results to the baseline: both growth and the skill premium show a non-monotonic, inverted-U pattern as high-skill concentration increases.

As such, my model is able to capture the empirical results on non-monotonicity while extending the results from step-by-step models to high-skill concentration at large firms. We can now use the model to assess two points. First, whether we are able to capture the results in Section 1.3.1 between skill supply, skill concentration, and growth. Second, whether a social planner can boost economic growth given this non-monotonicity.

1.5 Counterfactuals

We will now use the model in Section 2.5 to analyze two counterfactual scenarios. In the first one, we consider the effect on growth from an increase in the aggregate supply of high-skilled labor, showing that it can lead to lower growth. In the second one, we analyze how Brazil could have improved its growth rate from the additional high-skill supply by propping up innovation catch-up through a labor subsidy.

[47] To avoid corner solutions for some values of parameters, I add 0.015 to aggregate demand. I show in Figure A15 that results are robust to using 0.005 instead, showing that this adjustment is largely innocuous.

1.5.1 Counterfactual 1: Increase in high-skilled labor Supply

As discussed in Section 3.1, one would usually expect the correlation between skill supply and economic growth to be positive. This is both the consensus in public policy and the expected result in several endogenous growth models (c.f. Section A.1 in the Appendix for an example). In the case of Brazil, skill supply has soared: the population share with a college degree went from 5.75% in 1991 to 16.8% in 2019 (UNDP et al., 2024). However, this increase in skill supply did not seem to have consistently boosted GDP growth, a surprising result. Moreover, college course quality did not change significantly in the period, as shown in Figure A18, nor did the student composition by degree subject, as shown in Figure A19.[48]

We can, then, ask what happens to the growth rate in our model when we increase skill supply and whether it matches our evidence on the effect of college creation in Section 1.3.1. First, it is important to clarify how total labor supply is affected. More high-skilled labor is usually linked to education which effectively converts low-skilled workers into high-skilled ones, possibly making it harder to hire the former. I, then, assess two scenarios. In one, labeled "external supply," I make no changes to low-skill hiring and high-skill supply grows regardless (e.g. from outside sources). In the other one, labeled "internal supply," low-skill hiring becomes more expensive as education reduces low-skilled labor supply.[49] To make comparisons easier with the data I re-estimate the model matching empirical moments for the initial period between 1999 and 2004.[50]

I then show results for the growth rate in Figure 1.5 along with the evolution of high-skill concentration. We see that the growth rate is not linearly increasing in L_{HS} as it eventually becomes flat and even declines slightly in the external-supply case, a surprising result as total population increases 1:1 which would normally result in higher growth mechanically. To compare this result with the empirical evidence, we can remove the part of growth due to population growth from a higher labor supply, resulting in per-capita values.[51] The per-capita growth curve has an inverted-U shape as the additional high-skill supply is not being put to

[48]The lack of change in quality might come as a surprise as we might expect the quality of the marginal student entering college to decrease with more supply. Some reasons to why that is not the case include financial constraints and a strong distaste for distance, both of which are not necessarily correlated with student talent, and an improvement in college quality from more competition (Cordeiro & Cox, 2023).

[49]For the internal supply case, I assume low-skilled labor becomes more expensive to keep supply constant. In reality, part of this supply also comes from population growth and internal migration. Hence, I assume each factor accounts for 1/3 of the additional high-skilled labor supply and I use 1.66 as the wage elasticity of low-skilled labor supply (Vick, 2017, taking into account the labor force participation rate by sex).

[50]When I do not have data for the 1999-2004 period, I use the moment values for the whole sample. For high-skill concentration, I target the in-sample average in 2000. I show model fit results in Table A.7.

[51]I adjust the share of high-skilled people using the in-sample high-skilled population share.

Figure 1.5: *Effect of increasing high-skilled labor supply on growth and high-skill concentration for different increases in aggregate labor supply*

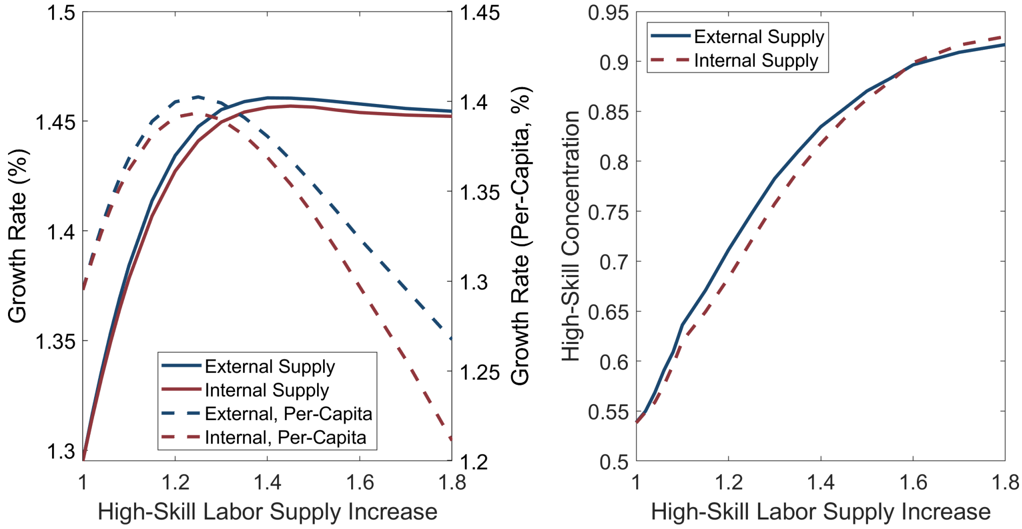

Note: External Supply case corresponds to the scenario where high-skill supply increases from outside sources. Internal Supply case corresponds to the scenario where low-skill hiring becomes more expensive as education reduces low-skilled labor supply. Per-Capita values remove the part of growth due to population growth from a higher labor supply.

use in R&D. This is due to the increase in skill concentration and the average gap s as most of the economy is now in the region where the leader innovates less ("lazy monopolist"). I show this in Figure A21 where I compare the cross-sectional growth rates and the distribution of gaps s for $L_{HS} = 1$ (baseline) and $L_{HS} = 1.5$. The gap distribution shifts to the right with the increase in high-skill supply as the leading firm benefits more from the decline in labor market tightness due to the labor supply increase.

There are two reasons why the leader hires relatively more high-skilled labor when supply increases. First, since the constant catch-up term in the follower's R&D production function is independent of high-skilled labor supply, the relative increase in the follower's R&D effort is lower than the leader's. I show this in Figure A22 (left-hand side plot). If we assess the follower's R&D effort without the catch-up terms ("ex-h"), the lagging firm actually increases R&D effort by more than the leader at low gap levels. However, the constant catch-up term does not fully explain the increase in high-skill concentration. Even if we counterfactually increase the follower's total innovation output to match the growth of the non-catch-up part of its innovation production function, aggregate high-skill concentration still rises to 80.1% (from 53.6%).[52] We, then, still need to understand why the leader's incentive to innovate increases by more than the follower's overall.

To do so, we can split up the firm's value function into two parts. In the first one, the firm derives higher value moving from s to $s+1$ from being able to

[52]To make the comparison as favorable as possible to the follower, I attribute zero growth to its R&D output where the non-catch-up term declines.

charge a higher markup due to a technological innovation. In the second one, value stems from being in a better defensive position as it now takes one additional step, relative to before, for the follower to surpass the leading firm. This split applies analogously to the follower with the important remark that only the dynamic part of the value function matters for the laggard as it does not receive profits. Importantly, both parts ("profit-only" and "dynamic") represent the incentive a firm has to move a step further in the technology ladder.

We can then analyze what happens to these two parts once skill supply increases. As high-skilled labor becomes easier to find, hiring costs as a share of profits decline. Moreover, there is an increase in aggregate demand as an indirect effect of the increased hiring, increasing profits. The relative rise in profits, however, is lower than the relative increase in total R&D hiring. This is because the increase in profits also depends on both low-skilled labor demand in production and high-skill demand at the non-innovative sector, which are only affected indirectly by the increase in skill supply. This implies that at a low gap level where competition is intense, the dynamic part of the value function increases by more than the profit-only part. Hence, the follower's incentive to catch-up rises by more than the leader's as the former only depends on the dynamic part. As such, the leader's position is at a higher threat which implies the dynamic part of its value function can actually decline. We see this in Figure A22 (right-hand plot, "L, Dynamic" vs. "F, Total") as the change in total incentives for the follower is larger than that for the leader at low s.

However, results invert at higher gap levels. As shown in Figure A22, the change in the dynamic part of the leader's value function surges at the point where the leader has the largest incentive to escape the follower's competition, i.e. the frontier between the escape-competition and the lazy-monopolist regions. This is intuitive: as the follower becomes more competitive at low s, the leader wants to avoid reductions of the gap more intensively. As the gap increases from that point onward, both firms see a reduction in the dynamic incentives. For the follower, catching up becomes harder as the gap is larger and any reduction of it implies a strong competitive response from the leader. This induces in the leader a lazy-monopolist effect due to the lower competitive threat. However, the profit-only incentive, which is exclusive to the leading firm, remains as the leader still benefits from improvements in its marginal cost. I show the breakdown between the profit-only and the dynamic parts of the change in the leader's incentive to innovate in Figure A23. We see that the profit-only component is the reason why the leader increases R&D output at higher gap levels. As such, the change in the leader's incentives from an increase in high-skill supply remains positive even at higher values of s, further lowering the follower's incentives to catch up. This

Figure 1.6: *Decomposition of the effect of an increase in skill supply on growth*

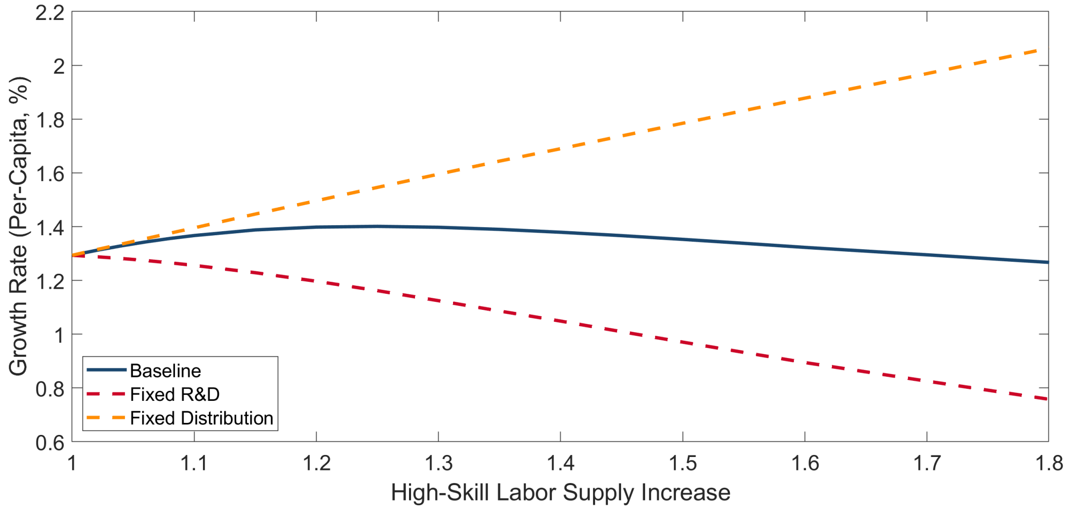

Note: Baseline refers to the baseline estimation results. Fixed Distribution corresponds to the partial effect on growth from fixing the initial distribution of gaps s at $L_{HS} = 1$ and allowing R&D effort to adjust with the larger skill supply. Fixed R&D corresponds to the partial effect on growth from fixing the initial R&D effort and allowing the gap distribution to vary.

explains the rise in high-skill concentration.[53]

Although initially surprising, the decline in growth from an increase in human capital when skill concentration is high can be understood as the net effect of two channels. The first one is the boost to R&D effort, which implies higher growth, when we lower the cost of innovation via a higher supply of skills. This is the usual relationship between human capital and growth in the literature. The second one is the effect on growth from shifting the gap distribution to the right, i.e. the overall increase in the distance between the two firms. I show both channels in Figure 1.6 by either fixing the initial distribution of gaps s at $L_{HS} = 1$ and allowing R&D to adjust with a larger skill supply ("Fixed Distribution"), or by fixing the initial R&D effort and allowing the gap distribution to vary ("Fixed R&D"). The total effect on growth ("Baseline") is the net contribution of each individual channel. This decomposition exercise makes it clear that the model does account for the usual positive effect of human capital and growth. However, it also shows how a high level of skill concentration can lead to a stronger skill concentration channel, which can more than offset the positive effect from lower R&D costs.

We can, then, assess how our model compares with the empirical results in Section 1.3.1. Recall that the increase in high-skill supply from college creation led to a relative decline in GDP growth at highly concentrated municipalities of 10%. This is due to a positive short-term boost to growth where skill concentration was low and a negative long-term decline where it was high. While the

[53]A quick way to see this is to notice that at a large enough gap, both firms have essentially no strategic incentive to innovate, yet the leading firm can still make small profit gains from improvements in marginal cost. This implies a skill concentration of 1 as the leader hires marginally while the follower does not hire.

model did not target these results, it is in good measure to compare the model predictions with our empirical estimates. First, note that the model-generated growth curve in Figure 1.5 is conditional on the initial level of high-skill concentration. This is intuitive: at low levels of skill concentration, more human capital boosts economic growth. At high levels, however, we observe a decline. As such, had high-skill concentration in Brazil been lower (higher), the positive-slope (negative-slope) part of the growth curve would have been longer. I show this in Figure A24. Although skill concentration is an endogenous variable in the model, we can increase or decrease its value by changing the constant catch-up parameter h_c.

We can use Figure A24 to assess whether the model captures the empirical estimates. We proceed in the following way. First, I set the initial levels of skill concentration for the "Low" and "High" scenarios to match the in-sample averages for the low and high skill concentration groups defined in Section 1.3.1, respectively. Second, we know from Figure A4 that a new college increases local skill supply, on average, by around 2pp. Relative to pre-treatment averages, this increase in skill supply represents around a 2.11- and a 1.77-times increase in local skill supply for the low and high skill concentration subsamples, respectively. We can then move along the growth curves to understand the reduced-form results. In the low skill concentration case, growth initially rises which is captured in significant and positive coefficients in Figure 1.2 though model results are relatively larger.[54] As skill concentration continues to increase, however, growth declines. At a 2.11-times increase in supply, the corresponding local growth rate is around 7% higher than the initial value, which produces non-significant coefficient estimates. As for the high skill concentration group, growth starts to decline at a significantly lower level of skill supply. For most of the curve, however, growth remains nearly flat, capturing the non-significant results for places with higher skill concentration in Figure 1.2. At a 1.77-times increase, results imply a decline of around 7.4% relative to initial conditions, matching the long-term decline shown in my reduced-form results. Finally, the model-generated relative difference in growth between "High" and "Low" is a 14.4% decline, a reasonably close value to the reduced-form estimate of around 10%. Overall, results are reassuring as the model is able to broadly capture the untargeted results from Section 1.3.1.[55]

Figure A24 also makes it explicit that the effect of human capital on growth depends crucially on the level of skill concentration. On the one hand, in comparing

[54] The difference in levels for the short-term estimates could be due to differences between steady-state and transition dynamics as the model assumes the former while the reduced-form results capture the latter.

[55] The model also does a good job capturing the effect for all municipalities. On average, new colleges lead to a 1.9-times increase in the local share of high-skilled workers, which corresponds in the model to a 4.85% decline in growth relative to $L_{HS} = 1$ (vs. around a 6% decline, as shown in Figure A6). I assess the transition dynamics of the increase in high-skilled labor supply in Section A.9 in the Appendix.

the “Low” with the “Baseline” case we observe that an increase in skill supply can have the expected positive effect on economic growth for a larger increase in skill supply if skill concentration is low. On the other hand, for a high enough level of skill concentration (“‘Very High” case), per-capita growth is monotonically decreasing in skill supply. We can conclude that the increase in skill supply in Brazil had a negative impact on long-term growth due to a combination of two things: the magnitude of the increase in skill supply and the initial level of skill concentration. As a corollary, it is clear that targeting high-skill concentration becomes an important policy lever to boost growth, a point which I assess in Section 1.5.2.

The increase in high-skill supply also produces other effects in the model that we observe in the data. Regarding the skill premium and high-skill unemployment, there is a remarkable difference between the external-supply scenario and the more realistic case where hiring low-skilled workers becomes harder. I show this in Figure A25. In the former, the skill premium goes up while unemployment declines. This is due to the boost to aggregate demand from the rise in population: as skill supply grows, aggregate demand increases which raises profits. This, in turn, raises low-skill and high-skill hiring, and high-skill wages. Hence, high-skill unemployment declines due to population growth.

In the internal-supply case, however, skill premium declines while high-skill unemployment initially declines to then go up. Importantly, the leader does not absorb the increase in L_{HS} in its entirety as incentives to innovate decline because skill concentration increases. Note that even though a decline in skill premium is expected from an increase in skill supply, the crucial point here is that this decline is not only due to higher skill supply but also lower skill demand. This “double-whammy,” whose breakdown I show in Figure A26, is important to deliver a significant decline in the skill premium which follows what we observe in the data (Figure A27). We also see that the demand-driven partial effect becomes increasingly more relevant in explaining the trend in skill premium as the skill concentration channel becomes stronger, with its share of the total effect increasing from 50% to 60%. Results on unemployment are also reflected in the data and are in line with the increase in high-skill underemployment, as shown in Figure A28.

Finally, the model also links the increase in human capital to lower innovation diffusion. As skill supply pushes high-skill concentration up, the laggard engages less in active R&D imitation even though there are more skills available in the economy. I show this in Figure A29 for both the external and internal-supply cases. This highlights how improving knowledge diffusion from the R&D frontier to followers, either directly or indirectly via high-skilled labor, is an effective measure to increase economic growth.

This analysis shows how we can achieve a non-monotonic relationship between economic growth and high-skilled labor supply through skill concentration. As leading firms benefit more from the increase in skill supply, they increase their gap relative to followers. Once the gap is high enough, incentives to innovate decline which offsets the boost to growth from a larger high-skill supply and leads to an oversupply of high-skilled workers. Results are in line with both aggregate-level data in Brazil and reduced-form estimates in Section 3.3 on the link between high-skill concentration, growth, and the skill premium.

1.5.2 Counterfactual 2: Social Planner

After analyzing the effects of an increase in skill supply on growth, we can ask how a social planner could do better. A shortage of high-skilled labor has been deemed one of the main obstacles to long-term growth in Brazil, prompting a government-induced increase in supply. What I showed, through both causal evidence and the model, is that a larger supply of high-skilled workers does not necessarily lead to more economic growth as skill concentration intensifies. However, I show next that by targeting innovation at lagging firms a planner is able to increase the growth rate by lowering high-skill concentration, weakening the negative effect of the skill concentration channel.

As high-skill supply and concentration increase, and sectors move to the lazy-monopolist state, lagging firms stop actively engaging in R&D, lowering investment and high-skill hiring as incentives to innovate are low. We can then consider the scenario where the social planner provides the follower with innovation inputs by taxing the leading firm and directly sponsoring high-skilled workers at laggards. Specifically, we make the following adjustment to the leader's value function (analogous for the follower):

$$rJ_s = \max_{\lambda_s, l_{s,HS}} \pi_s - \rho\frac{\lambda_s^2}{2} - w_{s,HS}l_{s,HS}(1+\tau) - \kappa\frac{v_s^2}{2} + [A_\lambda\lambda_{-s} + A_l(l_{-s,HS} + l_{s,HS}\tau)^\alpha +$$
$$h_l(l_{-s,HS} + l_{s,HS}\tau)^\alpha + h_c](J_{s-1} - J_s) + [A_\lambda\lambda_s + A_l l_{s,HS}^\alpha](J_{s+1} - J_s) \tag{1.25}$$

where $\tau \in [0,1]$ is the tax on the leader's total high-skilled labor costs, which are then used to finance $\tau l_{s,HS}$ workers at the follower.

I plot model results using the baseline estimation and $\tau = 1\%$ in Figure 1.7. We see that a labor subsidy is quite effective in boosting growth: at an 80% increase in skill supply, growth goes from around 1.25% to 1.6%. Note how the subsidized curve keeps a positive slope for longer, highlighting how the planner can recover the positive relationship between human capital and growth. This is because the subsidy helps followers "fight back" which lowers the average gap between themselves and leading firms. I show this in Figure A30 where I plot the firms' R&D effort for both the baseline and the subsidy cases when aggregate high-skill supply is 1.5. The orange dots determine points of convergence in the

Figure 1.7: *Effect of subsidizing high-skilled labor at the laggard firm*

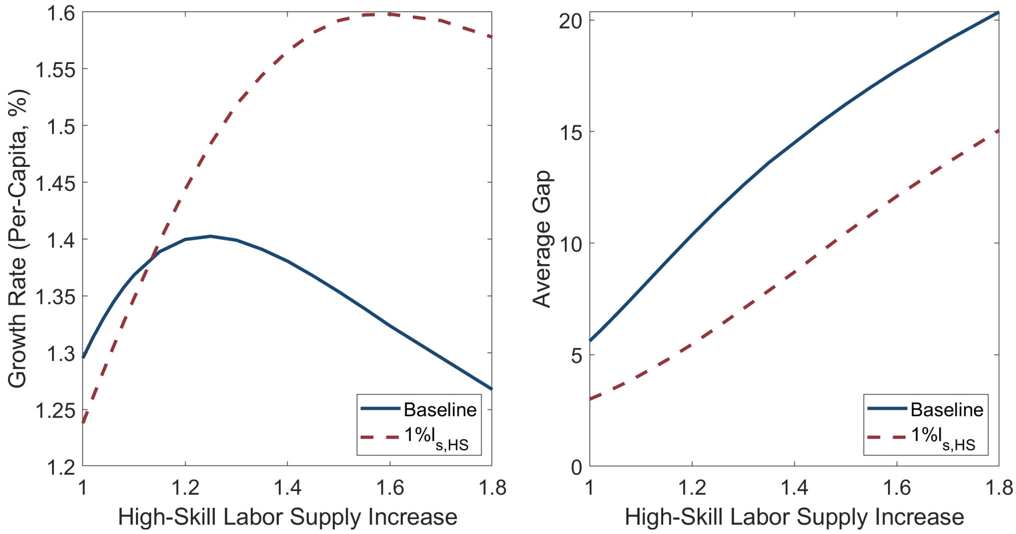

Note: $1\% l_{s,HS}$ refers to the case where $\tau = 1\%$.

gap distribution: while at gap levels below the intersection the leader innovates relatively more which pushes s up, at levels above the intersection the follower innovates relatively more, bringing s down. With the subsidy, this point of intersection moves left (from 1 to 2) to a lower gap level, indicating a lower level of skill concentration. R&D effort, and hence growth, goes up as more intense catch-up increases incentives for the leader to keep innovating so as to escape competition. This is the case even though the leader's R&D output is lower at low levels of the gap due to the tax disincentive.

While the increase in the growth rate from a 1% tax rate looks impressive, note that the change relative to the baseline depends on s. This highlights that the increase in the growth rate when the skill supply increase is large comes from pent-up skill supply. That is, the around 28% increase in growth (vs. a baseline 2% decline) when $L_{HS} = 1.8$ relative to the case where $L_{HS} = 1$ is due to the 80% increase in skill supply that is more appropriately being employed in innovation. Moreover, the tax is applied to total high-skilled labor costs of the leader and it is most effective when high-skill concentration is quite high. To assess the magnitude of this tax increase, we can conduct a back-of-the-envelope calculation using US data on tax revenues and public R&D subsidies.[56] In 2019, the US government spent around $175.5 billion on R&D tax incentives and government-financed innovation. Assuming that the top 25% of the income distribution is representative of high-skilled workers, an extra 1% increase in income tax amounts to around $82 billion, or almost half of all federal spending in R&D. The equivalent share calculated for Brazil would likely be higher as the Brazilian government spends less in R&D relative to GDP (0.82 for the US in 2019 vs. 0.4 in Brazil). Hence, the

[56] Along with being an easier reference, the US tax data is more easily available. Data comes from the Tax Foundation and the OECD for public R&D spending.

increase in R&D support would be substantial, though results show that a labor subsidy to innovation at lagging firms that takes into account skill concentration can be quite effective in boosting growth.

Importantly, this analysis points towards a different direction regarding education policy in places where skill concentration is high. An ever-increasing high-skilled labor supply, in itself, is not a recipe for higher growth rates once the skill concentration channel dominates the positive effect of human capital on growth. What is key to this conclusion is understanding the interaction between the high-skilled labor market and how innovative firms compete in the R&D space. As such, calls for a higher supply of skills should be understood within the context of high-skilled labor concentration at large firms. Along with boosting skill supply, government should also focus on competition policy.[57]

1.6 Conclusion

This paper shows how the effect of human capital on economic growth depends on high-skill concentration at large firms. I start by showing causal evidence in a difference-in-differences research design that increases in local skill supply from college creation had a negative and significant effect on GDP growth in municipalities where skill concentration was high. Results are robust to different specifications, changes to the sample, and show a relative decline of around 10% in local GDP growth between places with high and low skill concentration in the long term.

I then proceed to establish the role of skill concentration in the link between human capital and growth. First, I leverage the same difference-in-differences design to show that the increase in skill supply led to an increase in skill concentration of around 12%. Second, I build an SSIV using data on public loans to firms to show that skill concentration is non-monotonically related to GDP growth. While increasing skill concentration from a low level boosts economic growth, if skill concentration keeps increasing the relationship inverts and growth starts to decline. I further show causal evidence using the same SSIV that local skill concentration also has a non-monotonic relationship with local skill premium. My identification strategy passes the recommended tests in the SSIV literature and estimates are robust to several changes in the specification.

I then rationalize results in a model with step-by-step innovation and high-skilled labor demand and search. When firms are close in the technology ladder, competition is intense which raises the growth rate. Once a leader is significantly far ahead, it reduces its innovation effort as the threat of competition is lower

[57] Figure A31 shows that a policy that targets innovation catch-up through changes in the parameter h_c can also deliver higher growth. This can be achieved, for instance, through changes in the patent system.

and the likelihood of a lagging firm catching up is low. All the while, I show that high-skill concentration at the leader is monotonically increasing in the R&D gap. Thus, the model is able to reproduce the non-monotonic relationships observed in the data between skill concentration and growth.

With the model in hand, I analyze the effect of an increase in skill supply on growth. I show that this effect can be decomposed into two parts: one positive, due to the boost to R&D effort from lower high-skill hiring costs, and one negative, due to the increase in skill concentration across firms. I further show that the negative effect more than offsets the positive one when the level of skill concentration is high enough, leading to a decline in growth. Results also match the causal evidence on the relative decline in growth in highly concentrated municipalities from an increase in skill supply. The model also captures the decline in aggregate skill premium and the rise of high-skill unemployment in Brazil.

I then assess the role of a social planner in boosting growth in places where skill concentration is high after increasing high-skilled labor supply. I show that once the planner helps the lagging firm catch up through a subsidy to high-skilled labor hiring, they can effectively counteract the high-skill concentration channel and increase growth. This is relevant as it highlights the important role that firm dynamics and interaction should play in education policy as increasing skill supply when skill concentration is high can effectively backfire as the policy ends up helping large firms grow even larger. As such, both education and competition policies should go hand-in-hand.

Results, then, show that raising high-skill supply increases skill concentration at large firms and can lead to lower growth. Moreover, results are able to explain several of the observed empirical regularities in Brazil. By focusing on high-skilled labor concentration, I am able to explain the puzzling observation that a three-fold increase in high-skilled labor supply did not produce an increase in growth trends in Brazil between the late 1990's and the 2010's. My model also proposes a micro-foundation to the low business dynamism observed in Akcigit & Ates (2023) for the US. As firms away from the technology frontier require high-skilled labor to catch up, an increase in labor market power at the leading firm could make it harder for a lagging firm to adopt frontier innovation. Crucial to this point is seeing high-skilled labor flows as a channel for knowledge diffusion between firms. This is related to the use of non-compete clauses in the US where a firm can block knowledge flows by blocking former employees from being hired by competitors.

While not in the scope of this paper, I leave two ideas for future work. First, this framework can be easily expanded to take into account inter-sector labor market competition. A sector leader who experiences an increase in its labor search productivity can reduce R&D effort in other sectors competing for similar workers. It would be interesting to understand the role of skill concentration and hiring competition in explaining structural shifts in the economy, for example from manufacturing to services. Moreover, the model can also be applied in the context of competition between a domestic ("laggard") and a foreign ("leader") firm. Through "brain-drain" where domestic high-skilled workers go to work at market leaders abroad, domestic firms may find themselves unable to keep up with the technological frontier. The same rationale can be applied within a country between two regions where high-skilled labor migrates from one region to another.

References

Abramovsky, L., Harrison, R., & Simpson, H. (2007). University research and the location of business r&d. *The Economic Journal*, *117*(519), C114-C141.

Acemoglu, D., & Akcigit, U. (2012). Intellectual property rights policy, competition and innovation. *Journal of the European Economic Association*, *10*(1), 1–42.

Adão, R., Kolesár, M., & Morales, E. (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, *134*(4), 1949-2010.

Aghion, P., Harris, C., Howitt, P., & Vickers, J. (2001). Competition, imitation and growth with step-by-step innovation. *The Review of Economic Studies*, *68*(3), 467-492.

Aghion, P., & Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, *60*(2), 323–351.

Akcigit, U., & Ates, S. T. (2023). What happened to us business dynamism? *Journal of Political Economy*, *131*(8), 2059-2124.

Akcigit, U., & Goldschlag, N. (2023). *Where have all the "creative talents" gone? employment dynamics of us inventors* (Working Paper No. 31085). National Bureau of Economic Research.

Arrow, K. J. (1962). The economic implications of learning by doing. *The Review of Economic Studies*, *29*(3), 155–173.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., & Van Reenen, J. (2020). The Fall of the Labor Share and the Rise of Superstar Firms. *The Quarterly Journal of Economics*, *135*(2), 645-709.

Azar, J., Marinescu, I., & Steinbaum, M. (2022). Labor market concentration. *Journal of Human Resources*, *57*(S), S167-S199.

Azoulay, P., Graff Zivin, J. S., Li, D., & Sampat, B. N. (2019). Public r&d investments and private-sector patenting: evidence from nih funding rules. *The Review of economic studies*, *86*(1), 117–152.

Baxter, M., & King, R. G. (1999). Measuring business cycles: Approximate band-pass filters for economic time series. *The Review of Economics and Statistics*, *81*(4), 575–593.

Becker, G. S. (1962). Investment in human capital: A theoretical analysis. *Journal of Political Economy*, *70*(5), 9–49.

Benhabib, J., & Spiegel, M. M. (2005). Human Capital and Technology Diffusion. In P. Aghion & S. Durlauf (Eds.), *Handbook of Economic Growth* (Vol. 1, p. 935-966). Elsevier.

Bils, M., & Klenow, P. J. (2000). Does schooling cause growth? *American economic review*, *90*(5), 1160–1183.

Black, S. E., & Lynch, L. M. (1996). Human-capital investments and productivity. *The American Economic Review*, *86*(2), 263–267.

Borusyak, K., Hull, P., & Jaravel, X. (2021). Quasi-experimental shift-share research designs. *The Review of Economic Studies*, *89*(1), 181-213.

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200-230.

Christiano, L. J., & Fitzgerald, T. J. (2003). The band pass filter. *International Economic Review*, *44*(2), 435-465.

Cohen, W. M., & Levinthal, D. A. (1989). Innovation and learning: The two faces of r&d. *The Economic Journal*, *99*(397), 569–596.

Cordeiro, F., & Cox, A. (2023). College quality and tuition subsidies in equilibrium. *Working Paper*.

Cox, A. (2024). From classroom to prosperity: Fostering development through higher education. *Working Paper*.

De Loecker, J., Eeckhout, J., & Mongey, S. (2021). *Quantifying market power and business dynamism in the macroeconomy* (Tech. Rep.). National Bureau of Economic Research.

De Ridder, M. (2024). Market power and innovation in the intangible economy. *American Economic Review*, *114*(1), 199–251.

Diamond, P. A. (1982). Wage determination and efficiency in search equilibrium. *The Review of Economic Studies*, *49*(2), 217–227.

Dix-Carneiro, R., Goldberg, P. K., Meghir, C., & Ulyssea, G. (2021). *Trade and informality in the presence of labor market frictions and regulations* (Working Paper No. 28391). National Bureau of Economic Research.

Dix-Carneiro, R., & Kovak, B. K. (2015). Trade liberalization and the skill premium: A local labor markets approach. *American Economic Review*, *105*(5), 551-57.

Döttling, R., Gutierrez Gallardo, G., & Philippon, T. (2017). Is there an investment gap in advanced economies? if so, why? *ECB Forum on Central Banking*.

Durlauf, S. N., Johnson, P. A., & Temple, J. R. (2005). Chapter 8 growth econometrics. In P. Aghion & S. N. Durlauf (Eds.), *Handbook of economic growth* (Vol. 1, p. 555-677). Elsevier.

Feenstra, R. C., Inklaar, R., & Timmer, M. P. (2015). The next generation of the penn world table. *American Economic Review*, *105*(10), 3150-82.

Felix, M. (2022). Trade, labor market concentration, and wages. *Job Market Paper*.

Growiec, J., McAdam, P., & Muck, J. (2023). *R&D Capital and the Idea Production Function* (Research Working Paper No. RWP 23-05). Federal Reserve Bank of Kansas City.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, *50*(4), 1029–1054.

Harrigan, J., Reshef, A., & Toubal, F. (2023). *Techies and firm level productivity* (Working Paper No. 31341). National Bureau of Economic Research.

Hausman, J. A., Newey, W. K., Woutersen, T., Chao, J. C., & Swanson, N. R. (2012). Instrumental variable estimation with heteroskedasticity and many instruments. *Quantitative Economics*, *3*(2), 211-255.

Hausman, N. (2022). University Innovation and Local Economic Growth. *The Review of Economics and Statistics*, *104*(4), 718-735.

Lameiras, M. A. P., & Vasconcelos, L. S. (2018). *A evolução da população ocupada com nível superior no mercado de trabalho* (Technical Report No. 41). IPEA.

Liu, E., Mian, A., & Sufi, A. (2022). Low interest rates, market power, and productivity growth. *Econometrica*, *90*(1), 193-221.

Lucas, R. E. (1988). On the mechanics of economic development. *Journal of Monetary Economics*, *22*(1), 3-42.

Manera, A. (2022). Competing for inventors: Market concentration and the misallocation of innovative talent. *Available at SSRN 4212815*.

Mankiw, N., Romer, D., & Weil, D. (1992). A contribution to the empirics of economic growth. *Quarterly Journal of Economics*, *107*(May), 407-437.

Moody, C. E. (1974). The measurement of capital services by electrical energy. *Oxford Bulletin of Economics and Statistics*, *36*(1), 45–52.

Mortensen, D. (1982). The matching process as a noncooperative bargaining game. In *The economics of information and uncertainty* (p. 233-258). National Bureau of Economic Research, Inc.

Nelson, R. R., & Phelps, E. S. (1966). Investment in humans, technological diffusion, and economic growth. *The American Economic Review*, *56*(1/2), 69–75.

Nimier-David, E. (2023). Local human capital and firm creation: Evidence from the massification of higher education in france. *Working Paper*.

Olea, J. L. M., & Pflueger, C. (2013). A robust test for weak instruments. *Journal of Business & Economic Statistics*, *31*(3), 358–369.

Olmstead-Rumsey, J. (2022). *Market Concentration and the Productivity Slowdown* (Manuscript).

Pissarides, C. A. (1985). Short-run equilibrium dynamics of unemployment, vacancies, and real wages. *The American Economic Review*, *75*(4), 676–690.

Pritchett, L. (2006). Chapter 11 does learning to add up add up? the returns to schooling in aggregate data. In E. Hanushek & F. Welch (Eds.), *Handbook of the economics of education* (Vol. 1, p. 635-695). Elsevier.

Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, *98*(5), S71–S102.

Roseth, V. V., Valerio, A., Gutiérrez, M., et al. (2016). Education, skills and labor market outcomes. *World Bank Group. Retrieved February*, *15*, 2017.

Schubert, G., Stansbury, A., & Taska, B. (2024). Employer concentration and outside options [Working Paper]. *Available at SSRN 3599454*.

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, *225*(2), 175-199.

Taylor, J. (1967). A surrogate for regional estimates of capital stock. *Bulletin of the Oxford University Institute of Economics & Statistics*, *29*(3), 289-299.

Toivanen, O., & Vaananen, L. (2016). Education and invention. *The Review of Economics and Statistics*, *98*(2), 382-396.

Ulyssea, G. (2010). Regulation of entry, labor market institutions and the informal sector. *Journal of Development Economics*, *91*(1), 87-99.

UNDP, IPEA, & FJP. (2024). *Atlas do desenvolvimento humano no brasil.* http://www.atlasbrasil.org.br/. (Accessed June 12, 2024.)

Valero, A., & Van Reenen, J. (2019). The economic impact of universities: Evidence from across the globe. *Economics of Education Review*, *68*, 53-67.

Vick, B. (2017). Measuring links between labor monopsony and the gender pay gap in brazil. *IZA Journal of Development and Migration*, *7*(1), 10.

Chapter 2

Heterogeneous Local Fiscal Multipliers: New Shift-Share Evidence From The UK

2.1 Introduction

Fiscal policy discussions usually revolve around precise estimates of the fiscal multiplier, particularly whether it is above or below one. Since the Great Recession, however, the literature on national multipliers has done significant work analyzing the variability of estimates and their state dependence. The literature on local fiscal multipliers, on the other hand, has paid less attention to multiplier heterogeneity at the local level. This is possibly because local estimation procedures do away with aggregate-level channels (e.g. monetary policy), whose effect on the fiscal multiplier is more well known. Identifying local multiplier variation, however, is important since local spending decisions, for instance whether a council should go through a fiscal expansion or contraction, can change significantly between councils if multipliers differ. It is, then, worthwhile to assess whether we have local heterogeneity in the fiscal multiplier and identify the drivers of such local variation.

To address this question, I take advantage of the unique fiscal setting in the UK. Local government (called councils or local authority districts, or LADs) service and capital spending in the UK is a significant share of total public spending (approximately 13% for the fiscal year ending in 2020 according to the Office for National Statistics, or ONS). Importantly, fiscal transfers are only a small part of total LAD spending (approximately 0.2% of local net current expenditure) so virtually all local spending is in the provision of a public good. LADs have different sources of funding, of which central government grants are a significant share. Changes to these grants affect local authorities differently as their reliance on different sources of revenue varies from one another. I show the distribution of per-capita grant awards and the grant share of local GDP in Figure 2.1 for the LAD-year pairs between 2009 and 2019. Clearly there is significant variation in grant award and reliance between local governments.

I leverage this variation in the reliance on central government grants at the

Figure 2.1: *LAD Distribution of Per-Capita Grant Awards and Grant Share of Local GDP*

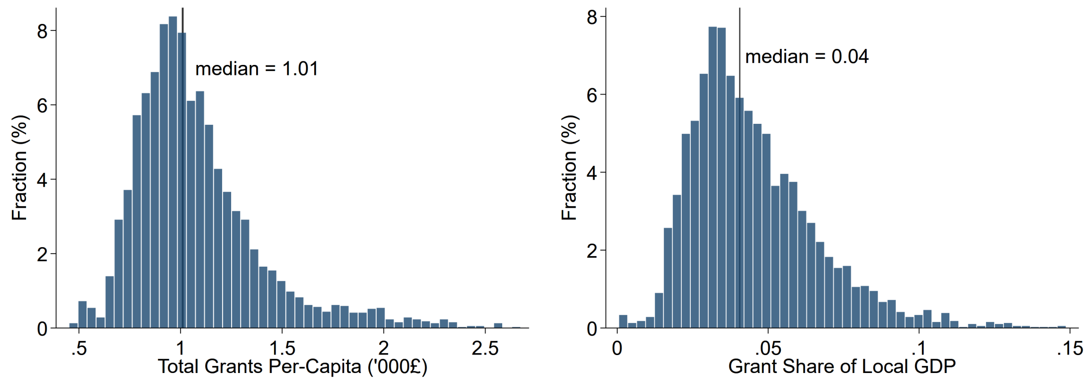

Note: Plot shows the distribution of local councils by total per-capita grants and by the grant share of local GDP. Data is from the ONS.

council level to identify the effect of local government spending on local GDP. Using a shift-share design, I find evidence of a positive average, short-term local fiscal multiplier of 1.69 for local government services and a multiplier of 1.71 for capital spending. Results are robust to controlling for local-level parameters, changes to the specification and the SSIV structure, and to running the estimation at the commuting zone level. I find no evidence of a statistically significant employment multiplier even though local authorities can use the additional funding to hire labor directly. To identify the fiscal multiplier estimate, I rely on the exogeneity of the one-year lagged central government grant share of local GDP to deal with the standard endogeneity issue in estimating fiscal multipliers. Although we cannot verify this exogeneity directly, I show evidence of its validity via recommended falsification tests.

I then characterize the heterogeneity in the fiscal multiplier with respect to local labor market and demographic parameters. Results can be divided into two groups. In the first one, higher economic inactivity and a higher low-skill labor share decrease the service and capital spending fiscal multipliers, respectively. In the second one, having more people in inactivity who want a job and higher anxiety levels in the population increase the multiplier. This variation due to local heterogeneity is significant as local estimates can vary between 0.6 and 3 if local observables change by one standard deviation. I further present evidence that this variation cannot be explained by local heterogeneity in marginal propensities to consume (MPCs), especially since I do not observe multiplier variation due to variables closely related to MPCs such as child poverty and inequality. Moreover, I find that local spending is able to boost worker productivity while improving local social and health conditions. On heterogeneity due to type of spending, I show that the average fiscal multiplier is mainly driven by spending in social care.

These results indicate two things. First, that local-level characteristics are important factors in determining the effectiveness of fiscal policy in increasing GDP. Second, that social-care spending plays an important role in the mechanism behind the above-one multiplier, as we also see improvement in local socio-economic conditions. Optimal policymaking should, then, take into account local conditions in determining which councils should go through a fiscal expansion or contraction.

To be able to assess fiscal policy, I construct a model using one possible mechanism that can explain both the local multiplier heterogeneity and the role of social care. Taking a cue from the economics and psychology literatures, I assume high and low-skilled workers are subject to a low mental bandwidth shock that lowers their cognitive aptitude and, hence, their productivity. This shock represents all ways through which a person's cognitive load may be overtaxed, for example through poverty and sickness, which dampens their ability to be at their best productivity level. Workers can return to their normal productivity level once hit by a high-bandwidth shock which depends on public spending. Moreover, workers can be of high-risk type which captures less revertible shocks such as suddenly becoming a carer for a partner with a long-term condition. I assume public spending is less effective for high-risk individuals in bringing them back to their healthy, more productive state. Applying this model to places with different low-bandwidth shock rates, shares of high-skilled labor, and shares of high-risk individuals I am able to reproduce the observed heterogeneous trends in the fiscal multiplier. Results are intuitive: public spending is most effective when helping high-skilled (i.e. most productive) workers who are not high-risk individuals return to their healthy state. To the best of my knowledge, this is the first paper that links the effect of individual-level cognitive bandwidth capacity to the effectiveness of fiscal policy.

I then analyze potential gains from optimal fiscal policy. I split the analysis into two parts. In the first one, I show that by taking local heterogeneity in the multiplier into account, we can derive gains from current local fiscal policy in the UK even if we keep total spending constant. This is due to *fiscal misallocation*, i.e. the fact that high-multiplier councils do not receive the largest central government grant awards. While this misallocation declined between 2010 and 2019 in the UK, cumulative results are significant. If the central government had optimally awarded grants between 2010 and 2019, i.e. more grants to councils with high multipliers, it would have generated an extra 57.9 billion pounds in real terms. Conversely, if we assume the national government could control local budgets altogether (vs. only the grant allocation) then removing fiscal misallocation entirely since 2010 would have resulted in an extra 156 billion pounds, or 19.2% of the central government budget in 2019. These results highlight the importance of taking heterogeneity into account in fiscal policy.

In the second part, I present the social planner's optimal fiscal policy and compare it with actual spending by UK councils. Results have two main takeaways. First, there is considerable optimal spending heterogeneity between councils as the interquartile range for the ratio between optimal and actual spending is 0.84. Second, a social planner would want, on average, to increase local budgets by 28%, a result that reflects the 1.69 baseline fiscal multiplier estimate. I show then that the potential gains from adopting the planner's fiscal policy have increased over time, averaging 0.5 percentage points of per-capita GDP yearly and 0.65 percentage in 2019. The latter would raise the 2019 GDP per capita growth rate in the UK by more than 50%.

Related Literature

This paper relates to different strands of literature. On the state dependency of local multipliers, Basso & Rachedi (2021) find evidence in the US that a higher share of young people in the population increases the fiscal multiplier of military spending. They rationalize results in a New Keynesian framework with credit market frictions. Similarly, Morita (2022) is a recent attempt[58] at explaining low fiscal multipliers in Japan with the aging of its population, now from a VAR-narrative perspective. Finally, Brandao-Roll et al. (2024) show local heterogeneity of multipliers from Pell Grants with respect to recessions and local type of college. I make two important additions to these results. First, I show further evidence of labor-market and demographic-driven local multiplier heterogeneity. Second, I propose a simple model that rationalizes results on local heterogeneity which does not require a New Keynesian framework to explain variation in the local fiscal multiplier.[59]

On the literature on local fiscal multiplier estimates, Chodorow-Reich et al. (2012), Fishback & Kachanovskaya (2015), and A. Auerbach et al. (2022)[60] provide different estimates for the local multiplier in the US, particularly during downturns. These papers, however, do not analyze the effects of state dependency at the local level which is my focus. There is evidence, on the other hand, that local fiscal multiplier estimates are larger during recessions relative to expansions (Nakamura & Steinsson, 2014, Shoag, 2016, Berge et al., 2021). I expand this result on two fronts. First, I show that local public spending is also state dependent relative to local-area characteristics. I then propose a mechanism to rationalize my empirical findings that makes use of individual-level cognitive load capacity. It is important to notice that a large part of the literature on multipliers deals with

[58] Building on Yoshino & Miyamoto (2017).

[59] C.f. Muratori et al. (2023) for evidence on multiplier heterogeneity due to differences in government purchases, and Gibbons et al. (2019) and Gibbons & Wu (2019) for analyses on the differential impact of road and airport infrastructure investments, respectively.

[60] See Chodorow-Reich (2019) for an interesting review of local fiscal multipliers.

the effect of direct transfers and/or purchases by governments. In my case, local authority spending in the UK is virtually all about the provision of services and capital. As such, the patterns of heterogeneity that I highlight along with the mechanism I propose do not involve heterogeneity due to variation in local MPCs, which is usually the explanation behind business cycle variation in multipliers from fiscal transfers.

On the aggregate level, there is significant evidence of fiscal multiplier heterogeneity. Different authors have shown state dependence of multipliers regarding interest rates (Christiano et al., 2011, Ramey & Zubairy, 2018), expansions and contractions (A. J. Auerbach & Gorodnichenko, 2012), exchange rate regime, debt level, and trade openness (Ilzetzki et al., 2013, Corsetti et al., 2012). There is also evidence that fiscal multipliers from changes in spending and taxes are different (Caldara & Kamps, 2017), along with variation due to methodological choice (Gechert, 2015, Capek & Crespo Cuaresma, 2020). Finally, there is work on heterogeneity with respect to spending type (Pappa, 2009, Boehm, 2016).[61] While the aggregate multiplier literature paints a rich picture vis-a-vis the variation in estimates, less is known at the local level. I intend to show that we also observe significant heterogeneity in local-level estimates which is not due to the usual channels analyzed at the aggregate level. Local-level heterogeneity has to rely on a micro-level mechanism, of which I show evidence.

Finally, I take inspiration from the literature at the intersection of psychology and economics on the psychological toll to one's mental bandwidth. Schilbach et al. (2016) is an interesting summary of how poverty, by imposing a cognitive load, can tax a person's bandwidth resulting in lower productivity and changes to rational behavior. Kaur et al. (2021) show evidence from a field experiment that increasing cash-on-hand raises the productivity of poor workers. Similarly, Schultz & Edington (2007) review results showing the toll of poor health on worker performance.[62] I expand these results by linking the low-bandwidth toll to fiscal policy as a mechanism that creates heterogeneity in the local fiscal multiplier depending on the local conditions of individuals. I also show direct evidence that fiscal policy both boosts worker productivity and improves local social and health conditions, results that can be naturally linked via the low-bandwidth mechanism.

The remainder of the paper is organized as follows. Section 2.2 describes the data. Section 2.3 describes the shift-share design that I use to calculate estimates. Section 3.3 presents my empirical results. Section 2.5 rationalizes results via a theoretical model and compares actual with optimal spending. Finally, Section

[61] C.f. Ramey (2019) for a more extensive review.

[62] See W. Burton et al. (2001) for evidence on allergies, W. N. Burton et al. (2003) on the positive effect of drugs on the productivity of sick workers, and Goetzel et al. (2004) on the costs of presenteeism.

3.5 concludes.

2.2 Data

I rely on several council-level sources to pin down the effect of local characteristics on local fiscal multipliers. The main public budget data comes from the ONS which has local public accounts information for England since the fiscal year of 2007-2008. These accounts hold information about service and capital expenditure on education, transportation, social care, healthcare, housing, cultural and environmental activities, law enforcement, planning and development, and general expenses. As the data covers the UK fiscal year which goes from April to March of the following year, I adjust all variables to match the chronological year.

Local authority data are also adjusted to account for differences in service provision between counties and districts. LADs can be categorized as metropolitan districts, London boroughs, unitary authorities, districts, and councils. For our analysis, it is important to note that non-metropolitan districts are part of a larger county.[63] Both of these local entities split the scope of local service that they provide: for instance, while non-metropolitan districts run environmental services, the encompassing council is responsible for social care. As councils receive central government grants for services provided in all of their districts, this creates a problem not only for the identification of the fiscal multiplier but also of cross-correlation between observations. To deal with the former, I split a council's spending and grants between its districts according to their population shares within the council. I then exclude LADs that are councils from the sample. As for the cross-correlation between non-metropolitan districts within the same council, I cluster standard errors by their parent council if they are districts.

Although councils help administer the majority of the welfare-related direct transfers in the UK, they are not directly responsible for such programs. Those transfers are classified as either mandatory or discretionary. The former, such as Housing Benefit (now called Universal Credit), is set and paid by the central government, while the latter is funded by local authorities (with some help from the central government) usually as an additional benefit in case a household requires further assistance. Discretionary transfers are only a small part of total LAD spending (approximately 0.2% of local net current expenditure). Hence, most of councils' public spending is in the form of government services. I also include data from the local authority capital accounts which relate to investments in fixed assets.

Local authority accounts also show the sources of funding through local taxes and central government grants. LADs in the UK can generate revenue via Coun-

[63]For example, the district of Cambridge is within the council of Cambridgeshire.

cil Tax, which is a property tax levied on residential properties, Business Rates, a property tax on businesses, central government grants, and local fees and fines. Council tax rates are set up by LADs, but from 2012 to 2018 they could not be raised more than 2% for most councils without a public referendum. Business rates are set up by the central government but since 2013 local authorities get to retain 50% of what they collect locally, while the other half is redistributed back to councils as a grant. Prior to 2013, the central government decided the redistribution of 100% of the business rate income. Finally, central government grants are funded by the national government and can be of two types: general grants, which can be used freely by the LADs though could be for a specific spending category such as education, and earmarked grants, where the LAD only acts as a "middle-man" by transferring the grant funds either to people or to a third-party who runs a specific service. The most relevant general grants are called Specific Grants Within the AEF (Aggregate External Finance), which is an umbrella for several smaller grants, Formula Grants, and Revenue Support Grants. Earmarked grants are called Specific Grants Outside the AEF. For the shift-share approach, we will use data on non-earmarked grants as earmarked ones are mainly for mandatory rent rebates (i.e. transfers) and, hence, do not relate to service spending.

Aside from local government spending, I use several local-level controls and auxiliary variables. Council-level demographic data comes from the ONS and the National Archive. Labor market data is from the Annual Population Survey. Local political control data comes from the Open Council Data UK. Finally, aggregate disease levels were measured using the DALY (disability-adjusted life year) available from the IHME (Institute of Health Metrics and Evaluation). This is a measure of aggregate disease burden defined as the sum of the number of years lived with a disability and the number of years lost due to early death calculated using life expectancy. To focus on disease factors that are more closely related to local public health and social care, I restrict the DALY to health changes due to risk factors which include environmental and occupational risks, behavioral risks (e.g. malnutrition), and metabolic risks (e.g. high cholesterol). Although not without its flaws, the DALY is an important measure in the public health literature which allows policymakers to compare, on the aggregate level, different disease risks by their impact on the population. I show in Table B.1 in the Appendix the summary statistics for the main variables of interest.

2.3 Research Design: The Shift-Share Approach

The identification strategy exploits the heterogeneous reliance of local councils on central government grants to pin down the local fiscal multiplier and its heterogeneity due to variation in local characteristics. The fiscal policy setting in the UK creates a framework where central government transfers to local governments affect each council differently depending on their reliance on these funds. Hence, I propose a shift-share IV approach to deal with the usual endogeneity issue in estimating the impact of public spending on local GDP.

A Bartik-style instrument exploits how an aggregate shock affects local areas differently through variation in local shares. In this case, I rely on how changes to central government grants at the national level affect local councils differently given their heterogeneous exposure to grants measured via the council-level grant-to-GDP share. I follow the "shares-approach" framework developed by Goldsmith-Pinkham et al. (2020) where the identification strategy relies on the exogeneity of the lagged grant shares conditional on observables. As such, consistency lies on exogenous exposure to common shocks. The main assumption behind a shares-based approach is that past exposure to a policy (i.e. the grant share of local GDP) is conditionally exogenous to growth in local GDP.

Formally, for council l at time t I estimate the local fiscal multiplier as follows:

$$\frac{Y_{l,t+1} - Y_{lt}}{Y_{l,t-1}} = \beta \frac{G_{lt} - G_{l,t-1}}{Y_{l,t-1}} + \gamma X_{l,t-1} + \phi_l + \psi_t + \epsilon_{lt} \tag{2.1}$$

where Y_{lt} is real GDP level per capita at the council level, G_{lt} is the local government real net spending per capita, $X_{l,t-1}$ are lagged controls, ϕ_l are council fixed-effects, ψ_t are year fixed-effects, and ϵ_{lt} is the residual. I opt for local GDP growth one period ahead to avoid issues with fiscal year reporting since local authority accounts are reported for periods between April and March of the following year. However, for β to have a direct fiscal multiplier interpretation, I scale both the dependent variable and the main regressor by the same variable $Y_{l,t-1}$.

Given the counter-cyclical nature of local spending, the OLS estimate of β is biased. To address this endogeneity issue, I instrument local government spending growth with the shift-share IV B_{lt} defined as:

$$B_{lt} = \sum_k g_{kt} s_{lk,t-1} \tag{2.2}$$

where $s_{lk,t-1}$ is the share of funding source k in the l council's GDP at time $t-1$ and g_{kt} is the national growth rate of funding k at time t. To be clear, if the central government offers funding to councils through two types of grants ($k = 1, 2$), the share $s_{1,10,0}$ is the share of grant 1 in council 10's local GDP at time 0 and $s_{2,10,0}$ is the share of grant 2. Identification comes from the exogeneity of $s_{lk,t-1}$ with respect to changes in local GDP.

Formally, the "shares-approach" for identification in a shift-share estimation requires both relevance and validity conditions to hold. Following Goldsmith-Pinkham et al. (2020), for T time periods, K grants, and L councils, the difference between the 2SLS estimator and the parameter of interest is:

$$\hat{\beta} - \beta = \frac{\sum_{t=1}^{T} \sum_{k=1}^{K} g_{kt} \sum_{l=1}^{L} s_{lk,t-1} \epsilon_{lt}^{\perp}}{\sum_{t=1}^{T} \sum_{k=1}^{K} g_{kt} \sum_{l=1}^{L} s_{lk,t-1} \Delta G_{lt}^{\perp}} \tag{2.3}$$

where ΔG_{lt} is the change in local fiscal spending scaled by lagged local GDP (as shown in Equation 2.1) and the $\perp$ superscript indicates the corresponding residualized variable after controlling for $X_{l,t-1}$ and the fixed-effects.

The relevance condition requires that the denominator in Equation 2.3 must converge to a non-zero value, i.e. that the grant shares hold predictive power over the local spending growth ΔG_{lt} conditional on controls and that the aggregate growth rates g_{kt} do not weight the covariates in a way that the sum cancels out. This is easily verified by regressing local spending growth on the instrument.

As for the validity condition, we require that the numerator in Equation 2.3 must converge to zero. This happens when the grant shares are uncorrelated with the error term conditional on the controls, i.e. when $\mathbb{E}[\epsilon_{lt} s_{lk,t-1} | X_{l,t-1}, \phi_l, \psi_t] = 0$ for all k where $g_{kt} \neq 0$. The assumption of grant share exogeneity rests on the idea that councils that rely on the central government for funds with different intensities are differently exposed to policy shocks affecting local grants. This is akin to a difference-in-differences counterfactual: in the absence of aggregate shocks to central government grants, high and low-dependent councils would have behaved similarly in terms of local GDP growth. As highlighted in Borusyak et al. (2021), the share exogeneity assumption is also appropriate when we use tailored exposure shares in the SSIV, which is the case in this framework.

As usual, the validity condition cannot be verified directly though I run recommended falsification tests. The IV validity may not hold if local authorities with different grant shares have other characteristics that can explain trends in local GDP growth other than through local spending. While this cannot be tested directly, I run falsification tests, as recommended by Goldsmith-Pinkham et al. (2020), in Section 2.4.2 that partially assess the plausibility of the assumption. First, I perform a balance check by separately regressing the grant shares and local GDP growth on local-level observables. The test consists of checking whether each observable correlates significantly and simultaneously with both the shares and the dependent variable. Even though a significant coefficient at this stage in both regressions is not a problem per se given that the validity assumption is conditional on controls it could point towards an omitted variable problem. In a second test, I instrument the baseline specification with the lagged grant shares interacted with year fixed-effects as separate instruments in a many-IV 2SLS setting, i.e. instrumenting with $s_{lk,t-1}$ interacted with time fixed-effects instead of

B_{lt}. This procedure is based on the fact shown by Goldsmith-Pinkham et al. (2020) that the Bartik estimator is equivalent to a GMM estimation that uses lagged local shares as instruments and a specific weight matrix whose components are called "Rotemberg weights." As such, this shares-directly 2SLS produces an unweighted estimate that should be similar to the baseline SSIV estimate under homogeneous effects. Fiscal multipliers, however, are known to vary with the business cycle implying that the test will fail. Nonetheless, it is still useful to analyze the Rotemberg weights to check whether the heterogeneity pattern makes sense.

2.4 Results

In this section, I present evidence of how local heterogeneity affects the fiscal multiplier. First, I pin down an average estimate of the multiplier using the baseline specification for both revenue and capital spending. I then run robustness checks for the SSIV design and show evidence to support the identification strategy. Next, I assess fiscal multiplier heterogeneity by spending category. Finally, I identify a set of local parameters that can explain the variation in the local fiscal multiplier between LADs and show direct evidence that the underlying mechanism is not variation in local MPCs.

2.4.1 Fiscal Multiplier

Before analyzing possible sources of heterogeneity, it is important to pin-down the average local fiscal multiplier to put the effects of heterogeneity into perspective. I proceed using the benchmark specification shown in Equation 2.1 where I use change in real per-capita local spending as my main regressor and local GDP per-capita growth as the dependent variable. Regarding the central government grants listed in Section 2.2 I only use non-earmarked grants, which consist of the majority of the central government funding, to construct the instrument shares. I also do not use grants that are linked to changes in local business taxation as those are likely not valid as instruments.[64] I then construct the shift-share IV in Equation 2.2 using the lagged within AEF specific grants share of local GDP.[65]

While I do have disaggregated data on the individual grants that comprise the within AEF grant bin, I choose to work with a single aggregated grant bin, i.e. I sum all grants within the AEF. This is due to significant noise at the grant level. Using individual grants separately to calculate the SSIV in Equation 2.2 is impractical given the frequent policy changes how grants are labeled. Grants are frequently created, renamed, split, merged, and ended depending on what the

[64]This excludes the Revenue Support Grant. I further exclude the Police Grant which is awarded directly to local police bodies. These are treated as a separate local entity in the local spending accounts and are excluded from my sample.

[65]In certain years, I also add temporary general grants called Local Services Support Grant and Area-Based Grant as those were created from a relabeling of previous specific grants.

central government is focusing on in a given year. Take, for instance, the "Early Intervention Grant" (EIG), an early education grant that was created in 2010 by bringing together several smaller grants. Its creation also included changes in how the grant was allocated and how much money the central government was willing to spend on it. In 2013 it was decided that the EIG grant would no longer be paid as a separate grant. It would instead show up under different grants such as the Dedicated Schools Grant (DSG). Since the underlying purpose of the funding remains the same (both EIG and DSG are "within" AEF grants), I can use the yearly aggregated sum to capture movements in government funding while avoiding changes that are essentially in form. This cleans the SSIV of much of the noise generated by single grants life cycle while effectively capturing the weight of the central government on local budgets.[66]

It is also important to highlight a few differences with the benchmark case in Goldsmith-Pinkham et al. (2020). First, I am dealing with an "incomplete shares" case where the local GDP share of all central government grants is not one. As shown in Borusyak et al. (2021), this would require controlling for the total sum of shares in each locality. However, since I am aggregating grants into a single bin we have $k = 1$ which implies that the total sum of shares is already factored in. In our case, however, the threat to identification comes from the earmarked grants which were not included in the SSIV calculation. If earmarked grant awards correlate with the non-earmarked grants used in the SSIV, our estimate will be biased. To deal with this issue, I show fiscal multiplier results where I also control for the local per-capita amount of grants outside the AEF. To further strengthen our identification, I add to the control set the other LAD funding sources, i.e. council tax, non-domestic rates charged from businesses (though set-up nationally), the per-capita amount of people receiving fiscal transfers, and the stock of reserves held by the local authority.

I now proceed with the estimation of the benchmark setting. I show in Table 2.1 the 2SLS results using next-period GDP growth and next-period employment change as my dependent variables.[67] Columns (1)-(4) report the next-period fiscal multiplier while columns (5)-(8) report the employment fiscal multiplier. Columns (1) and (5) have no local-level controls, columns (2) and (6) add several council-level controls, columns (3) and (7) control for the one-year lagged outside AEF per-capita grant amount, and columns (4) and (8) add local controls together with the lagged outside AEF grant amount. Controlling for grants not included

[66] I show estimation results using grant bins aggregated by large spending categories (e.g. education, social care) in Table B.2. Point-estimates are statistically indistinguishable from baseline ones.

[67] Figure B1 in the Appendix shows the binned scatter plots for the first-stage and the reduced form results of the next-period GDP growth specification controlling for year and council fixed-effects. It shows strong IV relevance and a positive correlation between the SSIV and local GDP growth, both of which do not seem to be driven by outliers.

in the SSIV allows us to compare councils that receive the same amount of central government grants not included in the SSIV while exploiting grant heterogeneity in non-earmarked grants, effectively accounting for total reliance on central government funding. Finally, standard errors are clustered at the county level for non-metropolitan districts and at the LAD-level for the other observations in all specifications, and the heteroskedastic-robust F-statistic for the instrument, which is reported at the bottom of the table, is well above the usual threshold level for the weak-IV test.

Table 2.1: Local Service Spending Fiscal Multiplier Estimates

	GDP_{t+1}				Emp_{t+1}			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Multiplier	1.738**	1.666**	1.769**	1.693**	0.140	1.725	0.0417	1.813
	(0.791)	(0.785)	(0.806)	(0.798)	(1.434)	(1.321)	(1.471)	(1.194)
N	3,235	3,235	3,235	3,235	3,235	3,235	3,235	3,235
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Council FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes		Yes		Yes		Yes
Outside AEF Grants			Yes	Yes			Yes	Yes
Robust F-statistic	100.8	108.9	101.0	109.8	100.8	102.8	101.0	108.9

Notes: Main regressor corresponds to growth in real local authority total service expenditure per-capita. GDP_{t+1} corresponds to local GDP per-capita growth one period ahead scaled by one-year lagged GDP per-capita ($(Y_{l,t+1} - Y_{lt})/Y_{l,t-1}$). Emp_{t+1} corresponds to local employment per-capita growth from $t-1$ to $t+1$ ($(L_{l,t+1} - L_{l,t-1})/L_{l,t-1}$). Local-level controls (one-year lagged): share of NVQ3 awards (equivalent to a high-school diploma), unemployment rate, median age, child poverty rate, average full-time wage, per-capita number of people receiving central government benefits (for Disability Living Allowance, Incapacity Benefits, Housing Benefit, Universal Credit, Personal Independence Payment, and Child Benefit) where each claimant counts by the number of individual benefits they receive, dummies for the political party controlling the local council, LAD reserves per-capita, per-capita amount of funds from non-domestic rates, and average council tax. Outside AEF Grants is the per-capita amount of all grants outside the AEF which are not included in the SSIV. Standard errors are clustered by counties for non-metropolitan districts and by individual LAD for the rest. *, **, *** denote significance at 10%, 5%, and 1% levels, respectively.

Starting with columns (1)-(4), we observe that point-estimates are larger than one and significantly different from zero. Adding the lagged earmarked grant funding as a control does not affect estimates significantly, nor does adding local-level controls and controlling for other sources of LAD funding (i.e. fiscal reserves, non-business rates, and council tax). My preferred fiscal multiplier estimate of 1.69 is close to the 1.9 median estimate in the literature on regional multipliers (Chodorow-Reich, 2019) although the literature is mainly about the multiplier of direct fiscal transfers. This is evidence that service-based fiscal multipliers can have a similar magnitude to those calculated from fiscal transfers. On a side note, although point-estimates are above 1 they are not significantly different from it. As the point of the analysis later on is to show multiplier heterogeneity, I will not focus on the average estimate.

I show results are robust to changes in the benchmark specification. I run different yet similar specifications to Equation 2.1 in Table B.2 in the Appendix. Results are robust to running a weighted specification weighting by the one-year lagged logarithm of LAD population (columns (1) and (2)), to using two-year lagged grant shares when constructing the SSIV, that is when using $s_{lk,t-2}$ in Equation 2.2(columns (3) and (4)), and to using fixed initial grant shares, i.e. $s_{lk,0}$, which follows the SSIV convention (columns (5) and (6)).[68] In the initial-shares case all controls are fixed at the same time period as the shares and interacted with year fixed-effects. As initial-period shares become less relevant in later periods, there is a significant drop in first-stage relevance when controls are added. Moreover, results remain unchanged if instead of using $(Y_{l,t+1} - Y_{lt})/Y_{l,t-1}$ as our dependent variable we use $(Y_{l,t+1} - Y_{lt})/Y_{lt}$ (columns (7) and (8)). In all specifications, point-estimates are statistically indistinguishable from benchmark ones.

As an additional step, I show results are robust to calculating the SSIV at a more disaggregated grant level. As aforementioned, I aggregate all non-earmarked central government grants that are not related to local business taxation funds when calculating the SSIV. We can, however, aggregate grants by broad spending category k (c.f. Equation 2.2) which requires matching individual grants to a category. I do so in columns (9) and (10) of Table B.2 in the Appendix for the following categories: education, social care, local development, housing, and healthcare (leftover grants are binned together as "other"). Point-estimates are higher though they remain statistically indistinguishable from baseline ones and significantly different from zero. As expected, we observe a drop in first-stage relevance. It is also worth noticing that, differently from the baseline case, here we have to control for the one-year lagged sum of non-earmarked grants. This is because the specification deals with "incomplete shares" where individual spending category shares do not add up to a constant (Borusyak et al., 2021). As such, I add the total share as a control in columns (9) and (10).

It is important to highlight at this point what the local multiplier results mean. Since I rely on local council variation to pin down the local multiplier, results imply that a 1% increase in an LAD's spending *relative* to other local councils increases *relative* local GDP by 1.69%. As such, results do not take into account spillover effects where spending in one LAD might affect a neighboring local authority, nor general equilibrium effects acting, for instance, via inter-council migration and commuting patterns. One way to assess these effects is to run the specification at a different geographical level. I do so at the level of "travel-to-work areas" (TTWA) which are commuting regions in the UK.[69] I show in Table

[68] Typically, shares are fixed in a period before a policy in which the instrument is based on comes into effect (i.e. a pre-period). Since my sample starts in the fiscal year of 2007-08, there is no pre-period. Nonetheless, the validity assumption works with any amount of time lag.

[69] While there are 398 councils in the UK, there are only 228 travel-to-work areas.

B.3 in the Appendix estimates for the fiscal multiplier at the TTWA level. Since LADs are not associated 1:1 with TTWAs, I link each local authority to the TTWA that contains the majority of its postcodes. Notwithstanding the lower number of observations and the fact that local grants are not awarded at the travel-to-work level, two facts that explain the drop in instrument relevance, estimates are close to the ones calculated at the council level. This provides evidence that results at the LAD-level are possibly relevant at larger geographical levels and are not significantly affected by commuting and migration patterns.

I further find no effect of local spending on employment. As shown in columns (5)-(8) of Table 2.1, the estimate for employment is not statistically significant.[70] The estimation uses Equation 2.1 except that now the dependent variable is the two-year change in employment (i.e. $(L_{l,t+1} - L_{l,t-1})/L_{l,t-1}$, where $L_{l,t}$ is employment per capita). As such, an increase in local spending does not seem to generate new jobs, despite local authorities being able to hire directly in the labor market.

We can also analyze the short-term multiplier effect of capital spending. Although results so far have been about public service provision, which corresponds to most local spending, local authorities also invest in fixed assets such as schools, vehicles, and intangibles.[71] As with services, councils receive central government grants which can be used for local capital investing. We can, then, apply a similar specification to Equation 2.1 to estimate the fiscal multiplier of public capital spending where I instrument local government spending with an SSIV calculated using central government capital grants.

Results for capital spending are shown in Table 2.2. In columns (1) and (2), I report the one-period-ahead fiscal multiplier for capital spending growth calculated over two years to take into account possible adjustment costs in capital investing. As with services, the capital multiplier estimate is above one and statistically significant, implying that a 1% increase in a council's capital investing relative to other local councils increases relative local GDP by 1.71%. This estimate is close to the services one of 1.69. We can, then, ask whether service and capital spending are confounding each other's fiscal multiplier estimates as spending patterns may be correlated. I assess this point in columns (3)-(6) where I regress local GDP growth on both service and capital spending, each instrumented with their respective central government grant SSIV.[72] Although instrument relevance is weaker given the more demanding specification, results show that there is little bias from regressing each type of spending separately as es-

[70]To get employment multipliers, we have to multiply the table coefficients by the employment-to-GDP ratio.

[71]Fixed asset investment corresponds to around 16% of total local public expenditure in sample.

[72]Grant shares for services are two-year lagged, i.e. $s_{lk,t-2}$, to match capital shares which are two-year lagged since the change in capital spending is over two years.

timates barely change, especially for capital spending whose fiscal multiplier is identified more precisely.[73] This is reassuring from the point of view of instrument validity.

Table 2.2: Local Capital and Services Spending Fiscal Multiplier Estimates

	GDP_{t+1}					
	(1)	(2)	(3)	(4)	(5)	(6)
Capital Spending$_{2y}$	1.574***	1.706***	1.502***	1.613***	1.616***	1.673***
	(0.540)	(0.561)	(0.552)	(0.571)	(0.548)	(0.564)
Service Spending			1.091	1.313	1.168	1.345
			(1.012)	(1.051)	(1.032)	(1.066)
N	2,938	2,938	2,938	2,938	2,938	2,938
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Council FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes		Yes		Yes
Outside AEF Grants					Yes	Yes
Robust F-statistic	69.4	63.8	35.7	32.8	36.0	33.7

Notes: Main regressors correspond to growth in real local authority capital and service expenditure per-capita. SSIV for service spending uses two-year lagged grant shares to match the timing of the capital grant shares. Subscript $2y$ indicates that the change is over two years. GDP_{t+1} corresponds to local GDP per-capita growth one period ahead. Local-level controls (two-year lagged): share of NVQ3 awards (equivalent to a high-school diploma), unemployment rate, median age, child poverty rate, average full-time wage, per-capita number of people receiving central government benefits (for Disability Living Allowance, Incapacity Benefits, Housing Benefit, Universal Credit, Personal Independence Payment, and Child Benefit) where each claimant counts by the number of individual benefits they receive, dummies for the political party controlling the local council, LAD reserves per-capita, per-capita amount of funds from non-domestic rates, and average council tax. Outside AEF Grants is the per-capita amount of all grants outside the AEF which are not included in the SSIV. Standard errors are clustered by counties for non-metropolitan districts and by individual LAD for the rest. *, **, *** denote significance at 10%, 5%, and 1% levels, respectively.

Results, then, show that public spending on services and capital have, on average, multipliers whose point-estimates are above one. It is not clear, however, whether such multipliers present heterogeneity between councils and, if so, whether it is related to usual explanations involving differences in individual-level MPCs. I analyze this point further in Sections 2.4.3 and 2.5 when discussing a mechanism behind the larger-than-one multiplier and its underlying local heterogeneity.

2.4.2 Robustness of the SSIV Design

The main concern in a "shares-based" shift-share framework is the violation of the conditional exogeneity assumption regarding local shares. For that to happen, unobservable correlates of the lagged grant shares need to have some explanatory power over the outcome variable, i.e. local-level real GDP growth. While this as-

[73] I report multipliers for total local spending, i.e. services and capital expenditure combined, in Table B.4 for both one-year and two-year changes in spending. Estimates are statistically significant in all specifications and indistinguishable from benchmark ones.

sumption cannot be verified directly it is useful to understand how shares are correlated with observables and how each observable might be correlated with the outcome variable. While I control for these observables in the main specification, it could still be the case that places with different values of observables have systematically different unobservables which would violate the validity condition of the shift-share instrument.

To assess this point, I run a balance test in Table B.5 in the Appendix by regressing the local-level observables on local grant shares and local GDP growth. I also show results for the largest grant inside the AEF, i.e. the Dedicated Schools Grant (DSG). The SSIV passes the balance test if there is no simultaneous significant coefficient in columns (1) and (2), or columns (3) and (4). All variables are demeaned and normalized to have unit variance so that coefficients are more easily interpretable. We observe that there are no balance issues with the DSG share as no covariate has a significant coefficient with both the DSG share and GDP growth. However, the local council being under the control of the Liberal Democrats (LD) poses a threat to identification with respect to the aggregated grant share as the coefficients are statistically significant in both columns (3) and (4). The magnitude of the correlation is relatively high as being controlled by the LD relative to the Conservatives is associated with a decrease in the grant share by 13% of its standard deviation and a decrease in growth by 23% of its standard deviation. However, only 4.5% of all LAD-year observations have a local council controlled by the Liberal Democrats. As such, the imbalance between high-share and low-share councils regarding political control by the LD does not seem to be substantial enough to affect our multiplier results in Table 2.1.[74]

Since the SSIV combines council-level grant shares in different years to construct a single instrument, it is also useful to analyze the contribution of each individual grant share (i.e. $s_{lk,t-1}$ for each t) to the estimation. As shown in Goldsmith-Pinkham et al. (2020), the shift-share IV estimation is equivalent to a GMM estimation using moment conditions on the shares and a specific weight matrix. These (Rotemberg) weights, which are calculated for each year t, tell us how sensitive our baseline estimate is to misspecification (i.e. endogeneity) in a particular year and they depend on the covariance between the first-stage fitted value calculated for each year using $s_{lk,t-1}$ as an instrument and the endogenous variable.[75] As such, we can run a 2SLS estimation using each yearly grant share separately as an instrument by interacting them with year fixed-effects to calculate unweighted estimates. Specifically, let T be the total number of years in the

[74]Removing LAD-year observations where the LD have control of the council results in a multiplier of 1.66 in the specification with all controls vs. 1.69 for the benchmark.

[75]In the simplest example where the individual, yearly instruments are all orthogonal to each other, the weights are simply the ratio between the just-identified first-stage R^2 and the full SSIV first-stage R^2.

sample. We then run the following first stage:

$$\frac{G_{lt} - G_{l,t-1}}{Y_{l,t-1}} = \sum_{t \in T} \delta_{kt} s_{lk,t-1} + \xi X_{l,t-1} + \phi_l + \psi_t + \eta_{lt} \tag{2.4}$$

where the second stage is the one shown in Equation 2.1.

The fiscal multiplier estimate calculated using this many-IV 2SLS should coincide with the ones in Table 2.1 in a setting with homogeneous treatment effects over time. However, since we know that fiscal multipliers calculated during expansions and recessions differ, we should expect the estimates calculated instrumenting with grant shares directly to be different from our benchmark ones.[76]

I show in Table B.6 in the Appendix the Limited Information Maximum Likelihood (LIML) results instrumenting with the lagged grant shares directly in a many-IV setting. The choice for the LIML estimator is due to its better small-sample properties in settings with many instruments where some instruments are weak. As expected, the coefficients are not close to the ones in Table 2.1 although grant shares show sufficient instrumental relevance. Moreover, we reject the null hypothesis of the overidentification tests. To understand this failure in our context, we can calculate year-specific fiscal multipliers $\hat{\beta}_t$ by instrumenting Equation 2.1 with the year-specific shares. If we assume heterogeneous effects between the fiscal multipliers $\hat{\beta}_t$ calculated over different years, the failure of the overidentification tests does not point toward instrument misspecification since we expect some dispersion in the $\hat{\beta}_t$'s caused by heterogeneity.

We, then, analyze the heterogeneity in the individual $\hat{\beta}_t$ by calculating their respective Rotemberg weights. I plot the heterogeneity of $\hat{\beta}_t$ in Figure B2 in the Appendix where the size of each point is scaled by its Rotemberg weight and where I use the specification in column (8) of Table 2.1. The figure shows why the overidentification test failed: there is significant dispersion in the $\hat{\beta}_t$'s, particularly between 2011 and 2012 which are the two years with the largest weights.

I report in Table B.7 in the Appendix the summary statistics for the Rotemberg weights along with details on the years with the largest weights. The main takeaway is that the two years with the highest weights (i.e. 2011 and 2012, accounting for more than 85% of the yearly weight) were years of large reductions in central government grant funding, as seen in Panel C. In 2011 and 2012, 97% and 90% of LADs saw a reduction in their grant funding, respectively, with a cumulative aggregate decline of around 0.6% of GDP in real terms. We can then explain the heterogeneity in $\hat{\beta}_t$ between 2011 (0.95) and 2012 (4.81) which is behind the failure of the overidentification test in Table B.6. While the first year of fiscal austerity seems to have led to some increase in service efficiency, causing

[76] C.f. Nakamura & Steinsson (2014), A. J. Auerbach & Gorodnichenko (2012). The presence of heterogeneous effects when using a shift-share instrument to calculate fiscal multipliers was also noted in Brandao-Roll et al. (2024).

the multiplier estimate to be lower than the average estimate of 1.69 as local GDP did not as much as spending, the subsequent reduction in funding was linked to a large drop in local GDP as the multiplier estimate becomes larger. This yearly heterogeneity is similar to the one observed between periods of fiscal expansion and contraction,[77] although my results highlight the importance of understanding the context when fiscal changes happen as one year of contraction might not lead to the same conclusion as two years of fiscal austerity.

As such, we can attribute the failure of the overidentification test to the yearly heterogeneity stemming from two sequential periods of fiscal contraction. It is then natural that the overidentification tests in the many-IV 2SLS estimation failed and should not be seen as a sign of misspecification. Moreover, as seen in Figure B2 the years with large weights also have large first-stage F-statistics which is reassuring from a small-sample bias point of view. Finally, although we observe years with negative weights they are not relevant for the overall result as the combined weight of positive-weighted years corresponds to around 96% of the overall weight share as shown in Panel A of Table B.7.

2.4.3 Underlying Channels and Local Multiplier Heterogeneity

In this section, I analyze both the channels through which local public spending increases local GDP and the underling fiscal multiplier heterogeneity between local authorities with respect to type of spending, local labor market, and demographic characteristics. The goal is to understand how the fiscal multiplier varies according to local-level aspects and to show that this local heterogeneity cannot be explained by the expected variation in marginal propensities to consume.

To assess whether we observe heterogeneity in the fiscal multiplier at the LAD level, I use the following specification:

$$\begin{aligned} \frac{Y_{l,t+1} - Y_{lt}}{Y_{l,t-1}} &= \beta \frac{G_{lt} - G_{l,t-1}}{Y_{l,t-1}} + \delta_2 \frac{G_{lt} - G_{l,t-1}}{Y_{l,t-1}} \text{ x } D_{l,t-1} + \\ & \delta_1 D_{l,t-1} + \gamma X_{l,t-1} + \phi_l + \psi_t + \epsilon_{lt} \end{aligned} \tag{2.5}$$

where $D_{l,t-1}$ is the council-level characteristic of interest at time $t-1$. As in Section 3.3, I instrument local spending growth with the shift-share instrument based on the lagged grant share of local GDP, both by itself and interacted with $D_{l,t-1}$.

I show results using the preferred benchmark specification in Table 2.3. I analyze the fiscal multiplier heterogeneity using the following local variables, all of which have been demeaned and normalized to have unit variance: per-capita economic inactivity,[78] share of those who are inactive who want to work, average

[77] A point made in Riera-Crichton et al. (2015), Jordà & Taylor (2016), and Pragidis et al. (2018).

[78] Inactivity is defined as people not in employment who have not been seeking work in the previous 4 weeks and/or are unable to start work within the next 2 weeks.

anxiety level,[79] low-skilled labor per capita, child-poverty rate, per-capita number of people receiving benefits, average wage of full-time workers, and wage inequality measured as the wage ratio between the 60^{th} and the 20^{th} percentiles. We assess this table in two parts. First, columns (1)-(4) show considerable heterogeneity in the fiscal multiplier with respect to local variables. In columns (1) and (4) the coefficient of the interaction is negative and statistically significant. This suggests that higher economic inactivity and low-skill labor share decrease the service and capital spending fiscal multipliers, respectively. Columns (2) and (3), on the other hand, report the opposite: the coefficient on the interaction is positive. This implies that having more people in inactivity who want to work and higher anxiety levels increase the local fiscal multiplier. All specifications include time and council fixed-effects, council-level controls (including the interaction term), and I am also controlling for the local authority spending share of different spending categories to show that results are not driven by heterogeneity in spending categories.[80] First-stage robust F-statistics are all above the usual threshold for weak instruments.

Figure 2.2 summarizes the effect of the local variables on the fiscal multiplier. The arrows show how the baseline multiplier (around 1.7 for both service and capital spending) changes given a one-standard-deviation increase to the mean of each local characteristic in Table 2.3, i.e. each arrow starts at the baseline multiplier estimate $\hat{\beta}_{base}$ and ends at $\hat{\beta}_{base} + \delta_2$ all else constant, where δ_2 is the interaction coefficient in Equation 2.5 and Table 2.3. For example, as we increase by one standard deviation the per-capita inactivity level from its mean value, the estimated multiplier decreases from 1.7 to 1.1, all else being equal, which corresponds to a 35% decrease in the multiplier. This is evidence that there is significant fiscal multiplier heterogeneity as a function of local-level parameters, as those are able to shift the fiscal multiplier around a range of values from 0.6 to 3.0.

We now analyze the second part of Table 2.3. As has been shown in the literature on fiscal multipliers, we would expect variation in estimates driven by MPC variation at the local level if fiscal spending here were about transfers.[81] I argue through three points that the heterogeneity observed in Table 2.3 cannot be entirely explained by differences in local MPCs. First, as aforementioned I do not include fiscal transfers as part of local public spending. Although public services and capital can still increase GDP via MPCs by boosting disposable incomes as

[79] Anxiety levels are measured via well-being a survey conducted by the ONS and vary on a scale from 0 to 10.

[80] Spending categories shares included in the set of controls are transportation, education, social-care, housing, cultural, planning, central, and environmental.

[81] C.f. Kaplan & Violante (2018) for a review of models that can generate significant MPC heterogeneity. On the literature of fiscal multiplier heterogeneity from MPCs, c.f. Anderson et al. (2016), Brinca et al. (2016), Carroll et al. (2017), to cite just a few.

Table 2.3: Local-Level Fiscal Multiplier Heterogeneity

	GDP$_{t+1}$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Service Spending	2.324**	1.606**	1.795		1.754*	1.460	1.521*	1.536
	(0.902)	(0.812)	(1.824)		(0.983)	(0.958)	(0.807)	(1.011)
Service Spending $\times$ D	-0.605*	0.439*	1.320**		-0.052	0.104	-0.211	-0.118
	(0.314)	(0.227)	(0.588)		(0.253)	(0.229)	(0.264)	(0.522)
Capital Spending				1.644***				
				(0.530)				
Capital Spending $\times$ D				-1.136*				
				(0.625)				
Interaction	INA	WTW	ANX	LSK	CPV	BEN	INE	WAG
N	3,228	3,228	2,369	2,702	3,228	3,228	3,228	3,228
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Council FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Spending Category	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Outside AEF Grants	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Robust F-test	46.4	49.0	12.6	14.2	40.5	37.9	44.3	30.2

Notes: Interaction variables: INA - per-capita inactivity, WTW - 3-year rolling average of the share of those in inactivity who want to work, ANX - average anxiety level, LSK - low-skilled labor per capita, defined as those with a fail GCSE grade mark (or equivalent) or lower, CPV - child-poverty rate, BEN - per-capita number of people receiving central government benefits, INE - wage inequality measured as the ratio of the 60^{th} percentile over the 20^{th} percentile, WAG - real average full-time wage. Main regressors correspond to growth in real local authority service and capital expenditures per-capita. GDP$_{t+1}$ corresponds to local GDP per-capita growth one period ahead. Local-level controls (one-year lagged): D, share of NVQ3 awards (equivalent to a high-school diploma), unemployment rate, median age, child poverty rate, average full-time wage, per-capita number of people receiving central government benefits (for Disability Living Allowance, Incapacity Benefits, Housing Benefit, Universal Credit, Personal Independence Payment, and Child Benefit) where each claimant counts by the number of individual benefits they receive, dummies for the political party controlling the local council, LAD reserves per-capita, per-capita amount of funds from non-domestic rates, average council tax, and 60-to-20^{th} wage inequality. Outside AEF Grants is the per-capita amount of all grants outside the AEF which are not included in the SSIV. Spending Category controls for the one-year lagged share of spending in transportation, education, social-care, housing, cultural, planning, central, and environmental. Standard errors are clustered by counties for non-metropolitan districts and by individual LAD for the rest. *, **, *** denote significance at 10%, 5%, and 1% levels, respectively.

LADs provide free services that may substitute those offered by private firms that charge for them, this can only happen indirectly here as funds have to be spent to provide an actual service.[82] Second, I use the specification in Equation 2.5 to assess the effect on the multiplier of variables that are more commonly associated with variation in MPCs which I show in columns (5)-(8). These variables are: child-poverty rate, per-capita number of people receiving benefits, 60^{th}-to-20^{th} wage inequality, and average wage of full-time workers. Interestingly, none

[82]We can also consider the more direct MPC channel from public service wages as local authorities may hire more workers with their funds. However, not only the employment fiscal multiplier in Table 2.1 is not statistically significant but it is also unclear how this hiring would correlate with the local variables in Table 2.3 so as to explain the multiplier heterogeneity via public service employees' MPCs.

Figure 2.2: *Expected Change in the Fiscal Multiplier for Different Local Characteristics*

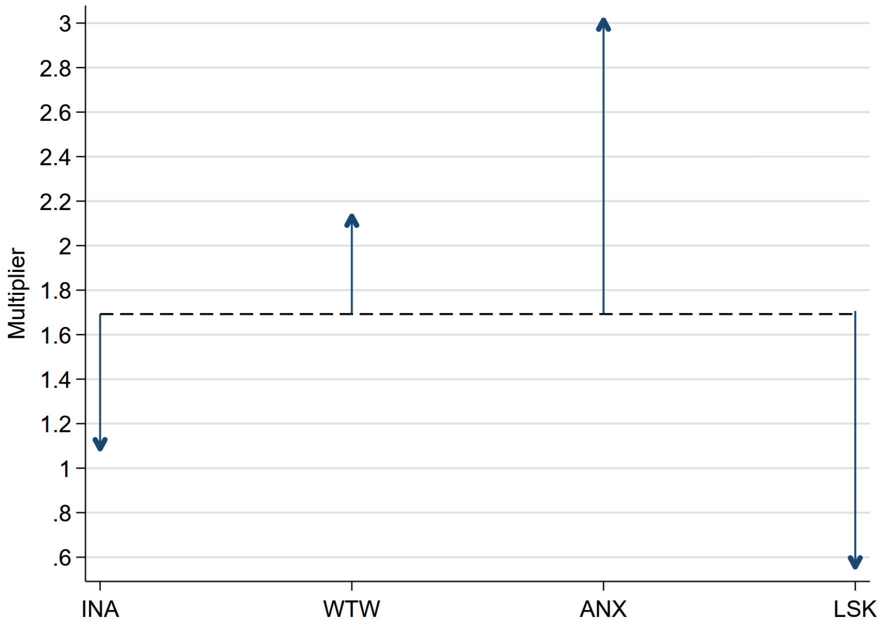

Notes: Expected change is calculated as the change in the baseline fiscal multiplier from Column (8) of Table 2.1 (dashed horizontal line) given a one-standard-deviation increase to the in-sample mean of each interaction variable in Table 2.3. Interaction variables: INA - per-capita inactivity, WTW - 3-year rolling average of the share of those in inactivity who want to work, ANX - average anxiety level, LSK - low-skilled labor per capita, defined as those with a fail GCSE grade mark (or equivalent) or lower.

of the results shows evidence of the MPC channel in explaining multiplier heterogeneity as all interaction coefficients are not statistically significant even though we would expect MPCs to correlate positively with the first three variables and negatively with average wages (controlling for inequality). As poor households usually have high MPCs, we would expect the interaction coefficients in columns (5)-(7) to be positive and significant, and negative and significant in column (8) if results were being driven by differences in marginal propensities to consume. Moreover, I already control in Table 2.3 for observables associated with MPC heterogeneity such as child poverty, unemployment, and inequality. While this evidence does not rule out completely that fiscal multiplier heterogeneity is being driven by differences in MPCs, it does show that my results are not being majorly driven by MPC heterogeneity.

I make a third argument against the MPC channel by providing direct evidence of alternative fiscal multiplier heterogeneity channels. I show in Table 2.4 two important sets of results. First, I show evidence of local heterogeneity in how fiscal spending affects labor productivity and employment. Columns (1)-(5) show evidence that the fiscal multiplier heterogeneity in Table 2.3 can be explained by how local variables change the relationship between local spending and labor market outcomes. While column (1) shows that higher inactivity low-

Table 2.4: Local-Level Fiscal Multiplier Heterogeneity Channels

	Hourly Productivity			Per-Job Productivity	Emp	DALY	Child Poverty	Inactivity But Wants Job
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Service Spending	2.070**	1.133	2.801			-1.004***	-5.519***	-19.56**
	(0.995)	(0.890)	(1.971)			(0.345)	(1.837)	(8.706)
Service Spending × D	-0.800**	0.542*	0.997*					
	(0.331)	(0.307)	(0.570)					
Capital Spending				2.741***	-0.886			
				(0.804)	(1.125)			
Capital Spending × D				-1.719*	-3.065*			
				(1.022)	(1.652)			
Interaction	INA	WTW	ANX	LSK	LSK			
N	3,228	3,228	2,369	2,702	2,702	2,938	3,235	2,938
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Council FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Spending Category	Yes	Yes	Yes	Yes	Yes			
Outside AEF Grants	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Robust F-test	46.4	49.0	12.6	12.7	12.7	106.5	108.4	111.7

Notes: Interaction variables: INA - per-capita inactivity, WTW - 3-year rolling average of the share of those in inactivity who want to work, ANX - average anxiety level, LSK - low-skilled labor per capita. Main regressors correspond to growth in real local authority service and capital expenditures per-capita. Dependent variables in columns (1)-(3) are one-year-ahead growth rates, while all others are calculated as the change from $t-1$ to $t+1$. Per-job and hourly productivities are the ratio between deflated local GDP and total number of local jobs and total number of hours worked, respectively. Productivity growth variables are smoothed with a 3-year moving average. Local-level controls (one-year lagged): D, share of NVQ3 awards (equivalent to a high-school diploma), unemployment rate, median age, child poverty rate, average full-time wage, per-capita number of people receiving central government benefits (for Disability Living Allowance, Incapacity Benefits, Housing Benefit, Universal Credit, Personal Independence Payment, and Child Benefit) where each claimant counts by the number of individual benefits they receive, dummies for the political party controlling the local council, LAD reserves per-capita, per-capita amount of funds from non-domestic rates, average council tax, and 60-to-20th wage inequality. Outside AEF Grants is the per-capita amount of all grants outside the AEF which are not included in the SSIV. Spending Category controls for the one-year lagged share of spending in transportation, education, social-care, housing, cultural, planning, central, and environmental. Standard errors are clustered by counties for non-metropolitan districts and by individual LAD for the rest. *, **, *** denote significance at 10%, 5%, and 1% levels, respectively.

ers the hourly productivity boost from local service spending, columns (2) and (3) show the opposite for the share of those in inactivity who want to work and the average anxiety level. As for capital spending, columns (4) and (5) show evidence that more low-skilled labor dampens the increase in per-job productivity and can lead to a negative employment multiplier. Second, Table 2.4 shows evidence using the benchmark specification that local public spending reduces the impact of disease on the population, child poverty, and inactivity for those who want to work.

Results, then, complement those in Table 2.3 as they show evidence of what is driving the heterogeneity in the fiscal multiplier. As both labor productivity and employment levels are affected differently by local spending depending on local characteristics, results show evidence that local services and capital have a direct effect on labor market parameters. Such results are hard to explain via the MPC channel.[83] Similarly, as fiscal spending lowers the incidence of disease risk factors, poverty, and economic inactivity through health and social care, we can associate the improvement in social conditions with better labor market outcomes.

Finally, I complement the previous result on labor market outcomes and improvements in local social conditions with evidence of heterogeneity by spending category. As local authorities provide services for several spending categories, we can assess whether each category affects local GDP differently. However, as I only have an instrument for total service expenditure I cannot control for simultaneous changes in spending for each category which raises the issue of potential bias in estimates as LADs can shift resources from one category to another. Instead of using Equation 2.1, I use the change in the category-specific share of total spending, i.e. $(G_{clt}/G_{lt}) - (G_{cl,t-1}/G_{l,t-1})$ where G_{clt} is local public spending in category c, as my main regressor. The resulting estimate can be interpreted as the increase in local GDP growth (in 100x percentage points) if the category share of spending increases by 1. While this potential increase in the spending share is unrealistic as it would imply a level of spending that is above total actual expenditure, estimates are informative of how spending categories compare relative to one another. I show results in Table B.8 in the Appendix for the categories where the first-stage robust F-statistic was higher than 10 or close to it. We can observe that increasing social-care and planning spending shares has a positive effect on local growth, whereas increasing the education spending share has a negative effect on GDP. Although estimates for the social-care and planning categories in columns (2) and (3) are statistically indistinguishable from each other, both are higher than the education coefficient. Along with the previous evidence on local outcomes and heterogeneity, as well as the fact that social care is the second largest spending category as seen in Figure B3 in the Appendix, this heterogeneity in how different spending categories affect GDP growth points towards a fiscal multiplier mechanism where social-care services play a major role. As a side note, the negative coefficient for education spending can possibly be explained by the fact that most of the return from education takes many years to realize and is likely more reliant on inter-council migration patterns. Nonetheless, from a short-term point of view social-care spending seems to be the most relevant in raising local GDP.

These results on fiscal multiplier heterogeneity make the case for moving the local fiscal multiplier discussion away from a debate around a single value. Het-

[83] Given the lack of data on capital utilization, it is not possible to disentangle the observed increase in labor productivity and an increase in the capital utilization rate. However, it is hard to see how a higher capital utilization rate supported by higher MPCs would explain the heterogeneity observed in Table 2.3, i.e. how capital utilization correlates with the local observables.

erogeneity in the labor market, demographic parameters, and spending categories is able to explain significant variation in estimated multipliers. Moreover, MPC heterogeneity seems unable to explain our results. If local fiscal spending can increase labor productivity, reduce economic inactivity, and improve well-being outcomes in health and poverty, the corollary is that local public spending can be thought of as being productive, and public sector productivity varies along local characteristics. I discuss a possible mechanism behind the observed heterogeneity in the fiscal multiplier in Section 2.5.

2.5 Heterogeneous Labor Model

To analyze the channel through which fiscal policy produces heterogeneous effects at the local level due to local demographic and labor market characteristics, I consider a simple model of labor heterogeneity and fiscal policy. As shown in Section 2.4.3, any mechanism trying to explain the fiscal multiplier heterogeneity must account for the role of social-care spending and local socio-economic conditions. Here, I propose that workers are subject to a low mental bandwidth shock which effectively lowers their productivity until a high bandwidth shock arrives. This low-bandwidth shock represents different situations where the agent's cognitive bandwidth becomes scarce, e.g. poverty, sickness, and/or having to take care of a relative. Once hit by a negative shock, worker productivity drops. It then only recovers through a high bandwidth shock that is a function of government spending and worker heterogeneity regarding how hard it is to bring them back to their normal cognitive state. Given model results, I show how a social planner would opt to spend differently depending on local conditions and how this fiscal allocation compares with local public spending in the UK.

We start the set-up with two types of workers: high- and low-skilled. We normalize total labor to one and assume there is an equal split between the two types in the aggregate. Workers are distributed between regions j and there is a single firm in each region that produces good c_{jt} at time t. Labor is inelastically supplied and we consider the symmetric case where there is the same amount of labor supply in each region which is entirely hired by the local firm though the shares of high- and low-skilled labor vary between regions. Agent i has the following utility function:

$$\begin{aligned} U_i &= \sum_{t=0}^{\infty} \beta^t ln(C_{it}) \\ ln(C_{it}) &= \int_0^1 ln(c_{ijt})dj \end{aligned} \tag{2.6}$$

where β is the discount factor. We normalize the price of the consumption good to 1. Firm j has the following production function:

$$Y_{jt} = (\gamma_{jt} L_{jt})^{\alpha}$$
$$\gamma_{jt} = \frac{1}{L_{jt}} \sum_{i} \gamma_{ijt} \tag{2.7}$$

where γ_{jt} is the average labor productivity of workers (whose individual productivity is γ_{ijt}) hired by firm j, L_{jt} is labor, and $\alpha \in (0,1)$ is a constant. Workers in each locality j are exposed to a low mental-bandwidth shock which arrives at a Poisson rate $\lambda_{jt,lb}$. Once in the low cognitive state, a worker may be hit by a high-bandwidth shock at a Poisson rate $\lambda_{jt,hb}$ which brings them back to their normal state. As such, average productivity γ_{jt} changes depending on how many workers are in their low-bandwidth state at any given time. Let $\gamma_{ijt} \in [\gamma_{hs,hb}, \gamma_{hs,lb}, \gamma_{ls,hb}, \gamma_{ls,lb}]$ where subscripts hb and lb mean high and low bandwidth, respectively, and hs and ls mean high and low skill, respectively. We assume $\gamma_{k,hb} \geq \gamma_{k,lb}$ for $k = hs, ls$ and $\gamma_{hs,k'} \geq \gamma_{ls,k'}$ for $k' = hb, lb$. Hence, workers switch between a state of low and high cognitive bandwidth, and these shocks affect labor productivity. A firm's average productivity will then depend not only on the share of high- and low-skilled labor in their region but also on the incidence of shocks. We assume that firms know when a worker is in their high- or low-bandwidth states so that it can change wages accordingly.

As for the local government, it funds itself with a lump-sum tax $T_{ag,jt}$ on households and uses its revenues $G_{ag,jt}$ (where $G_{ag,jt} = T_{ag,jt}$) to invest in social and health care. Hence:

$$\lambda_{jt,hb} = f(G_{ag,jt}) \ , \ f' > 0 \ , \ f'' < 0 \tag{2.8}$$

Next, we make an important distinction regarding how workers are affected by the low-bandwidth shock. We assume a share of workers is of "high-risk" type, i.e. when hit by a low-bandwidth shock they require more government spending, relative to the non-high-risk type, to be brought back to their normal cognitive state. The idea is to capture the difference between relatively easily revertible shocks (i.e. malnutrition) and more life-altering ones (i.e. becoming a carer for a partner with a long-term condition).[84] For high-risk individuals, the high-bandwidth shock arrives at a rate of $\theta f(G_{ag,jt})$, where $\theta \in (0,1)$. Moreover, notice that while $\lambda_{jt,hb}$ is determined by government spending, $\lambda_{js,lb}$ is exogenous.

In order not to mechanically over-tax low-skilled households relative to high-skilled ones and to take into account that high-earning agents have a larger tax

[84]Given inelastic labor supply, I do not allow workers to leave the labor force altogether when hit by a low-bandwidth shock. While that might be more realistic and explain heterogeneity with respect to the employment multiplier, it does not change the main results.

burden, we consider that high-skilled households pay a larger share of taxes which is proportional to how much they earn relative to low-skilled households. As such, given a high-skill share of φ_j in local area j, then:

$$T_{ag,jt} = \varphi_j T_{jt} \frac{w_{hs,hb}}{w_{ls,hb}} + (1 - \varphi_j) T_{jt} \tag{2.9}$$

where $w_{hs,hb}$ and $w_{ls,hb}$ are the wages of high- and low-skilled workers, respectively, at their high-bandwidth state. Given our assumption that firms adjust wages after the arrival of shocks, the variation in wages comes from both skills and cognitive states, i.e. $w_{jt} \in [w_{hs,hb}, w_{hs,lb}, w_{ls,hb}, w_{ls,lb}]$. Finally, agents have the following budget constraint in expectation:

$$C_{it} = \begin{cases} E_t[w_{jt} - T_{jt} \frac{w_{hs,hb}}{w_{ls,hb}}], \ if \ high - skill \\ E_t[w_{jt} - T_{jt}], \ if \ low - skill \end{cases} \tag{2.10}$$

We can, then, write the household problem. Since variables do not have time-dependency, we can solve the problem statically (I remove subscripts where their absence does not hinder interpretation):

$$\max_{\{C_k\}_{hb,lb}} E(U) = \max_{\{C_k\}_{hb,lb}} \begin{cases} \frac{f(G)}{f(G)+\lambda_{js,lb}} ln(C_{hb}) + \frac{\lambda_{js,lb}}{f(G)+\lambda_{js,lb}} ln(C_{lb}), \ if \ not \ high - risk \\ \frac{\theta f(G)}{\theta f(G)+\lambda_{js,lb}} ln(C_{hb}) + \frac{\lambda_{js,lb}}{\theta f(G)+\lambda_{js,lb}} ln(C_{lb}), \ if \ high - risk \end{cases} \tag{2.11}$$

Equation 2.11 implies that each agent chooses consumption to maximize expected utility which depends on the arrival rate of both the high- and low-bandwidth shocks, and whether they are high-risk individuals or not.

As for firms, they maximize per-period profits:

$$\max_{\{L_{ijt}\}_i} (\gamma_{jt} L_{jt})^\alpha - \sum_i L_{ijt} w_{ijt} \tag{2.12}$$

which implies that labor is paid at marginal productivity, i.e. $w_{ijt} = \alpha \gamma_{ijt} (\gamma_{ijt} L_{ijt})^{\alpha - 1}$.

Finally, we consider the problem of the local social planner who chooses G_{jt} to maximize aggregate welfare:

$$G_{SP} = \arg\max_{G_{jt}} ln(G_{ag,jt}) + \sum_i E(U_i) \tag{2.13}$$

where I make the important assumption that, along with being productive, government spending happens through the wages of LAD employees. I assume public servants are included in aggregate welfare though they are not subject to low-bandwidth shocks, for simplicity.

We can then solve Equations 2.10, 2.11, 2.12, and 2.13. To do so, I match model parameters and moments to values in the UK data. First, I define $\lambda_{jt,hb} = BG_{jt}^{\delta}$, where $\delta \in (0, 1)$ and B is a constant. Then, I calibrate the productivity parameters by matching the average wage rate for high- and low-skilled labor. Since I do not observe individual workers in the data, I define the high-skill (low-skill) wage rate as the average wage in councils where the share of workers at the NVQ4 education level (roughly equivalent to university degree holders) is in the top (bottom) quartile of the share distribution. I then choose $\alpha = 0.8$ and $\delta = 0.8$.

This leaves us with two parameters to estimate: $\{B, \theta\}$. I proceed with a GMM estimation using the following three moments: the average fiscal multiplier estimated with the preferred specification in Table 2.1 (column 4), the average real local GDP per capita, and the ratio between the average GDP of councils above and below the median value of per-capita economic inactivity of people who do not want to work. While the latter moment will help us pin down θ, the first two are directly influenced by B. The empirical low-cognitive shock incidence is calculated using the average of the standardized (i.e. demeaned and with unit variance) local DALY, unemployment rate, and child-poverty rate.[85] I then set the minimum value of the shock arrival rate at zero and translate all values accordingly. As for the share of high-risk people, I use the per-capita amount of people in inactivity who do not want to work.

We can then proceed with the estimation by taking into account a few important details. Since the fiscal multiplier estimate in Section 3.3 represents the relative effect, between councils, of an additional pound spent by the local government, I calculate the model-based multiplier as the increase in consumption from a one-pound increase in public spending without an equivalent increase in local taxation. We revert this procedure when we solve for the optimal fiscal policy from the central government's point of view. Finally, I consider that the low-bandwidth shock leads to a drop in average worker productivity of 20%.[86]

I show in Table 2.5 the parameter estimates as well as the moment fit. Overall, the fit is good. We can also check whether model-estimated multipliers show the same heterogeneity as the one observed empirically in Table 2.3 by regressing these multipliers on the low-bandwidth rate $\lambda_{jt,lb}$, the high-skill share, and one minus the high-risk share (so as to match Table 2.3). I do so in Table 2.6 which shows that model results match the empirical heterogeneity: while having a higher share of high-skilled workers and of those in inactivity who want a job

[85]Since data on anxiety levels only start in 2011 I do not use the estimated multiplier heterogeneity due to anxiety when calculating the expected local multiplier. Notice as well that the sample size for the capital spending multiplier estimation is smaller due to it being over two-year growth rates.

[86]This accounts for any form of psychological toll that reduces one's cognitive bandwidth, c.f. Schilbach et al. (2016) and Kaur et al. (2021) on poverty-induced stress, and Schultz & Edington (2007) for a summary on health. Although estimates for individual maladies are below 20%, the low-bandwidth shock captures the combined effect of cognitive tolls, including those leading to absenteeism which are harder to assess in experiments.

increases the fiscal multiplier, having a higher incidence of the low-bandwidth shock lowers the multiplier. Results can be explained by higher returns to public spending when resources are used to revert the negative shock to more productive workers. Similarly, having a higher share of workers who are not high-risk means a higher "bang for the government's buck."[87] As for the low-bandwidth shock, a higher $\lambda_{jt,lb}$ reduces the period of time a worker can expect to stay healthy, lowering the multiplier.

Table 2.5: Model Estimation and Moment Fit

Parameter	Value	Parameter	Value
γ	1.05	κ	1.34
b	0.62	h_l	1.90
ρ	3,084	h_c	0.31
A_l	2.23	ν	0.21
A_λ	29.6		
Moments		**Data**	**Model**
Growth Rate (%)		1.31	1.32
Skill Premium, Large Firms		2.76	2.77
Labor Market Tightness		0.48	0.48
High-Skill Wage, Non-Large Firms		0.58	0.54
High-Skill Concentration		0.59	0.59
Firm Profitability		0.20	0.21
R&D Investing-to-Sales Ratio (%)		0.19	0.21
Cost-per-Hire		0.12	0.11
High-Skill Unemployment		0.19	0.22
Share of High-Skill Concentration $\leq$ 50%		0.38	0.40
Non-Targeted Moment		**Data**	**Model**
R&D Worker Share (%)		0.91	0.70

This model, hence, is able to rationalize the local heterogeneity in the fiscal multiplier shown in Section 2.4.3. It also highlights the role of fiscal policy in targeting both health and social care, represented here by the low-bandwidth shock $\lambda_{jt,lb}$ and the share of high-risk individuals. As workers are affected by the mental (and at times physical) toll of poor health and social problems such as poverty, unexpected carer duties, and food insecurity, they become less productive. Government, then, has a role to play in choosing fiscal spending so as to provide social care.

We can finally analyze the social planner's optimal fiscal policy and compare it with actual local government grant allocation in the UK. Our planner takes the point of view of the central government who collects funds through taxation and distributes grants to local councils. Note, however, that the planner's problem described in Equation 2.13 corresponds to the local government's perspective. Taking now the point of view of the central government allows us to

[87] We can also interpret the high-risk share as the amount of public funds that government has to spend on welfare aspects that do not show up in the usual utility function, such as retiree healthcare, which are in the government's mandate. To be sure, the model does not take into account all aspects important for welfare.

Table 2.6: Model-Estimated Fiscal Multiplier Heterogeneity

	Multiplier
	(1)
High − Skill Share	0.242***
	(0.042)
$\lambda_{jt,lb}$	-0.060***
	(0.018)
Want To Work	0.191***
	(0.049)
N	2,410
Time FE	Yes
Council FE	Yes
Controls	Yes
Outside AEF Grants	Yes

Notes: Multiplier refers to the model-generated fiscal multipliers. High-skill share is the share of workers at the NVQ4 education level (roughly equivalent to university degree holders). $\lambda_{jt,lb}$ is the low-bandwidth rate estimated using the average of the standardized (i.e. demeaned and with unit variance) local DALY, unemployment rate, and child-poverty rate. Want To Work is the share of people in inactivity who want to work. Local-level controls: share of NVQ3 awards (equivalent to a high-school diploma), unemployment rate, median age, child poverty rate, average full-time wage, per-capita number of people receiving central government benefits (for Disability Living Allowance, Incapacity Benefits, Housing Benefit, Universal Credit, Personal Independence Payment, and Child Benefit) where each claimant counts by the number of individual benefits they receive, dummies for the political party controlling the local council, LAD reserves per-capita, per-capita amount of funds from non-domestic rates, and average council tax. Outside AEF Grants is the per-capita amount of all grants outside the AEF which are not included in the SSIV. Standard errors are clustered by counties for non-metropolitan districts and by individual LAD for the rest. *, **, *** denote significance at 10%, 5%, and 1% levels, respectively.

assess whether local characteristics are being taken into account in deciding local fiscal policy at the national level. Particularly, we distinguish two complementary levels of analysis: a relative one, where within a cross-section of LADs some local authorities should get more funds relative to others from a point of view of the relative fiscal multiplier between them, and an absolute one, which corresponds to optimal local fiscal policy as calculated in the model. We start with the former by calculating the average expected multiplier for each council-year pair using results from Table 2.3 given local inactivity and low-skill share. I show results for the distribution of average multipliers pre- and post-2014 in Figure B4 in the Appendix for England. Two points are worth noticing. First, there has been an overall increase in the multiplier over time as the average jumps from 1.69 pre-2014 to 1.98 post-2014 which, from the social planner's point of view, would imply higher public spending. Second, there is substantial cross-LAD heterogeneity with an interquartile range of 0.58 pre-2014 (0.56 post-2014). This heterogeneity points towards gains from optimally rearranging the central government grant allocation even if the total amount being awarded does not change.

We start by analyzing possible gains from a better allocation of fiscal support. The natural premise behind an optimal allocation of central government grants to LADs is that councils with a higher relative fiscal multiplier should receive a fiscal boost. We assume, then, that the national government can reallocate grant funds between local authorities although it cannot change the total amount spent. We do so for each year separately, for service spending only, and for the LAD-year pairs for which we have an estimate of the local multiplier. Naturally, the best unrestricted allocation is one that gives all funding to the LAD with the highest multiplier. I adopt a more realistic approach which assumes that the central government has to allocate grants to councils where the estimated fiscal multiplier is above 1 (a "spending that pays for itself" approach). I then redistribute the total amount spent on grants proportionally to the estimated fiscal multiplier and plot the yearly actual and optimal average multipliers in Figure 2.3, along with the difference between both. As mentioned before, we can see that both actual and optimal multipliers increase over time. There has also been an increase in allocative efficiency of central government grants as the difference between optimal and real multipliers drops from 0.19 in 2010 (12.8% increase in the actual multiplier) to around 0.11 in 2019 (5.4% increase). However, the loss in local GDP due to this *fiscal misallocation* between LADs is still significant and can be measured both in terms of central government grants, which are within the purview of the national government, or total local spending, assuming the central government could directly control local budgets. Had the fiscal misallocation been corrected since 2010 for grants, it would have generated an extra 57.9 billion pounds in real terms from the total grants, enough to revert the real cumulative reduction in total grant awards to local authorities since the austerity program started in 2010 (around 18 billion pounds) and double the 2019-20 total grant funding in real terms (around 39.9 billion pounds). In a more optimistic scenario where the central government could allocate the total amount spent by local budgets instead of just grants, the cumulative gains since 2010 total around 156 billion pounds in real terms, or 19.2% of the central government budget in 2019 in real terms.

We now assess fiscal spending through the local planner's solution. Although grant fiscal misallocation declined over time, the increase in multipliers for most LADs implies that a benevolent social planner would choose to increase local spending. We can, then, solve Equation 2.13 using the estimated model parameters and compare the optimal spending G_{SP} with actual grant and total local spending. Importantly, we are now taking into account that any extra pound spent needs to be balanced via taxation. I show results averaged in time for each LAD in Figure 2.4 for England where I show both the planner's solution and the ratio between G_{SP} and actual spending.[88] There are two main takeaways from results. First, there is considerable heterogeneity between councils as the interquartile range of the optimal to actual spending ratio is 0.84. This reflects our previous results on multiplier heterogeneity pointing at gains from better allocating grants

[88]A close-up of the London region is available in Figure B5 in the Appendix.

Figure 2.3: *Actual and Optimal Average Multipliers*

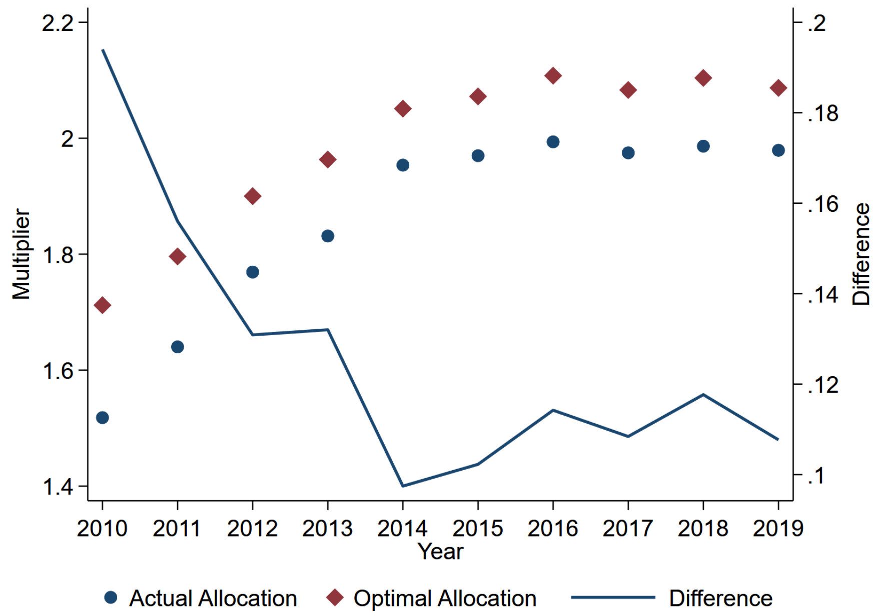

Note: Actual Allocation refers to the average local fiscal multiplier estimated in Table 2.3 using the local-level values for inactivity per capita, share of those in inactivity who want to work, and per-capita low-skilled labor. Optimal Allocation is the estimate of the fiscal multiplier if government were able to reallocate local grants only to councils where the estimated fiscal multiplier is above 1 and proportionally to the estimated fiscal multiplier. Difference refers to the difference between optimal and actual allocations.

and local spending. From a policy perspective, there is a significant gain from better allocating local grants as they have a low correlation (0.18) with the local planner's solution as shown visually in Figure B6 in the Appendix. Second, the social planner would prefer, on average, to increase budgets by 28% (median: 8.8%) since the average multiplier is above one.

Finally, we can gauge the overall impact on GDP from optimal local spending. Relative to the analysis of the fiscal misallocation which measured the effect of an optimal allocation keeping total spending constant and based on relative fiscal multipliers, we now assume government can determine local budgets, collects local taxes, and that fiscal spending acts through the high bandwidth shock. I show in Figure 2.5 the effect on the UK's yearly GDP from adopting the planner's optimal fiscal policy along with actual and optimal total spending. We observe that the gains from optimal fiscal policy have increased over time to a potential increase in per-capita GDP of around 0.65 percentage points in 2019, raising the 1.2% per-capita GDP growth of that year by more than 50%. This is due to the increase in the average local multiplier coupled with the reduction in local spending, increasing the gap between the planner's and actual policies even though the fiscal misallocation improved until 2014 as shown in Figure 2.3.

Results, then, imply that local fiscal policy should be seen through the lens of

Figure 2.4: *Optimal Spending and Optimal to Actual Spending Ratio*

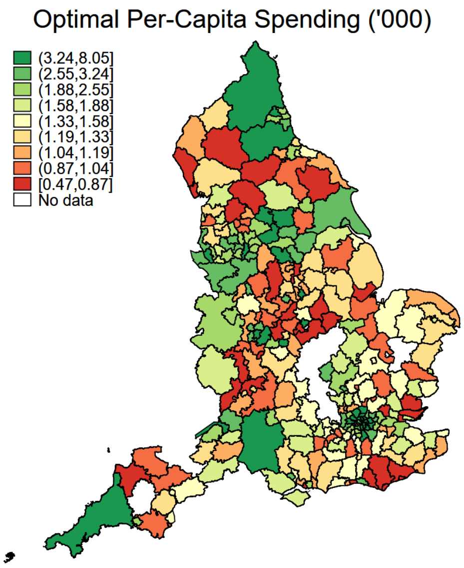

Note: Local areas in the map consist of LADs. Optimal Per-Capita Spending refers to the planner's solution in the model for optimal local spending. Optimal to Actual Spending Ratio is the ratio between the model-estimated optimal spending and actual spending by LADs. Data and model-generated estimates are for 2010 and in pounds.

Figure 2.5: *Actual and Optimal Total Spending, and GDP Increase from Optimal Spending*

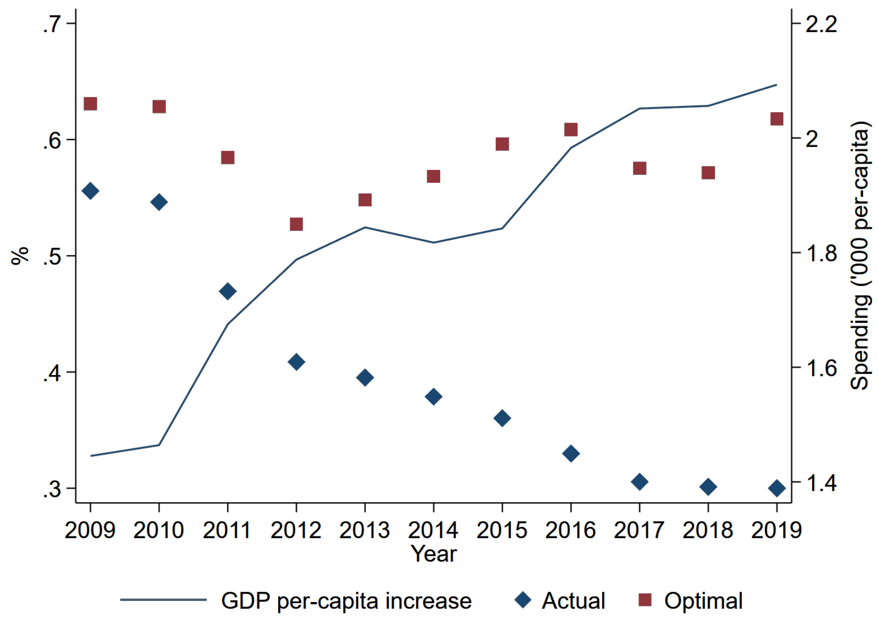

Note: Actual refers to actual per-capita total local spending. Optimal refers to total per-capita local spending from the planner's solution in the model. GDP per-capita increase refers to the weighted average model-estimated change in real per-capita GDP, weighted by local population, from adopting the planner's solution.

local heterogeneity and its causes. As the data show and the model rationalizes, labor market and demographic characteristics are important in driving the value of the fiscal multiplier as government action, particularly through social care, can boost worker productivity. Although in reality matching the planner's optimal spending might seem far-fetched given budgetary concerns, I show that there is scope for gains even in keeping total spending constant as long as an optimal allocation of central government grants reduces fiscal misallocation between local authorities. Importantly, results on LAD heterogeneity show that fiscal policy should not be reduced to a search for a single fiscal multiplier.

2.6 Conclusion

I exploit local variation in the reliance of local councils on central government grants in the UK to estimate the local fiscal multiplier and to understand the role of local heterogeneity in its magnitude. Assuming that the one-year lagged LAD-level grant share of local GDP is exogenous to GDP growth, changes in the aggregate disbursement of grants affect councils differently and these shares can be used in an SSIV to pin down the fiscal multiplier. I estimate an average service multiplier of 1.69 and a capital one of 1.71, both in line with other estimates in the literature on regional multipliers. I show results are robust to running a weighted regression, to changes in the SSIV, and to running the specification at the TTWA level. I also provide evidence to support the shift-share instrument validity assumption via a balance test and an analysis of heterogeneous effects which shows estimate heterogeneity due to two subsequent years of fiscal austerity.

I then proceed to analyze the dependence of the fiscal multiplier on demographic and labor market characteristics. Councils spending similar amounts but with different local characteristics show different GDP responses to local government spending. While a higher share of economic inactivity and low-skilled workers reduce the fiscal multiplier, having more people in inactivity who want a job and higher anxiety levels increase the multiplier. I argue through three points that this heterogeneity is not due to differences in local marginal propensities to consume. First, virtually all local spending that I consider excludes direct transfers. Second, I do not observe heterogeneity in the multiplier with respect to variables related to MPCs. Finally, I show that fiscal spending in services and capital boosts worker productivity and improve local socio-economic and health conditions. These findings point towards a mechanism of fiscal policy effectiveness that both relies on labor and demographic heterogeneity, and has at its core social care which I show is largely responsible for the above-one average multiplier.

I rationalize these results via a model of heterogeneous labor and a shock that affects worker productivity by lowering their cognitive bandwidth. The social planner can, then, use public spending to increase the rate at which workers return to their normal, healthy state. I match the model to the data and show that the model-estimated local multipliers replicate the heterogeneity observed in the empirical results regarding the multiplier. The model highlights a mechanism through which returns to social-care spending change depending on how negatively affected a population is by conditions such as poverty, sickness, and the psychological toll of having to care for someone. I then use both the empirical estimates and the model to assess the current local fiscal policy in the UK. I show that there are significant gains from reducing fiscal misallocation, i.e. the opportunity cost from having more money spent in regions where the multiplier is lower, even if we keep current local public spending unchanged: addressing the misallocation since 2010 would have generated an additional 57.9 billion pounds from an optimally allocated central government grant, or 156 billion pounds in a scenario where government could reallocate total local spending. Using results from the model and if we allow for total spending to change, I show that an optimal local fiscal policy would increase yearly GDP growth by 0.5 percentage points on average. This gap relative to potential output has increased over time due to the estimated increase in the overall fiscal multiplier, which implies a larger optimal fiscal policy and the reduction in real local spending in the UK.

The idea behind empirical estimates and the model is to show that local fiscal multipliers have significant heterogeneity. Results indicate that fiscal policy discussions should take into account the particular demographic and worker-level aspects of a locality before concluding on fiscal expansion or contraction. Given an average, country-wide local-level estimate of the fiscal multiplier, it is not clear whether a local authority should increase or decrease its spending. Moreover, my results suggest the possibility of gains from fiscal policy from a better allocation of local grants and/or local spending even if we keep aggregate expenditure constant. The role of fiscal misallocation between LADs only makes sense if we allow for local heterogeneity. An interesting venue for future research, then, is to consider local fiscal spending in a dynamic setting where current spending affects local parameters which, in turn, may change future fiscal multipliers. The social planner should, then, target not only static gains from optimal fiscal policy but also future increases in the multiplier, two goals that might be at odds and provide an interesting area to assess trade-offs.

References

Anderson, E., Inoue, A., & Rossi, B. (2016). Heterogeneous consumers and fiscal policy shocks. *Journal of Money, Credit and Banking*, *48*(8), 1877-1888.

Auerbach, A., Gorodnichenko, Y., McCrory, P. B., & Murphy, D. (2022). Fiscal multipliers in the covid19 recession. *Journal of International Money and Finance*, *126*, 102669.

Auerbach, A. J., & Gorodnichenko, Y. (2012). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, *4*(2), 1-27.

Basso, H. S., & Rachedi, O. (2021). The young, the old, and the government: Demographics and fiscal multipliers. *American Economic Journal: Macroeconomics*, *13*(4), 110-41.

Berge, T., De Ridder, M., & Pfajfar, D. (2021). When is the fiscal multiplier high? a comparison of four business cycle phases. *European Economic Review*, *138*, 103852.

Boehm, C. E. (2016). *Government Spending and Durable Goods* (CESifo Working Paper Series No. 6244). CESifo.

Borusyak, K., Hull, P., & Jaravel, X. (2021). Quasi-Experimental Shift-Share Research Designs. *The Review of Economic Studies*, *89*(1), 181-213.

Brandao-Roll, J., Ridder, M., Hannon, S., & Pfajfar, D. (2024). The fiscal multiplier of education expenditures [Working Paper].

Brinca, P., Holter, H. A., Krusell, P., & Malafry, L. (2016). Fiscal multipliers in the 21st century. *Journal of Monetary Economics*, *77*, 53-69. ("Inequality, Institutions, and Redistribution"held at the Stern School of Business,New York University,April 24-25, 2015)

Burton, W., Chen, C.-Y., Schultz, A., & Edington, D. (2001). The impact of allergies and allergy treatment on worker productivity. *Journal of Occupational and Environmental Medicine*, *43*, 64-71.

Burton, W. N., Morrison, A., & Wertheimer, A. I. (2003). Pharmaceuticals and worker productivity loss: A critical review of the literature. *Journal of Occupational and Environmental Medicine*, *45*(6), 610–621.

Caldara, D., & Kamps, C. (2017). The Analytics of SVARs: A Unified Framework to Measure Fiscal Multipliers. *The Review of Economic Studies*, *84*(3), 1015-1040.

Capek, J., & Crespo Cuaresma, J. (2020). We just estimated twenty million fiscal multipliers. *Oxford Bulletin of Economics and Statistics*, *82*(3), 483-502.

Carroll, C., Slacalek, J., Tokuoka, K., & White, M. N. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, *8*(3), 977-1020.

Chodorow-Reich, G. (2019). Geographic cross-sectional fiscal spending multipliers: What have we learned? *American Economic Journal: Economic Policy*, *11*(2), 1-34.

Chodorow-Reich, G., Feiveson, L., Liscow, Z., & Woolston, W. G. (2012). Does state fiscal relief during recessions increase employment? evidence from the american recovery and reinvestment act. *American Economic Journal: Economic Policy*, *4*(3).

Christiano, L., Eichenbaum, M., & Rebelo, S. (2011). When is the government spending multiplier large? *Journal of Political Economy*, *119*(1), 78–121.

Corsetti, G., Meier, A., Müller, G. J., & Devereux, M. B. (2012). What determines government spending multipliers? *Economic Policy*, *27*(72), 521–565.

Fishback, P., & Kachanovskaya, V. (2015). The multiplier for federal spending in the states during the great depression. *The Journal of Economic History*, *75*(1), 125–162.

Gechert, S. (2015). What fiscal policy is most effective? a meta-regression analysis. *Oxford Economic Papers*, *67*(3), 553-580.

Gibbons, S., Lyytikäinen, T., Overman, H. G., & Sanchis-Guarner, R. (2019). New road infrastructure: The effects on firms. *Journal of Urban Economics*, *110*, 35-50.

Gibbons, S., & Wu, W. (2019). Airports, access and local economic performance: evidence from China. *Journal of Economic Geography*, *20*(4), 903-937.

Goetzel, R., Ozminkowski, R., Hawkins, K., Wang, S., & Lynch, W. (2004). Health, absence, disability, and presenteeism cost estimates of certain physical and mental health conditions affecting u.s. employers. *Journal of occupational and environmental medicine / American College of Occupational and Environmental Medicine*, *46*, 398-412.

Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, *110*(8), 2586-2624.

Ilzetzki, E., Mendoza, E. G., & Végh, C. A. (2013). How big (small?) are fiscal multipliers? *Journal of Monetary Economics*, *60*(2), 239-254.

Jordà, O., & Taylor, A. M. (2016). The Time for Austerity: Estimating the Average Treatment Effect of Fiscal Policy. *The Economic Journal*, *126*(590), 219-255.

Kaplan, G., & Violante, G. L. (2018). Microeconomic heterogeneity and macroeconomic shocks. *Journal of Economic Perspectives*, *32*(3), 167-94. doi: 10.1257/jep.32.3.167

Kaur, S., Mullainathan, S., Oh, S., & Schilbach, F. (2021). *Do financial concerns make workers less productive?* (Working Paper No. 28338). National Bureau of Economic Research.

Morita, H. (2022). On the relationship between fiscal multipliers and population aging in japan: Theory and empirics. *Economic Modelling*, *108*, 105772.

Muratori, U., Juarros, P., & Valderrama, D. (2023). Heterogeneous spending, heterogeneous multipliers [IMF Working Paper]. (2023/052).

Nakamura, E., & Steinsson, J. (2014). Fiscal stimulus in a monetary union: Evidence from us regions. *American Economic Review*, *104*(3).

Pappa, E. (2009). The effects of fiscal shocks on employment and the real wage. *International Economic Review*, *50*(1), 217-244.

Pragidis, I., Tsintzos, P., & Plakandaras, B. (2018). Asymmetric effects of government spending shocks during the financial cycle. *Economic Modelling*, *68*, 372-387.

Ramey, V. A. (2019). Ten years after the financial crisis: What have we learned from the renaissance in fiscal research? *Journal of Economic Perspectives*, *33*(2), 89-114.

Ramey, V. A., & Zubairy, S. (2018). Government spending multipliers in good times and in bad: Evidence from us historical data. *Journal of Political Economy*, *126*(2), 850-901.

Riera-Crichton, D., Vegh, C. A., & Vuletin, G. (2015). Procyclical and countercyclical fiscal multipliers: Evidence from oecd countries. *Journal of International Money and Finance*, *52*, 15-31.

Schilbach, F., Schofield, H., & Mullainathan, S. (2016). The psychological lives of the poor. *American Economic Review*, *106*(5), 435-40.

Schultz, A., & Edington, D. (2007). Employee health and presenteeism: A systematic review. *Journal of occupational rehabilitation*, *17*, 547-79.

Shoag, D. (2016). *The Impact of Government Spending Shocks: Evidence on the Multiplier from State Pension Plan Returns* (Working Paper).

Yoshino, N., & Miyamoto, H. (2017). Declined effectiveness of fiscal and monetary policies faced with aging population in Japan. *Japan and the World Economy*, *42*(C), 32-44.

Chapter 3

The Fiscal Multiplier of Education Expenditures

Julio Brandao-Roll (LSE), Maarten De Ridder (LSE), Simona M. Hannon (Fed Board), Damjan Pfajfar (Cleveland Fed)

3.1 Introduction

Investments in education make up a significant part of government spending in advanced economies. In the United States, educational spending measured 5.4 percent of national income in 2020, which exceeds defense spending and spending on welfare programs. These investments are usually motivated by the well-documented effects that education has on well-being and economic growth in the long run (see, e.g., Barro, 1991, Benhabib & Spiegel, 1994, Bils & Klenow, 2000, and Manuelli & Seshadri, 2014). Like any other form of government spending, however, educational investments also have the potential to stimulate economic activity in the short run. Programs that reduce the cost of tuition or that involve direct transfers to students could, for example, increase purchasing power and therefore raise consumption and employment. They also unlock complementary sources of income that arise when students go to college, such as student loans. Such programs could be used to stimulate economic activity during recessions and serve as a tool for macroeconomic stabilization. Yet empirical evidence that establishes the magnitude of the short-run effects of educational investments on economic activity remains scant.

We quantify the effect of educational investments on economic growth in the short run. Specifically, we measure the impact of the Federal Pell Grant Program at the city (Metropolitan Statistical Area–MSA) level. Pell grants are need-based grants to low-income undergraduate and select post-baccalaureate students, designed to enable them to access post-secondary education. The Federal Pell Grant Program is the largest program to help low-income students attend college in the United States, with total awards exceeding 30 billion U.S. Dollars in 2023. To express the effect of this program on an MSA's economic growth, we estimate the program's "fiscal multiplier," the increase in income and economic activity for a

Figure 3.1: *The Pell Grant Program: Expenditures and Recipients*

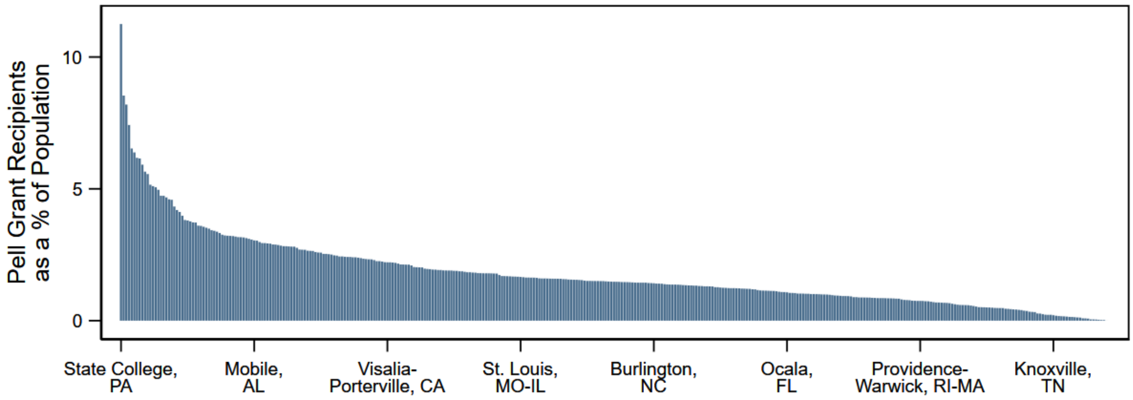

Notes: The figure plots the average fraction of a city's population that is eligible for Pell grants, in descending order. Data is obtained from the Title IV Program Volume Reports by the Department of Education.

given increase in the program's spending. Our estimate of the fiscal multiplier quantifies the effect of a *relative* increase in Pell grant disbursements on the *relative* increase of a city's aggregate income and employment. If the income from Pell grants acts as a substitute for other income, such as income from work, the multiplier of Pell grants will be between 0 and 1 for income and negative for employment. In contrast, if the multiplier of Pell grant enhances local economic activity, the multiplier will be above 1 for income and positive for employment.

We obtain a causal estimate of the multiplier of the Pell Grant Program using a shift-share instrumental variable approach that exploits cross-sectional variation in Pell-grant eligibility. Rather than directly relating increases in local Pell grants to local growth, we instrument local Pell grants by the interaction of national changes to Pell grant disbursements with the fraction of a city's population that received Pell grants *prior to* the change. This exploits the fact that there exists significant variation in the share of a population that is eligible for Pell grants (Figure 3.1). Our strategy is an example of the "shares" approach to shift-share instrumental variables, which enables a causal identification of local fiscal multipliers using cross-sectional variation—that is valid even if national changes in Pell grants are endogenous—as long as certain conditions are met. We employ using a series of validity tests proposed by Goldsmith-Pinkham et al. (2020), and show that our shift-share instrument passes these tests.

We find that Pell grants have a significantly positive effect on economic activity. Our main result is that the multiplier of Pell grants—the percentage increase in a city's relative income or employment from a relative increase in Pell grants by one percent of initial income—is 2.8 for our full sample on local income and 1.9 for employment. This means that a dollar spent on Pell grants creates (more than) twice as much relative economic activity. This estimate is robust to the inclusion of city and time fixed effects, city-size weighting, controls for spending by state governments, other main fiscal transfers to low-income households, and various other controls for the economic performance of a city. We find that schools increase expenditures when the Pell grant program becomes more generous, but that the main source of short-run economic gains is likely a rise in consumer spending. We find that, in part, the Pell grant fiscal multiplier operates through enabling students to attend college and acquiring students loans. An expansion in Pell grants increases student loan disbursements, which further relax students' budget constraints and enable these low-income students to spend more. This plausibly explains why the effect of Pell grants is larger than multipliers of other fiscal programs that do not directly unlock other sources of income.

Our estimate of the multiplier comes with uncertainty: our preferred estimate has a standard error of 1.5. This error is in line with multiplier estimates for other sources of spending, and our results for employment confirm that there is a significantly positive effect on economic activity when Pell grants increase. We do find that there is a high degree of variability in the magnitude of the multipliers, driven by the timing of the award and the types of schools that Pell grant awardees attend. When comparing the effect of Pell grants that are received by students at for-profit institutions to grants for students at non-profit institutions, we find that multipliers are lower at for-profit colleges. This appears to be because for-profit schools raise tuition fees in response to an increase in Pell grant generosity. It therefore appears that Pell grants are implicitly acting as subsidies for the for-profit university sector.[89] While we find that both for-profit and non-profit colleges raise some expenditures when Pell grants rise, we find no evidence that the majority of the positive effect of Pell grants operates through this channel. We also find that four-year institutions have significantly larger multipliers than two-year institutions. Finally, we assess whether the multiplier of Pell grants is higher during recessions, as has previously been found for military expenditure.[90] We find that Pell grants' effect on economic activity is larger in recessions than in expansions. While multipliers are only statistically significantly higher in recessions than expansions in the post 1999 sample, the point estimate is large and suggests that Pell grants can serve as a macroeconomic stabilizer during recessions.

Our estimates of the multiplier of Pell grants add to a vast literature that uses geographic cross-sectional variation in fiscal spending to estimate its short-run economic effects. The use of geographic variation became increasingly popular in the aftermath of the Great Recession. The advantage of using geographical cross-sectional data is that there is much greater variation in spending at the

[89] Note that since 2010 the "Gainful Employment" regulation has limited the Pell grants at certain for-profit colleges (see Cellini et al., 2016). In general, Turner (2017) estimates that 11-20 percent of Pell grants passes through to schools.

[90] See, e.g., Nakamura & Steinsson (2014), A. J. Auerbach & Gorodnichenko (2012), and Berge et al. (2021).

sub-national level, and more of this variation is plausibly exogenous. Like other sub-national estimations of multipliers, our results do have a particular interpretation: they measure the effect of Pell grants in one city on that city's relative economic performance, rather than the effect of Pell grants on economic performance at the national level. Local multipliers are estimated under different conditions than national multipliers. First, a city-level increase in Pell grants typically does not involve an increase in city-level fiscal deficits (and subsequent taxation), such that Pell grants do not crowd out private spending.[91] Second, a city-level increase in spending is typically sufficiently small such that it does not induce a monetary policy response. Third, local estimates may be affected by positive spillovers across local areas. A. Auerbach et al. (2020) note that local fiscal spending can have large spillovers to geographically close metropolitan areas, and find that these spillovers are typically positive.[92] Chodorow-Reich (2019) argues, on theoretical grounds, that the kind of geographical cross-sectional multiplier we estimate remains informative. It measures the national-level multiplier of fiscal spending when it is deficit financed and when monetary policy does not respond to the fiscal expansion, for example because interest rates are constrained by the effective lower bound. As these conditions often apply during recessions, our results give insight into the effectiveness of Pell grants as a tool to stimulate demand during downturns at the national level.

Our analysis of the multiplier for Pell grants takes into account changes in other local spending, such as unemployment insurance, housing benefits, food stamps, or state appropriations. Our main multiplier estimates thus evaluate the economic impact of an increase in Pell grants while keeping other fiscal spending constant. This approach allows us to factor out the possibility that other types of spending may automatically increase or decrease in response to changes in Pell grants, which could cloud our understanding of the direct effect of Pell grants. Our results indicate that these automatic responses are minimal; we observe similar multipliers even without controlling for other types of spending targeted to support low-income households. This evidence suggests that the positive effects of Pell grants on economic activity are neither diluted nor intensified by shifts in other fiscal spending.

The multiplier of 2.8 for Pell grants is higher than most estimates based on cross-sectional geographical variation of other forms of government spending. Early examples include Nakamura & Steinsson (2014), who estimate the state-level response of output to defense spending and find an average multiplier of 1.5. Acconcia et al. (2014) estimate multipliers from reductions in spending due

[91]Note that in a standard Neoclassical model, local multipliers are higher when they *are* locally financed, as the resultant tax increase would incentivize an increase in labor supply.

[92]The lack of local crowding-out is consistent with the idea that there is often excess capacity in production (Murphy 2017).

the expulsion of mafia-infiltrated city council members in Italy and find a multiplier of 1.9. Cross-sectional estimates of the multiplier were also frequently used to assess the effect of the American Recovery and Reinvestment Act (ARRA) (see, e.g., Chodorow-Reich et al. 2012, Chhabra et al. 2019, Conley & Dupor 2013, Dupor & Mehkari 2016, Feyrer & Sacerdote 2011) and the fiscal stimuli during the COVID-19 recession (see, e.g., A. Auerbach et al. 2022). Dupor & McCrory (2018), Suárez Serrato & Wingender (2016), and Hasna (2021) also study fiscal multipliers at a local level. Chodorow-Reich (2019) summarizes the literature on cross-sectional multipliers and finds that the mean estimated multiplier is 2.1 and the median estimated multiplier is 1.9. This evidence suggests that the multiplier for Pell grants is high compared to the multiplier of other forms of government spending and is, therefore, an effective tool to stimulate short-run economic activity. Furthermore, we also contribute to the literature on state dependence of fiscal multipliers by examining the differential effects across expansions and recessions (see, e.g., A. J. Auerbach & Gorodnichenko 2012, Berge et al. 2021, Barnichon et al. 2021, and Ghassibe & Zanetti 2022).

Why is the effect of Pell grants on local income so large? There are three likely drivers. First, Pell grants are a direct cash transfer and thus a part of personal income. This means that transfers themselves cause a one-for-one increase in personal income to begin with. Only if Pell grants "crowd out" other sources of private income, for example through a reduction in students' labor supply, would the multiplier ever be below one. Other forms of government spending, such as infrastructure investments or defence spending, only raise local income indirectly. As we find large effects of Pell grants on employment as well, however, this is unlikely to explain the full magnitude of Pell grants' multiplier. Second, Pell Grants are awarded to students from lower-income families. These students have limited borrowing capacity and are likely to have high marginal propensities to consume.[93] Fiscal stimuli that target households with high marginal propensities to consume generate higher multipliers (see, e.g., Johnson et al., 2006, Parker et al., 2013, and Jappelli & Pistaferri, 2014). Third, Pell grants enable many students to attend college (Dynarski, 2003), but are not enough to fully finance the cost of attending the college. These low-income students often rely on complementary financing sources, such as student loans. We find that large estimated multipliers can be reconciled, at least in part, with our finding that Pell grants increase cause an increase in student loans as well. These three factors may explain why Pell grants cause local income growth that significantly exceeds the initial cash transfer and produce fiscal multipliers that are higher than those estimated for other types of government spending.

In addition to providing evidence on the magnitude of the fiscal multiplier,

[93] For a literature review on the heterogeneity of marginal propensity to consume see, e.g., Jappelli & Pistaferri (2010).

this paper contributes to the literature on the Pell Grant Program. Previous work has documented several other positive effects, in particular in relation to education outcomes. Bettinger (2004) shows that receiving a Pell grant reduces college drop-out behavior. Pell grants also increase educational attainment, the probability of attending college, credit accumulation and has positive effects on students persistence and degree completion (Dynarski, 2003, Castleman & Long, 2016, and Fack & Grenet, 2015). Denning et al. (2019) show that eligibility for an additional Pell grant significantly increases the likelihood of degree receipt and raises earnings four years after the receipt of the degree. As higher earnings increase tax payments, they estimate that the government expenditures are fully repaid within 10 years. Dinerstein et al. (2014) do look at the short-term benefits of Pell grants as part of various federal transfers to post-secondary education during the Global Financial Crisis. Our analysis differs from Dinerstein et al.'s because we estimate the Pell grant's multiplier by exploiting the cross-sectional variation in the share of Pell grant recipients across cities. This means that our estimate of the multiplier is causal if these local exposure shares are exogenous—something we carefully assess in line with the recommendations by Goldsmith-Pinkham et al. (2020)—even if national changes in Pell grants are endogenous to national macroeconomic conditions. This new shift-share strategy also enables us to control for state and time-fixed effects. They find that counties which benefited from increases in the generosity of the Pell Grant Program did not have a significant increase in local income during the Crisis. They argue that one reason for this may be that students do not spend their grants in the immediate vicinity of their university. We conduct our analysis at the city (MSA) rather than the county level, which may be more appropriate given that consumers regularly travel outside their home county to consume.[94]

The remainder of this paper proceeds as follows. We begin by providing an overview of the Pell Grant Program in Section 3.2, in which we also explain our empirical approach. In Section 3.3 we discuss our main results, while in Section 3.4 we discuss how multipliers vary over the business cycle and compare multipliers at different types of colleges. Section 3.5 concludes.

3.2 Empirical Approach

This section outlines the empirical strategy to estimate the short-term economic effects of Pell grants. We start with a brief summary of the Pell grant program

[94]Using credit card data, Dunn & Gholizadeh (2023) show that consumers regularly consume outside their home county and that this consumption link across counties has important implications for economic measurement. In addition, Black et al. (2020) show that increased access to student loans increases college attainment and implies that these students do not have to rely on other sources of funding and do not have to work as much while in college.

and how grants are allocated to students in Section 3.2.1. Section 3.2.2 outlines the identification strategy while Section 3.2.3 summarizes the dataset.

3.2.1 Pell Grants: Background

The Federal Pell Grant Program was initiated in 1974 as the Basic Educational Opportunity Grant to provide a need-based grant to enable low-income students to attend college. It was renamed the Pell Grant Program after Senator Claiborne Pell in 1980. It started off as a program for 280 thousand students in 1974 with a total appropriation of $122 million, which increased to over 9 million recipients and a $30 billion appropriation by 2023. The program's size depends on the size of the cohort receiving Pell grants and on the maximum grant amount determined by the law. The program expanded particularly rapidly from the early 2000s to 2010. Since 2000, the U.S. has witnessed a substantial increase in enrollment at post-secondary institutions and a marked increase in college tuition, both reflected in the non-profit and the for-profit education sectors. Federal support for higher education was expanded in order to compensate for the increasing costs, as part of the College Cost Reduction and Access Act of 2007 and of the American Recovery and Reinvestment Act of 2009.[95]

The size of individual grants primarily depends on a student's family earnings. The largest share of Pell grant disbursements is typically received by students from families with an adjusted gross income of less than $60,000.[96] The grant amounts are conditional on the student's expected family contribution (EFC), the institutional cost of attendance, the student's enrollment status, and whether or not they attend a full academic year or less.[97] A full-time student is eligible for the following Pell grant award if the maximum Pell grant ($Pell^{MAX}$) is higher than the EFC:

$$Pell_{i,t} = max\left\{(Pell_t^{MAX} - EFC_{i,t}), Pell_t^{MIN}\right\}, \tag{3.1}$$

where $Pell^{MIN}$ is the minimum Pell grant.[98] Once the grant amount is determined, the institution at which the student is enrolled either credits the grant funds to the student's account, pays the student directly by check, or combines these methods. Grant recipients can enroll at various types of institutions, rang-

[95]A full summary of legislative changes is found in the Appendix.

[96]For example, 96.6 percent of Pell grant recipients in 2011-12 had an income of $65,995 or less (see Delisle, 2017).

[97]Financial need is determined by the Department of Education using a standard formula established by Congress to evaluate the to determine the EFC. The formula relies on the student's income (and assets for independent students), the parents' income and assets (for dependent students), the family's household size, and the number of family members (excluding parents) attending post-secondary education.

[98]Awards are rounded to the nearest $100. Part-time student awards are scaled by a factor of 0.5; scale factor is used for all determinants in eq. (3.1). Part-year students receive a prorated Pell grant.

ing from four-year colleges to those specialized in occupational training. Currently, about 5,000 post-secondary institutions participate in the program and more than 40 percent of all undergraduates are relying on this type of aid. A significant share of grant recipients are enrolled at public two-year schools and at for-profit institutions. Pell grants do not typically cover the entire cost of attendance and, as result most recipients supplement this type of aid with funds from other sources, such as federal and/or private student loans, personal savings, and 529 plan savings.

3.2.2 Strategy

Identification Problem We estimate the effect of Pell grants on short-run economic activity by exploiting variation in Pell grant disbursements across cities. By relying on regional variation, we enable causal identification of Pell grants' multiplier. At the national level, changes in Pell grants are highly endogenous to economic fluctuations. Enrollment in higher education is counter-cyclical, tending to increase when economic performance is poor, for example, causing an endogenously negative relationship between growth and the size of the Pell Grant Program. In fact, the 2009 increase in the level of individual Pell grants as part of the American Recovery and Reinvestment Act was expressly in response to poor economic performance during the Global Financial Crisis. This places a downward bias on national-level multiplier estimates.

We overcome this national-level limitation by analyzing the effect of an increase in generosity of the Pell Grant Program at the city (Metropolitan Statistical Area–MSA) level. While the generosity and conditionality of Pell grants are determined federally, there is significant variation in the extent to which subnational areas benefit from an increase in national-level Pell grant awards. This variation is driven by the fact that areas differ in the number of eligible students in post-secondary education. A city with a large number of universities benefits more from an increase than a city without universities, while a city where a small fraction of its student population is eligible (e.g., because of average income) benefits less than a city where a greater fraction is eligible, even if both cities have a similar number of students overall. Metropolitan areas are the appropriate level of analysis because a vast majority of U.S. college students resides locally where their school is located.[99] College students also tend to spend most of their in-

[99] According to the 2015 Digest of Education Statistics Table 309.10 covering student residence and migration, 82 percent of first-time degree-seeking undergraduate students attend college within their state of residence. Dunn & Gholizadeh (2023) show that consumers regularly consume outside their home county, so areas like MSAs or commuting zone may be more appropriate for the analysis of fiscal multipliers than counties. Additionally, there is more variation in spending across metropolitan areas than at other levels commonly used in the estimation of multipliers, like at the state level. Over the complete sample, the ratio of Pell grant spending to GDP is 0.16 percent across MSAs with a standard deviation of 0.17 percent, while that ratio is 0.12 percent at the state level with a standard deviation of just 0.08 percent.

come for basic household goods, such as groceries, housing, transportation, and health care. Usually, most spending for these categories occurs within the MSA the school is located in.

Shift-Share Instrument Approach We estimate the effect of an increase in Pell grants along:

$$\frac{\Delta y_{m,t;t-2}}{y_{m,t-2}} = \beta \cdot \frac{\Delta e_{m,t;t-2}}{y_{m,t-2}} + \phi_m + \psi_t + \gamma' x_{m,t-2} + \mu_{m,t}, \tag{3.2}$$

where β is the multiplier, $y_{m,t}$ is per-capita personal income in metropolitan area m in year t, while $\Delta y_{m,t;t-2}$ is its bi-annual change. $\Delta e_{m,t;t-2}$ is the bi-annual change in the per-capita transfer of Pell grants to students enrolled at schools in m. It follows that our estimates of the multiplier of the Federal Pell Grant Program, coefficient β in eq. (3.2), measures the *relative* increase in metropolitan area m's income when it achieves a relative increase in Pell grants of 1 percent of local income. $x_{m,t-2}$ is a vector of local control variables, while ϕ_m and ψ_t denote fixed effects for metropolitan areas and years, respectively. We measure economic activity through personal income, which is a measure that correlates highly with GDP.[100] We use biannual changes to mitigate the noise coming from the mismatch between calendar years and academic years, and to account for the fact that shocks to spending tend to precipitate in the second year. When estimating Pell grant's effect on employment, we replace the dependent variable by $\frac{\Delta L_{m,t;t-2}}{L_{m,t-2}}$ where $L_{m,t}$ is local employment.

To obtain a causal estimate of Pell grants' multiplier β, we must still address the possibility of endogeneity in changes to local Pell-grant awards. Increases in Pell-grant awards at the level of a metropolitan area may respond, for example, to an increase in local college enrollment that is driven by a deterioration of local economic conditions. This again puts a downward bias on the estimates of the multiplier.

We address this identification problem using shift-share instrumental variables (SSIV). We rely on the "shares-approach" identification strategy, as proposed by Goldsmith-Pinkham et al. (2020). In particular, we construct shift-share instrument $b_{m,t}$ that equals the interaction of the national growth in Pell grant disbursements and the fraction of a metropolitan area's population that received a grant two years prior:

$$b_{m,t} = \left(\frac{\Delta e_{t;t-2}}{y_{t-2}} \right) \cdot s_{m,t-2} \tag{3.3}$$

where $\Delta e_{t;t-2}$ is the bi-annual change in national Pell grants while y_{t-2} denotes

[100] MSA GDP is only available from 2001. The correlation between GDP and personal income is 0.997 in overlapping years.

twice-lagged national mean of personal income, and where $s_{m,t-2}$ denotes the share of a city's population that received a Pell grant two years prior. Our SSIV therefore leverages variation in the density of Pell grant recipients across cities prior to an increase in the program's generosity to identify the grants' effect on short-term growth.

Goldsmith-Pinkham et al. (2020) show that an SSIV along the one described in eq. (3.3) enables a causal identification of β as long as (i) the instrument $b_{m,t}$ is relevant and that (ii) the "shares" in the SSIV are orthogonal to the structural error term.[101] The first condition is straightforward to verify by regressing the instrument on changes in a city's Pell grant receipts. The second condition, which in our case requires that the shares $s_{m,t-2}$ are orthogonal to $\mu_{m,t}$ in eq. (3.2), cannot be verified directly. Instead, we conduct a series of falsification tests recommended by recent papers on the use of SSIVs (Goldsmith-Pinkham et al. 2020, Borusyak et al. 2022) to validate the strategy, and show that our SSIV passes these tests consistently.

3.2.3 Data

To estimate the short-run economic effects of the Pell Grant Program we analyze a sample of 367 metropolitan areas with data from 1990 to 2015.[102] Summary statistics are provided in Table 3.1. We obtain data on personal income from the Bureau of Economic Analysis (BEA) and our data on employment from Bureau of Labor Statistics (BLS). Our data on the Pell Grant Program comes from two sources: the Delta Cost Project—an independent, nonprofit organization, that provides estimates based on data from the Integrated Postsecondary Education Data Systems—and the Title IV Program Volume reports published by the Department of Education. These datasets provide information about Pell grant disbursements at the level of higher education institutions. We aggregate the data to the metropolitan area level, which we are able to do for around 87 percent of Pell grants.[103] We use Delta Cost as our primary source for Pell grant data because it covers the entire sample period, while data from the Department of Education is

[101]While we also expect the "shifts," i.e., the annual changes to the national amount spent with Pell grants, to be exogenous in this framework we do not adopt the "shifts-approach" in Borusyak et al. (2022) because we do not have enough time periods for a large-enough shock sample size.

[102]Our data starts from the universe of MSAs, from which we exclude areas that never receive Pell grants and MSA-years with bi-annual income changes of less than -5 percent. This assures that we do not include local "natural" disasters in the sample. If we were to include these natural disasters, the estimates of the multipliers would remain roughly the same, if anything they would increase a touch.

[103]The remainder of Pell grants is awarded to institutions in rural areas.

available only from 2000.[104] The latter also has data on the number of Pell grant recipients, which we use to construct the shift-share instrumental variable.[105]

Our control variables come from a variety of sources. From the Delta Cost and Department of Education datasets we obtain various characteristics of an MSA's higher education institutions. These include the number of undergraduate students enrolled and the average tuition fee they pay, information on the fraction of institutions that is for-profit, and whether institutions primarily offer two- or four-year degrees. We additionally obtain financial control variables from Equifax credit bureau data through the Federal Reserve Bank of New York's Consumer Credit Panel. We use this dataset to control for student and overall debt, median Equifax Risk Score, mortgage delinquency, and credit card utilization. Data is available for the post-1999 period at quarterly frequency, which we annualize by taking averages. Finally, we retrieve demographic control variables for race and average education levels from the Census Bureau.

We use Delta Cost to obtain data on state appropriations for higher education. As we aim to measure the fiscal multiplier effects of Pell grants, in some specifications we control for state appropriations as they may be substitutes or complements for Pell grants. This means that without controlling for appropriations, our estimated multipliers would not measure the effect holding other fiscal spending on higher education constant. For readers interested in understanding how Pell grants' multiplier may be amplified or dampened by changes in other spending, we also present results without this control.

We also use data on the main U.S. fiscal transfer programs for supporting low-income households to control for possible confounding shocks. Data on the Supplemental Nutrition Assistance Program (SNAP) comes from the USDA and is available at the county level for most states for a single month in the year (July). We then multiply SNAP disbursements by 12 to get yearly amounts. Housing assistance data to low-income households comes from the U.S. Department of Housing and Urban Development (HUD) which is available at the MSA level af-

[104] A small fraction of observations in Delta Cost is adjusted or imputed. To validate the Delta Cost data, we compare the MSA-level Delta Cost data with the available DoEd data aggregated at the same level. This comparison reveals 16 areas where Pell grants from Delta Cost differ erratically from the DoEd data, which we address in two ways on a case-by-case basis. First, for the cases when one year of data were missing or one MSA-year observation was considered suspicious, we used linear interpolation based on the Delta Cost data. Second, for the cases when multiple MSA-year observations were either missing or were questionable, we applied the growth rate observed in the DoEd data to Delta Cost data. From our sample of 367 MSAs we correct the path of Pell grants for 9 using interpolation and 10 using the DoEd growth rate.

[105] Given that Pell grant recipient data is not available between 1990 and 1999, we impute the Pell grant recipient share ($s_{m,t}$) for these years. To do so, we approximate the number of Pell grant recipients in an MSA by dividing an MSA's Pell grant disbursement by the maximum per-capita Pell grant that year, as most students receive the maximum amount. We then calculate $s_{m,t}$ as before. The correlation between the actual Pell grant recipient shares and the imputed shares is 0.925. To assure that our results are robust to not using the imputation, we always include specifications that only rely on data from 2000 to 2015.

Table 3.1: Summary Statistics

	Mean	St. Dev.	Obs	Min	Max	Source
Dependent Variable						
Δ Personal Income (Biannual)	0.036	0.038	8,436	-0.050	0.624	BEA
Δ Employment (Biannual)	0.007	0.028	8,436	-0.130	0.251	BLS
Pell Grants and SSIV						
Growth in Expenditure - MSA	0.016	0.067	8,436	-0.757	1.084	Delta Cost
Growth in Expenditure - National	0.012	0.031	8,436	-0.029	0.116	Delta Cost
Pell Recipients Share (% of Pop.)	1.743	1.691	8,436	0.000	22.681	DoE
Appropriations Share (% of Income)	1.005	1.701	8,436	0.000	15.210	Delta Cost
Loan Disbursements						
Growth in Loan Disbursements	0.062	0.221	4,447	-2.544	3.237	Delta Cost
Control Variables						
Growth in Appropriations	-0.012	0.237	8,436	-5.237	6.664	Delta Cost
Students (Log Change)	0.028	0.152	8,436	-4.222	4.352	Delta Cost
Tuition fee (Log)	8.606	0.835	8,436	4.745	10.768	Delta Cost
For Profit (%)	18.873	20.148	8,436	0.000	81.818	Delta Cost
Black (% of Population)	10.414	10.668	8,436	0.094	52.672	Census
Hispanic (% of Population)	9.661	14.142	8,436	0.284	95.745	Census
Bachelors Degree (% of Pop.)	9.342	2.646	8,436	3.200	21.018	Census
Credit Card Utilization Rate	26.744	6.033	4,447	8.004	65.020	CCP (post 1999)
Age(Median)	47.549	4.385	4,447	27.500	63.000	CCP (post 1999)
Risk Score (Median)	701.264	34.814	4,447	583.875	787.188	CCP (post 1999)
Mortgage Delinq. (%)	5.713	4.803	4,447	-7.601	64.438	CCP (post 1999)
Total Debt (% of Income)	4.002	3.331	4,447	0.043	42.245	CCP (post 1999)
Student Debt (% of Income)	0.278	0.329	4,447	0.001	4.812	CCP (post 1999)
SNAP Share (% of Income)	0.459	0.288	4,447	0.000	3.346	USDA
UI Share (% of Income)	0.406	0.383	4,447	0.030	7.444	DoL
HUD Share (% of Income)	0.200	0.104	4,447	0.001	0.799	HUD
Growth in Fiscal Transfers	0.000	0.004	4,447	-0.052	0.049	USDA, DoL, HUD, Delta Cost

Notes: Summary statistics for the merged sample. Data from 1990 to 2015 covering 376 metropolitan areas. CCP stands for Federal Reserve Bank of New York/Equifax Consumer Credit Panel. BEA stands for Bureau of Economic Analysis. USDA stands for US Department of Agriculture. DoL stands for US Department of Labor. HUD stands for US Department of Housing and Urban Development.

ter 2004. HUD data provide average government spending per unit per month which we multiply by the total number of units and by 12 to get yearly spending. Finally, unemployment insurance (UI) spending data comes from the Department of Labor. This data is only available at the state level. To get MSA-level estimates, we assume MSA disbursements are proportional to the MSA's share of unemployed calculated at the state level. For MSAs that cross state lines, we further assume that their population is split between states proportionally to their population levels. Finally, we replace missing data with the respective yearly sample average for each fiscal transfer programs.

3.3 Results

We now proceed with the main exercise. Section 3.3.1 presents the multiplier estimates for Pell grants along (3.2) and shows that our instrument is relevant. Section 3.3.2 performs a series of validity tests for the shift-share instrument as prescribed by Goldsmith-Pinkham et al. (2020) and also provides robustness results using alternative estimators. Finally, Section 3.3.3 assesses Pell grants' effect on student loans.

3.3.1 Multiplier Estimates

Results for the main estimation of the multiplier of Pell grants are presented in Table 3.2. Panel A of this table shows the effect of Pell grants on local income per capita and panel B shows the effect on employment. The estimated coefficient β from equation (3.2) represents the multiplier.[106] All estimations control for MSA-level and time fixed effects. In some regressions we additionally control for changes in state appropriations for higher education. State appropriations are included because they interact with Pell grants. During the 2009 recession, for example, state appropriations fell by 29 cents for every dollar increase in federal research funds (Dinerstein et al. 2014). Some states even reduce appropriations proportionally to increases in Pell grants. This has a negative effect on economic activity and not controlling for appropriations would therefore lead to an underestimation of the ceteris paribus effect of Pell grants on short-term growth.[107] As state appropriations might also suffer from endogeneity with respect to economy growth, we instrument it with a SSIV that is analogous to the one we use for Pell grants, i.e., we calculate the shift-share instrument using national changes in appropriations spending and the twice lagged appropriations share of local income. Finally, we cluster standard errors by MSA.

Our results show that the Pell Grant Program has an economically and statistically significant multiplier effect. For the full data from 1990 to 2015, in panel A we find a multiplier of 2.7 with our standard controls. This means that when an MSA receives an increase in Pell grants of 1 percent of income, local income increases by 2.7 percent compared to other MSAs. We then progress in steps to add additional controls: controlling for MSA controls, which include change in undergraduate students, average tuition and the percentage of schools that is for profit, demographic controls consisting of the percentage of the population that is black, Hispanic, and the fraction that at least has a Bachelor's degrees, the estimated multiplier rises to 2.8. This barely affects the estimated multiplier effect. We can also see that the effect of state appropriations is small, suggesting that our multiplier effects are only marginally offset by the nature of state funding for higher education.

Columns (5) to (10) of Table 3.2 are for the shorter post-1999 sample for which we have financial controls from the Federal Reserve Bank of New York/Equifax

[106] Table C.1 in the Appendix uses our baseline methodology to calculate the one-year multiplier of Pell grants to local income.

[107] Reductions in state appropriations tend to have negative effects on students. Webber (2017) shows that for every \$1,000 per student state budget cut, the average student pays \$257 more in tuition and fees. Webber (2017) also shows that this trend has increased over time. State appropriations for higher education are also shown to have an impact on enrollment and borrowing. Goodman & Volz (2019) find that changes in appropriations induce students to substitute between public and for-profit colleges and have corresponding effects on student borrowing.

Table 3.2: Effect of Pell Grants on Local Income Per Capita and Employment

	Full Sample				Post 1999						Full Sample
	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS	(9) 2SLS	(10) 2SLS	(11) OLS
Panel A: Income Growth											
Multiplier	2.735*	2.656*	2.910**	2.796*	3.640**	3.615**	3.058*	2.982*	2.967*	3.126*	-1.672*
	(1.441)	(1.486)	(1.442)	(1.486)	(1.691)	(1.725)	(1.672)	(1.650)	(1.679)	(1.638)	(0.925)
Panel B: Employment Growth											
Multiplier	1.620	1.835	1.764	1.930*	3.120**	3.243**	2.855**	2.634**	2.771**	2.660**	-1.559**
	(1.123)	(1.132)	(1.125)	(1.143)	(1.312)	(1.322)	(1.277)	(1.258)	(1.264)	(1.256)	(0.772)
Observations	8,436	8,436	8,436	8,436	4,447	4,447	4,447	4,447	4,447	4,447	8,436
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes	Yes	Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test	107.9	100.0	107.9	100.0	78.0	74.2	73.7	77.2	73.5	80.9	-
Joint F-test	-	67.3	-	67.7	-	37.8	37.5	-	37.4	45.7	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population (see eq. 3.3). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test is the robust F-statistic of the first-stage regression of Pell grants. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Consumer Credit Panel. In columns (5) to (7) we reproduce the specifications from columns (1), (2), and (4), respectively, for the shorter sample. The estimated multipliers are about 0.9 point higher in columns (5) and (6) compared to columns (1) and (2) and 0.3 point higher in column (7) compared to column (4). It follows that there remains a strong positive effect of Pell grant receipts on local economic growth in the shorter sample. Columns (8) to (10) add the Equifax controls, consisting of credit card utilization rates, average age, Equifax Risk Score, mortgage delinquency, and both total and student debt as a percentage of income. Adding these controls yields a multipliers between 2.9 and 3.1. Note that controlling for appropriations has very little, if any, effect in the shorter sample, as can be seen by comparing results in columns (8) and (9). In column (10) we reproduce the specification of column (9) using the broader definition of changes in fiscal spending instead of just appropriations. These transfers include the state appropriations as well as SNAP, HUD, and UI, and thus capture all the major fiscal transfers in the U.S. to low-income households. This yields a 0.15 point increase in the multiplier, which is well within a standard error. Importantly, our estimate of the Pell grant multiplier does not fall when including these fiscal spending controls. Thus, our high estimates of the economic effects of Pell grants are not driven by a positive correlation between Pell grants and other fiscal spending.

Our results suggest that multipliers of Pell grants are large. Our estimates exceed the median (1.9) and the average (2.1) of multipliers found in previous studies relying on geographic cross-sectional variation in other forms of fiscal spending, as surveyed by Chodorow-Reich (2019). Our estimates also exceed the 1.3 to 2.5 range for military spending in Nakamura & Steinsson (2014). Importantly, all of our estimates exceed unity. As Pell grants are transfers, they cause one-for-one increases in the BEA's measured personal income—hence our multiplier estimates show that local activity increases with disbursements.

In panel B of Table 3.2 we display estimates the effect of Pell grant disbursements on local employment. This is useful as it enables a comparison with a part of the literature that—among others for reasons of data availability—uses this dependant variable. Estimates of the fiscal multiplier on employment suggest that the employment fiscal multipliers are between 1.8 and 3.2, with the preferred estimate around 1.9 for the full sample. We can observe that—similarly as for the evidence in Panel A—multipliers increase somewhat when we move to the shorter sample. In column (7) where we reproduce the estimates from column (4) for a post-1999 sample, the fiscal multiplier on employment increases by about 0.9 points. When including the financial controls in columns (8)–(10) the effect of Pell grant disbursements on local employment ranges between 2.6 and 2.8. It is important to note that the fiscal multipliers do not change materially when we include controls for the appropriations and other fiscal transfers as is the case for the fiscal multipliers on local income. However, we can note that fiscal multipliers on employment are more precisely estimated than fiscal multipliers on local income as can be seen by substantially lower standard errors in panel B compared to panel A. Thus, the employment growth multipliers are highly significant in the post-1999 sample and suggest that Pell grant disbursements have large and significant effects on local employment.

To compare our employment growth multipliers with other similar multipliers estimated in the literature, the same specification for military expenditure at the state level in Nakamura & Steinsson (2014) gives an employment growth rate multiplier of 1.3. In addition, due to spillover effects (McCrory, 2020) the employment growth multipliers are generally increasing in the economic geography. Thus, one would expect smaller multipliers for MSAs than for state-level data. Our estimate implies that the cost of creating a job through Pell grants is around \$30,500.[108] For a comparison, A. Auerbach et al. (2022) studies the effect of fiscal stimuli during the COVID-19 recession and find that at a core-based statistical

[108] This number is found from equation (3.2) using the change in employment rates as the dependent variable. The effect of Pell grants on employment count is given by $\partial L_t/\partial E_t = \widehat{\beta} \cdot L_{t-2}/Y_{t-2}$. Inserting the inverse of average personal income per employee in the sample and $\widehat{\beta} = 1.93$ gives 0.328 jobs per \$10,000.

area the cost of creating a job was about $50,000. At the commuting-zone level, Dupor & McCrory (2018) find that it takes between $67,000 and $100,000 to create a job-year, while at the county level, Suárez Serrato & Wingender (2016) estimate that about $30,000 in federal spending creates a new job-year. Our multipliers are very similar to those estimated in Suárez Serrato & Wingender (2016).[109] In addition to the effects on local income, our results show that transfers to education can also have significant effects on employment growth.[110]

The bottom row of Table 3.2 presents Kleibergen-Paap F-statistics for all instrument relevance jointly, while the row above presents robust F-statistics of the first-stage regression for Pell grants. We calculate critical values for the F-statistic using the Olea & Pflueger (2013) test and find critical values of 23.1 for a 10 percent worst case bias, which is the usual threshold. This critical value is comfortably exceeded in all columns.

As a robustness check, Table C.2 in the Appendix presents results where we check whether city size matters for the estimated multiplier (Shoag, 2013). Our weighting relies on the two-year lagged logarithm of MSA population. We find that the estimated multipliers are very similar to those in Table 3.2.

3.3.2 Shift-Share Validity Tests

We next show that our shift-share instrumental variable passes a series of validity tests. Our strategy uses the "shares approach" (Goldsmith-Pinkham et al., 2020), which means we assume that the share of Pell recipients in the MSA's population is exogenous to the error term in the second stage of the regression analysis. That is, we assume that MSAs with a large share of Pell recipients would have seen a similar level of economic growth as MSAs with low shares absent changes in aggregate Pell grant disbursements.

As with all instruments, this exclusion restriction cannot be tested directly. Instead, we follow the recommendation in Goldsmith-Pinkham et al. (2020) to conduct a series of validation tests. In Section 3.3.2 we perform a balance test, which involves regressing both the Pell grant recipient population share and the second stage's dependent variable on the covariates. In Section 3.3.2 we assess whether the endogenous dynamic bias identified in Jaeger et al. (2018) poses a threat to our results. In Section 3.3.2 we analyze which years in our sample are most important for our shift-share estimate of the fiscal multiplier. In Section 3.3.2 we use the Pell grant population shares directly as instruments. Finally, in Section 3.3.2 we assess whether our estimates are robust to using alternative IV estimators. We show that our instrument passes each of these checks. As our sample starts several years after Pell grant program was introduced, testing for pre-trends is not

[109] Chodorow-Reich (2019) surveys other estimates of local fiscal multipliers.

[110] Previously, Feyrer & Sacerdote (2011) argued that transfers for education have modest effects on employment.

possible without additional assumptions.

Balance Tests

The first validity test for our shift-share instrument is a balance test. The test consists of two regressions, one with either growth in personal income or employment growth as the dependent variable, and one with the instrument shares—Pell grant recipients as a percentage of the population—as the dependent variable. In each regression, all of the control variables of the main analysis are the regressors. The idea behind this exercise is to assess whether we observe any simultaneous correlation between *observables*, the recipient shares, and income growth. If so, we might worry about omitted variable bias from *unobservables*.[111] The robustness of our point estimates in Table 3.2 to the inclusion of various alternative control variables suggests that such simultaneous correlation is unlikely; in this section we present a formal test that confirms that such simultaneous correlation is absent.

We report the results of the balance test in Table 3.3. As we use two time samples, we present results for the balance test separately for 1990–2015, columns (1) and (3), and for post 1999, columns (5) and (7). We include the largest possible set of observables (control variables) for the time sample as we do in the preferred specifications (4) and (9) in Table 3.2. To make coefficients easier to interpret, all variables are demeaned and normalized to have unit variance. The balance test falsifies the instrumental variable if one of the outcome variables—growth in personal income or employment growth—and the shares variable—Pell grant recipients as a percentage of the population—significantly correlate with an observable. If there are no simultaneous significant correlations for shares and either dependent variables, the balance test passes.[112] We mark these cases with a "✓" in columns (4) and (8). As shown in Table 3.3, all of our observable control variables pass the balance test. This makes it less likely that an unobserved confounder would correlate with both the income growth (or employment growth) and the share of the population that ex-ante receives Pell grants. We show in Table C.3 in the Appendix that our SSIV also passes a balance test for the specification in column (10) of Table 3.2.[113]

We further assess the validity of the appropriations and fiscal transfer SSIVs by running their respective balance tests. These are reported in Tables C.4 and C.5 in the Appendix, respectively. The balance test for appropriations is a pass for all but three control variables, the local share of Hispanics and people with at least a bachelor degree, and the share of for-profit colleges. While this is potentially

[111] A similar argument is put forward in Oster (2019).

[112] Naturally, the test passes if we have significant correlations for both dependent variables so long as it is not significant for shares.

[113] We also run a balance test adding real college spending per capita as an additional observable and verify that our identification strategy still passes the test. Results are available upon request.

Table 3.3: Balance Test

	Full Sample				Post 1999			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Shares	Income Growth	Empl. Growth	Pass	Shares	Income Growth	Empl. Growth	Pass
Approp. Growth	0.003	0.032***	0.051***	✓	0.001	0.033***	0.047***	✓
	(0.014)	(0.010)	(0.014)		(0.008)	(0.011)	(0.013)	
Log(Tuition)	-0.006	-0.070	-0.023	✓	-0.022	-0.240***	-0.124*	✓
	(0.048)	(0.051)	(0.050)		(0.049)	(0.062)	(0.064)	
D.Log(Students)	0.017*	0.001	-0.006	✓	0.021*	0.001	-0.009	✓
	(0.010)	(0.009)	(0.009)		(0.011)	(0.011)	(0.007)	
For Profit	0.052	-0.003	0.025	✓	0.006	0.031	0.100**	✓
	(0.035)	(0.040)	(0.028)		(0.028)	(0.057)	(0.047)	
Share Black	0.157	-0.028	0.241	✓	0.594	0.044	0.423	✓
	(0.233)	(0.159)	(0.157)		(0.388)	(0.521)	(0.357)	
Share Hisp.	-0.172	0.548***	0.763***	✓	-0.287	0.358	0.799***	✓
	(0.141)	(0.144)	(0.117)		(0.223)	(0.362)	(0.279)	
Share Bach.	-0.116	0.074	0.273***	✓	0.028	0.664***	0.792***	✓
	(0.133)	(0.100)	(0.080)		(0.135)	(0.163)	(0.145)	
Risk Score					-0.041	-0.024	-0.094**	✓
					(0.029)	(0.042)	(0.048)	
Age					0.002	-0.013	0.006	✓
					(0.037)	(0.041)	(0.038)	
Debt to Income					-0.052	0.238	-0.082*	✓
					(0.032)	(0.191)	(0.045)	
Card Util.					0.026*	-0.039	-0.017	✓
					(0.016)	(0.030)	(0.024)	
Mort. Delinq.					0.013	-0.009	-0.022	✓
					(0.015)	(0.019)	(0.019)	
Observations	8,436	8,436	8,436		4,447	4,447	4,447	
R-square	0.7	0.3	0.5		0.8	0.4	0.6	
Time FE	Yes	Yes	Yes		Yes	Yes	Yes	
MSA FE	Yes	Yes	Yes		Yes	Yes	Yes	

Notes: Independent variables are twice lagged in columns (2), (3), (6), and (7) except for appropriation growth. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

worrying, results in Table 3.2 for the specifications without state appropriations (columns (3) and (8)) show that our results are not significantly affected by removing appropriations altogether. As for the combined fiscal transfers variable, the balance test fails for the local share of Hispanics and people with bachelor degrees, and for the debt to income ratio. This result, however, seems to be affected by the missing data imputation procedure that we describe in Section 3.2.3. Indeed, we do not observe a simultaneous significant correlation for these control variables and income growth when we restrict the sample to observations with complete data availability (see Table C.5 in the Appendix). As for the specification with employment, the failure of the balance test for the debt to income ratio is small in magnitude: a standard deviation change in the observable only affects shares by 4.1% of their standard deviation. Moreover, the estimate for the specification in column (10) in Table 3.2 does not deviate from the other estimates significantly, which is evidence that any potential underlying bias is not qualitatively affecting results.

We next assess the balance of the geographical distribution of the population share of Pell grant recipients. It is well established that local income growth and employment growth varies strongly across U.S. regions, with poorer areas growing faster than richer areas (see, e.g., Barro et al., 1991). If Pell grant recipient

Figure 3.2: *Geographical Distribution of the Ratio of Pell Grant Recipients over Population*

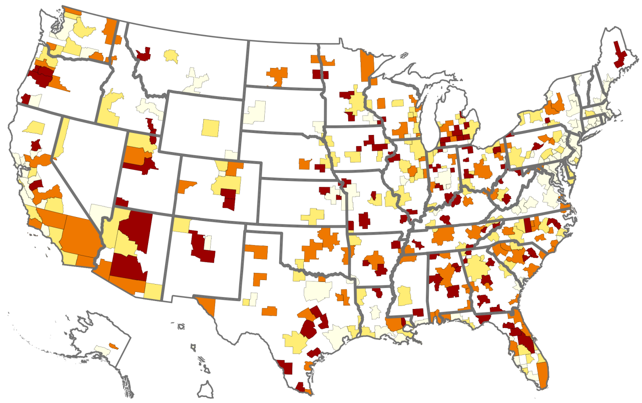

Notes: The figure plots $s_{m,2010}$ from equation (3.3). Blank areas fall outside metropolitan areas or, in rare cases, are areas that never receive Pell grants.

shares are persistently higher in areas with low or high income growth (employment growth), that would be an example of an observable variable that correlates with both our shares and dependent variable, violating the SSIV balance test proposed by Goldsmith-Pinkham et al. (2020). As in informal investigation, Figure 3.2 plots the share of Pell grant recipients in an MSA's population for 2010.[114] Dark areas have a larger population share of Pell recipients and are more sensitive to national changes in the amount disbursed, while light areas have a lower share. The figure shows that the distribution of Pell recipients does not exhibit geographical clustering. Most states contain both MSAs with above- and below-average Pell grant recipient shares, in support of the requirement that there is a balanced distribution of our instrument across regions. Insofar as geographical clustering makes a correlation between our instrument shares and dependent variables more likely, Figure 3.2 thus provides an additional successful balance test.

Bias from Persistent Effects of Pell Grants

Next we study the potential identification issues arising from endogenous dynamic responses to SSIVs. Jaeger et al. (2018) shows that our estimates in Table 3.2 may be biased if two conditions are met. First, the SSIV is serially correlated. Second, the multiplier effects induced by the Pell disbursements affect future local GDP, as the economy takes some time to adjust to changes in the Pell grants. The latter can happen if, for example, Pell grants boost local demand which prompts

[114]Results for other years, available on request, are similar to those for 2010.

firms to invest in capital, an effect that might take longer than a year to show up. If these conditions are met, our estimates capture not only the short-term multiplier effect of Pell grants but also their long-term, persistent effects, confounding the identification approach in the paper. The solution proposed in Jaeger et al. (2018) entails controlling for one-year lagged Pell grants and instrument them with one-year lagged SSIV. Table 3.4 reports results for both income and employment growth. As the coefficient on this lagged regressor is not significant for any of our specifications and the coefficient on current Pell grants is statistically indistinguishable from those in Table 3.2, we find no evidence of a significant dynamic bias in our coefficients.[115]

Rotemberg Weights

In a third validation exercise of our shift-share instrument, we analyze whether particular years drive the identification of our multiplier estimates. To do so, we calculate the Rotemberg weights, which tell us how sensitive the overall estimate of the fiscal multiplier is to endogeneity (misspecification) in the instrument. Goldsmith-Pinkham et al. (2020) show that the shift-share IV estimates can be written as a weighted sum of a GMM estimates using the recipient shares interacted with time fixed-effects as separate instruments. Weights in this decomposition are called "Rotemberg weights" and they sum up to 1. These weights depend on the covariance between the specific instrument's fitted value of the endogenous variable and the endogenous variable itself.[116] In our case, the SSIV estimate can be interpreted as a weighted sum of just-identified estimates calculated using the MSA-level Pell recipients population share for each year ($s_{m,t-2}$) separately.

Rotemberg weights indicate which years are most important for both our estimate and the identification strategy. If the instrument is misspecified (endogenous) in a specific year and that year has a high Rotemberg weight, our estimate of the multiplier could be significantly biased. Ideally, years with a high Rotemberg weight and a high first stage F-statistics should be close to the overall estimate of the fiscal multiplier. We plot the Rotemberg weights in Figure 3.3, which shows the just-identified yearly estimates of the multiplier $\hat{\beta}_t$ with respect to the first-stage F-statistics calculated using the specification in column (4) of Table 3.2.

[115] In this paper, we estimate two-year multipliers to avoid any potential issues regarding short-run dynamics of key variables of interest. To further demonstrate the robustness of our results, we also provide estimates of fiscal multipliers lagging the SSIV shares one extra period—three periods—to show that (any remaining) serial correlation of local income (employment) growth and/or correlation between the (twice-lagged) SSIV shares and past growth does not influence our results. Indeed, we find that in Table C.6 in the Appendix that our results are qualitatively the same.

[116] In the simplest example where the individual, yearly instruments are all orthogonal, the weights are simply the ratio between the just-identified first-stage R^2 and the full SSIV first-stage R^2.

Table 3.4: Effect of Pell Grants on Local Income Per Capita Controlling for Lagged Pell Grants

	Full Sample				Post 1999						Full Sample
	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS	(9) 2SLS	(10) 2SLS	(11) OLS
Panel A: Income Growth											
Multiplier$_t$	4.125**	4.149**	3.824**	3.851**	4.813**	4.849**	4.194**	4.198**	4.233**	4.059**	-1.922**
	(1.913)	(1.910)	(1.886)	(1.883)	(2.195)	(2.190)	(2.078)	(2.066)	(2.059)	(2.057)	(0.920)
Multiplier$_{t-1}$	-1.789	-1.967	-1.180	-1.393	-1.514	-1.637	-1.504	-1.572	-1.680	-1.215	0.460
	(1.819)	(1.847)	(1.763)	(1.788)	(1.992)	(2.046)	(1.968)	(1.953)	(2.009)	(1.992)	(0.859)
Panel B: Employment Growth											
Multiplier$_t$	2.448*	2.407**	2.086*	2.055*	3.201**	3.124**	2.500*	2.544*	2.444*	2.516*	-2.586***
	(1.262)	(1.207)	(1.226)	(1.181)	(1.496)	(1.429)	(1.331)	(1.358)	(1.283)	(1.334)	(0.912)
Multiplier$_{t-1}$	-1.065	-0.753	-0.415	-0.166	-0.105	0.158	0.471	0.117	0.434	0.188	1.883***
	(1.166)	(1.115)	(1.119)	(1.085)	(1.256)	(1.268)	(1.222)	(1.210)	(1.235)	(1.219)	(0.618)
Observations	8,436	8,436	8,436	8,436	4,447	4,447	4,447	4,447	4,447	4,447	8,436
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test	74.8	76.3	74.7	76.2	52.5	54.7	53.6	51.6	54.1	53.2	-
Joint F-test	35.3	35.8	35.3	35.9	24.5	21.2	20.9	-	20.8	32.5	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population (see eq. 3.3). SSIV strategy for the one-year lagged Pell grants regressor uses the three-times-lagged share of recipients in MSA population (see eq. 3.3). SSIV strategy for appropriations uses the two-year lagged appropriation share of income. Controls are one-year lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test is the robust F-statistic of the first-stage regression of Pell grants. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

The size of each point is proportional to the magnitude of each weight and the horizontal dashed line represents the baseline estimate in column (4) of Table 3.2 for both income and employment growth. As we can see in both panels of Figure 3.3, the high-weight years (2011 and 2010) are close to the baseline estimate and have a high F-statistic. This suggests that our estimate of the fiscal multiplier is not noticeably biased due to endogeneity of the instruments in a particular year: Even if there is some misspecification in a particular year this does not significantly affect the overall estimate. Additionally, while we do observe some dispersion between the $\hat{\beta}_t$ most of the outliers carry a low weight and a relatively low F-statistic. The large weights in 2011 and 2010 are expected given the large increases in Pell disbursements in those years (c.f. Section C.1 of the appendix). Moreover, while there are points with negative weights, these account for little in the overall weight.

Figure 3.3: *Multiplier Estimates from Single-Year Instruments and Rotemberg Weights*

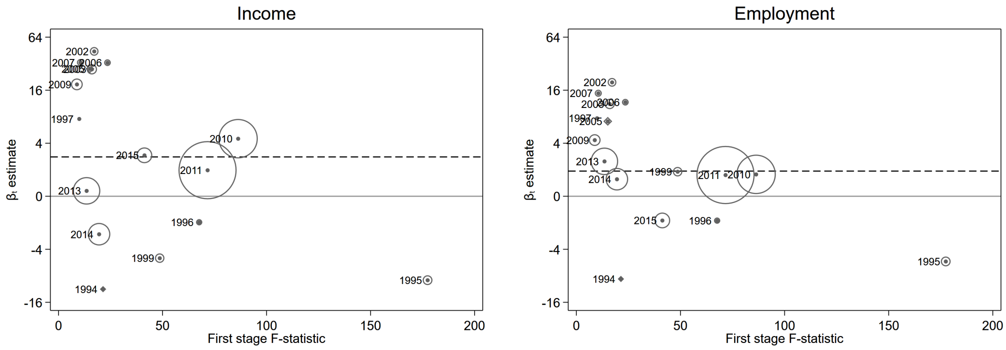

Notes: The figure plots each $\hat{\beta}_t$ as a function of their first-stage F-statistics for the full sample. The $\hat{\beta}_t$ are calculated using the specification of column (4) in Table 3.2 using both income and employment growth rates. The size of the circles and diamonds is scaled by the magnitude of the respective Rotemberg weight. Circles denote positive weights while diamonds denote negative weights. The horizontal dashed line shows the overall $\hat{\beta}_{SSIV}$ of column (4) in Table 3.2. The figure excludes instruments with first-stage F-statistics below 5.

Analysis using Shares Directly

In a fourth validity check, we replace the instrument in eq. (3.3) with the interaction of our shares variable—the lagged share of Pell-grant receivers as a fraction of the population—and time fixed effects. The idea behind this validity check, as proposed by Goldsmith-Pinkham et al. (2020), is to assure that our estimates of the multiplier of Pell grants are driven by variation in the shares variable, and not primarily by the shocks to the national Pell grants program. Estimates of the multiplier with the alternative instrument should therefore be similar to the main estimates. This check is important, as changes in the program are partially endogenous to economic growth. Indeed, the main expansions of the program that we discussed in previous sections were explicitly in response to the Great Recession.

To implement the "shares-directly" instrumental variable specification, we run:

$$\frac{\Delta e_{m,t;t-2}}{y_{m,t-2}} = \sum_{t \in T} \alpha_t \cdot s_{m,t-2} + \phi_m + \psi_t + \gamma' x_{m,t-2} + \mu_{m,t}, \tag{3.4}$$

where T denotes the total number of years in our sample. The second stage of the regression is unchanged. By interacting $s_{m,t-2}$ with time fixed effects rather than changes in the size of the Pell grant program at the national level, we reweigh the importance of years. With the new instrumental variables, the analysis now derives the estimated multipliers from an unweighted average of the points in Figure 3.3.

The multiplier estimates from the shares-directly approach are presented in Table 3.5. Results for the full sample are in columns (1)–(4), while columns (5)–(10) report the results for the post-1999 sample. The result in column (1) is a touch

Table 3.5: Effect of Pell Grants on Local Income Per Capita (using shares directly)

	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
Panel A: Income Growth											
Multiplier	1.762	1.890	2.144*	2.103	4.418**	4.351**	3.811**	3.854**	3.769**	3.880**	-1.672*
	(1.308)	(1.348)	(1.286)	(1.338)	(1.726)	(1.759)	(1.699)	(1.670)	(1.698)	(1.664)	(0.925)
Panel B: Employment Growth											
Multiplier	0.325	1.022	0.628	1.144	2.707**	2.888**	2.465**	2.207*	2.354**	2.227*	-1.559**
	(0.978)	(0.986)	(0.960)	(0.988)	(1.236)	(1.260)	(1.220)	(1.182)	(1.194)	(1.187)	(0.772)
Observations	8,436	8,436	8,436	8,436	4,447	4,447	4,447	4,447	4,447	4,447	8,436
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes	Yes	Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test	107.6	72.3	107.0	71.3	12.8	13.6	13.3	12.8	13.1	14.0	-
Joint F-test	-	17.8	-	17.9	-	8.3	8.4	-	8.5	11.9	-
AR Test p-value, Income	0.00	0.00	0.00	0.00	0.02	0.02	0.04	0.02	0.03	0.01	-
J-test, p-value, Income	0.00	0.00	0.00	0.00	0.02	0.02	0.03	0.01	0.01	0.01	-
AR Test p-value, Employment	0.00	0.00	0.00	0.00	0.01	0.01	0.04	0.05	0.05	0.06	-
J-test, p-value, Employment	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population directly as instruments. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test is the robust F-statistic of the first-stage regression of Pell grants. Joint F-test is the robust F-statistic of the joint IV set. AR Test p-value shows the p-value for the Anderson-Rubin weak-IV test. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

smaller than the companion in Table 3.2, though not significantly different from zero. The multipliers in columns (3) and (4) increase slightly to about 2.1, but still remain below their counterparts in Table 3.2, potentially because the proxy variable for the Pell recipients that we use in the first part of the full sample is constructed using MSA-year variation. As such, the fixed effects absorb some of variation in the proxy. When we consider the post-1999 sample, where we do not have to rely on a proxy for the Pell recipients, the estimates in columns (5) to (10) are slightly higher in panel A and slightly lower in panel B, but overall close to the corresponding estimates in Table 3.2.

In the final row of Table 3.5, we report the p-value of an over-identification J test. The shares-directly regression enables such a test because the regression now relies on T instruments rather than a single instrument. The J-test rejects that all of our instruments give rise to similar estimates of the multiplier—as was already evident from Figure 3.3. In some settings this may be worrisome, because it means that sub-samples of the data imply heterogeneous multipliers. In

Table 3.6: Effect of Pell Grants on Local Income Per Capita Using Different Estimators

Post 1999	Income Growth			Over ID test	Employment Growth			Over ID test
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
2SLS (SSIV)	3.615**	3.058*	2.967*	.	3.243**	2.855**	2.771**	.
	(1.725)	(1.672)	(1.679)	.	(1.322)	(1.277)	(1.264)	.
2SLS	4.351**	3.924**	3.769**	26.758	2.888**	2.312**	2.354*	29.588
	(1.759)	(1.671)	(1.698)	[0.013]	(1.260)	(1.203)	(1.194)	[0.005]
LIML	4.454**	4.024**	3.856**	26.777	3.121**	2.481**	2.510**	29.543
	(1.813)	(1.716)	(1.746)	[0.013]	(1.350)	(1.261)	(1.255)	[0.005]
HFUL	4.778***	4.237***	4.107***	156.115	3.316***	2.438***	2.620**	237.354
	(1.494)	(1.476)	(1.499)	[0.000]	(1.031)	(0.965)	(0.980)	[0.000]
Observations	4,447	4,447	4,447	4,447	4,447	4,447	4,447	4,447
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA Controls		Yes	Yes	Yes		Yes	Yes	Yes
Financial Controls			Yes	Yes			Yes	Yes

Notes: 2SLS uses each yearly share as a separate IV. LIML uses the limited information maximum likelihood estimation with the same set of instruments. Finally, HFUL uses the estimator from Hausman et al. (2012) also with the same set of instruments. Controls are contemporaneous to the respective timing of shares. Overidentification tests in column (5) refer to the specification in column (4). We use the Sargan test (Sargan, 1958) for the 2SLS and LIML estimators, and the overidentification test from Chao et al. (2014) for the HFUL estimator. P-values for the overidentification tests are in brackets. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

our case, however, heterogeneity in multipliers across different years is expected. There is a wide literature, for example, that suggests that multipliers are different across various stages of the business cycle (see, e.g., Berge et al., 2021). We study this heterogeneity in Section 3.4.1 and confirm that also Pell grants multipliers differ substantially across different business cycle stages. Therefore, it is not a surprise that we reject the null hypothesis of the overidentification test.

Alternative Instrumental Variable Estimators

In a fifth and final validity check, we replace the standard two-stage least squares (2SLS) estimator with two alternatives: the Limited Information Maximum Likelihood (LIML) and the heteroskedasticity-robust Fuller (1977) (HFUL) estimator proposed by Hausman et al. (2012). As Goldsmith-Pinkham et al. (2020) note, estimates from the LIML and HFUL estimator may differ from the 2SLS estimator because they rely on different identification assumptions. The LIML estimator reduces the small-sample bias from weak instruments though it is inconsistent under heteroskedasticity. The HFUL estimator, on the other hand, is consistent under heteroskedasticity and many instruments though it can be slightly more biased than LIML under homoskedasticity. Given these different properties, if point-estimates are similar under these approaches it is less likely that our strategy is misspecified.

Table 3.6 presents the results. The alternative estimators all use the separate instruments for every year in the data, which means they are most comparable to our shares-directly estimates. As these shares are only available for the post-1999 sample, we focus on those specifications. The first two rows in the table reproduce the baseline SSIV estimator (from Table 3.2). The third and fourth row present the fiscal multiplier estimates using the Limited Information Maximum Likelihood (LIML) estimator and the HFUL estimator from Hausman et al. (2012). Columns progressively add controls as in the previous tables.[117] Comparing the estimates, it is clear that the alternative estimators imply similar multipliers of the Pell Grant Program. All show similarly large positive economic effects of the program, and the difference between specifications within columns is not statistically significant.

3.3.3 Understanding the Pell Grant Multiplier

The preceding three sections show that the multiplier of the Pell Grant Program is large and causally estimated. In the remainder of this section, we posit a hypothesis for the driver of the multiplier: Pell grants enable students to raise consumption, as they both increase students' income directly and give them access to student loans.[118] According to the National Postsecondary Student Aid study, in 2015-16 school year 56% of Pell grants recipients supplemented Pell grants with student loans. If Pell grants enable students to attend college, they therefore also enable recipients to acquire student loans, further easing their budget constraint.

We test this assertion with two additional analyses. We first look at the effect of an expansion of Pell grants on the disbursement of student loans. Data on student loans is available through Delta Cost at the school level for the post-1999 sample. We aggregate the school-level data to the city level to estimate a regression akin to eq. (3.2) using the ratio of changes in student loan disbursement over lagged personal income as the dependent variable. Results are presented in the first three columns of Table 3.7. The point estimates average around 2.0, which means that an increase in Pell grants causes an increase in student loan disbursements that is around twice the size of the initial Pell grants rise. This means that, besides directly raising income, Pell grants indeed enable students to increase their borrowing.

In the second analysis, we assess whether the combined increase in Pell grants and student loans is sufficiently large to explain the overall increase in local activity after exogenous increases in Pell grants. To do so, we re-estimate the original regression in eq. (3.2) using the change in the sum of Pell grants and student loans as the independent variable, leaving the SSIV unchanged. If the increase in local

[117] Table C.7 in the Appendix reports the specifications without controlling for appropriations.

[118] Student loan disbursements and number of students receiving student loans are for first-time degree/certificate-seeking undergraduate students who received student loans. Loans to students are defined as any monies that must be repaid to the lending institution for which the student is the designated borrower. They include all Title IV loans and all institutionally- and privately-sponsored loans but do not include PLUS and other loans made directly to parents.

Table 3.7: Pell Grants and Student Loans

Post 1999 Sample	Student Loan Growth			Income Growth			Employment Growth		
	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS	(9) 2SLS
Δ Pell Grants	1.990***	1.978***	1.977***						
	(0.456)	(0.457)	(0.453)						
Δ (Pell Grants + Loans)				1.209**	1.027*	0.997*	1.085***	0.959**	0.931**
				(0.549)	(0.541)	(0.549)	(0.393)	(0.393)	(0.390)
Observations	4,447	4,447	4,447	4,447	4,447	4,447	4,447	4,447	4,447
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA Controls		Yes	Yes		Yes	Yes		Yes	Yes
Financial Controls			Yes			Yes			Yes
Fiscal Transfers									
Δ Pell Grants F-test	74.2	73.7	73.5	71.0	70.6	71.0	71.0	70.6	71.0
Joint F-test	37.8	37.5	37.4	30.8	30.4	30.5	30.8	30.4	30.5

Notes: SSIV strategy for the Pell grants regressor and the sum of Pell grants and loan disbursements uses the twice-lagged share of recipients in MSA population (see eq. 3.3). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test is the robust F-statistic of the first-stage regression of Pell grants. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

economic activity is proportional to the increase in the sum of Pell grants and student loans, we expect an estimated β of exactly 1 in this regression. The final six columns in Table 3.7 present the estimates for both local income growth and employment growth, which are indeed all around 1.[119] In the absence of individual consumption data of Pell grant recipients, this is the closest test of consumption hypothesis we are able to provide. Note that the lower estimates for β in Table 3.7 compared to Table 3.2 above do not mean that we overestimate the fiscal multiplier of Pell grants, as student loans are not a fiscal transfer. The fiscal multiplier is the ratio of the change in personal income and the change in the Pell grants, and is thus given by the estimates in Table 3.2.

It may come as a surprise that Pell grants and student loans are complementary, as evidenced by the results in the first three columns of Table 3.7. In practice, all students who may be eligible for any type of student aid apply through a single application known as the Free Application for Federal Student Aid or the FAFSA form. Then colleges, considering the student aid disbursement criteria discussed in section 3.2.1, send out financial aid offers that include the financing aid package to cover the difference between the cost of attendance and the expected family contribution (EFC). Lower-income students often rely both on Pell grants and different types of student loans, as they are often constrained in

[119] Table C.8 in the Appendix reports the specifications without controlling for appropriations.

their ability to raise EFC. Thus, it is not surprising that Marx & Turner (2018) find that Pell grants reduce potential borrowing (student loans), suggesting that Pell grants and student loans act as *substitutes* for students who have to meet the cost of attendance threshold. This can be reconciled with our results, however, by the fact that we measure the relationship between Pell grants and student loans at the city level. As long as Pell grants enable students to attend college when they otherwise would not have, a positive correlation between changes in Pell grants and student loans should arise.

3.4 When Are Pell Grants Most Effective?

We next assess under what conditions the effect of an increase in Pell grant disbursements on local economic activity is the largest. To do so, we look at how the multiplier effect of Pell grants varies across recessions and expansions, and whether the effect of grants depends on the type of institution that students attend.

3.4.1 Multipliers in Recessions and Expansions

We first compare the multiplier of Pell grants during episodes when the economy is in expansion to when it is in recession. Recent evidence suggests that fiscal spending generally has a greater effect on output when the economy is in recession.[120] If this holds for Pell grants, they could form a particularly effective tool to stabilize macroeconomic activity. We estimate the following equation to test this:

$$\begin{aligned}\frac{\Delta y_{m,t;t-2}}{y_{m,t-2}} &= F(z_{m,t-2})\left[\alpha_E + \beta_E \frac{\Delta e_{m,t;t-2}}{y_{m,t-2}}\right] + \left[1 - F(z_{m,t-2})\right]\left[\alpha_R + \beta_R \frac{\Delta e_{m,t;t-2}}{y_{m,t-2}}\right] \\ &\quad + \phi_m + \psi_t + \gamma' x_{m,t-2} + \mu_{m,t}, \qquad (3.5)\end{aligned}$$

where β_R and β_E respectively capture the multiplier in recessions and expansions, while $F(z_{m,t-2})$ is a continuous function that strictly increases with a moving average of lagged biannual growth, or employment for the specification on employment growth, $z_{m,t-2}$.[121]

This equation is also known as a smooth transition model, which we borrow from the literature on the state-dependent effect of fiscal and monetary policy on economic activity.[122] The specification assigns weights to observations based on

[120] Examples include A. J. Auerbach & Gorodnichenko (2012), Corsetti et al. (2012), Ilzetzki et al. (2013), Blanchard & Leigh (2013), Jordà & Taylor (2016), and Berge et al. (2021). Ghassibe & Zanetti (2022) further stress that the source of fluctuations matter as well. Ramey & Zubairy (2018) do not find state-dependence in a historical sample with news shocks about defense spending.

[121] Specifically, $z_{m,t-2}$ is defined as the moving average of two-year lagged local income growth. We adopt a moving average with weights of 0.5 for $t-3$, 1 for $t-2$, and 0.5 for $t-1$.

[122] Similar specifications are used by A. J. Auerbach & Gorodnichenko (2012), Ramey & Zubairy (2018), Tenreyro & Thwaites (2016), and De Ridder & Pfajfar (2017).

whether the economy is in recession or expansion. If two-year lagged growth was relatively high, the observation weights towards β_E while it weights more towards β_R if lagged growth was low. Following Tenreyro & Thwaites (2016), $F(z_{m,t})$ is a logistic function:

$$F(z_{m,t}) = \frac{\exp\left(\theta \frac{[z_{m,t}-\mu_m]}{\sigma_m}\right)}{1+\exp\left(\theta \frac{[z_{m,t}-\mu_m]}{\sigma_m}\right)}, \tag{3.6}$$

where μ_m determines the fraction of the sample in which the metropolitan area is in recession, σ_m gives the standard deviation of biannual growth and θ determines how stark the demarcation between recessions and expansions are (e.g., for a lower θ, the weight of observations is more equally split between β_E and β_R). μ_m is calibrated such that each area is in recession 20 percent of the sample, which matches the percent of quarters that the economy is in recession at the national level according to the NBER. We calibrate θ to 3 in line with Tenreyro & Thwaites (2016). We estimate equation (3.5) using two-stage least squares, where Pell grant disbursements at the MSA level are instrumented using the same instruments as in our main regressions, but multiplied by $F(z_{m,t-2})$ for the expansion term and $1 - F(z_{m,t-2})$ for the recession term.

Results are presented in Table 3.8. Recession multipliers represent β_R in eq. (3.5) while expansion multipliers represent β_E. The recession (expansion) should be interpreted as the two-year effect of a relative increase in Pell grants on relative income growth if growth is initially at its *lowest* (highest) level in the dataset. Thus, the multipliers in the table are for these extreme observations. The actual multiplier of an increase in Pell grant disbursements depends on how close growth is to either of these levels. As before, panel A reports state-dependant fiscal multipliers for personal income growth and panel B for employment growth. Column (1) contains the specification that controls for metropolitan and year fixed effects. Column (2) adds appropriations, while column (3) adds the area-specific controls, but does not control for appropriations. Column (4) controls for both appropriations and the area-specific controls. Columns (5)-(7) repeat the regressions in columns (1), (2), and (5) on the post-1999 sample, while columns (8)-(10) add the financial control variables and explore the relevance of appropriations and other major fiscal transfers to low-income households.

All columns show considerably larger multipliers in recessions than in expansions, where the differences between the two multipliers are most of the time statistically significant for the post-1999 sample. This result holds for both income growth and employment growth. In our preferred estimate on the full sample with all controls, column (4), the multiplier for income growth is 2.4 in expansions and 3.2 in recessions. While the differences between the two multipliers are

Table 3.8: State-Dependence of Education Spending Multiplier

	Full Sample				Post 1999						Full Sample
	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS	(9) 2SLS	(10) 2SLS	(11) OLS
Panel A: Income Growth											
Recession Multiplier	3.494*	3.440*	3.272*	3.173	6.514**	6.506**	6.727***	6.238**	6.234**	6.265***	-0.470
	(1.867)	(1.918)	(1.912)	(1.969)	(2.521)	(2.538)	(2.471)	(2.425)	(2.441)	(2.397)	(1.395)
Expansion Multiplier	1.990	1.952	2.494	2.429	1.969	1.955	1.065	1.137	1.131	1.361	-1.533
	(1.795)	(1.815)	(1.802)	(1.817)	(1.997)	(2.034)	(2.009)	(1.940)	(1.973)	(1.931)	(1.056)
Panel B: Employment Growth											
Recession Multiplier	2.057	2.263	2.227	2.394*	5.382***	5.432***	5.140***	5.016***	5.073***	5.045***	0.639
	(1.403)	(1.395)	(1.408)	(1.406)	(1.794)	(1.800)	(1.742)	(1.731)	(1.736)	(1.739)	(0.845)
Expansion Multiplier	1.472	1.803	1.431	1.705	0.183	0.403	-0.162	-0.422	-0.188	-0.400	-2.492*
	(1.840)	(1.848)	(1.811)	(1.835)	(1.815)	(1.835)	(1.887)	(1.845)	(1.860)	(1.842)	(1.287)
Difference, Income	-1.504	-1.488	-0.778	-0.744	-4.545	-4.550	-5.661**	-5.101*	-5.103*	-4.905*	-1.063
Std. Error, Income	(2.331)	(2.336)	(2.428)	(2.434)	(2.924)	(2.927)	(2.866)	(2.752)	(2.755)	(2.728)	(1.674)
Difference, Employment	-0.585	-0.460	-0.796	-0.688	-5.199**	-5.029**	-5.301**	-5.438**	-5.261**	-5.446**	-3.131*
Std. Error, Employment	(2.239)	(2.221)	(2.215)	(2.206)	(2.493)	(2.493)	(2.565)	(2.541)	(2.538)	(2.549)	(1.604)
Observations	8,435	8,435	8,435	8,435	4,446	4,446	4,446	4,446	4,446	4,446	8,435
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, Recession	216.2	214.0	217.0	214.8	253.7	254.7	259.0	260.9	263.6	267.7	-
Δ Pell Grants F-test, Expansion	141.4	142.7	140.9	142.6	92.7	83.4	81.2	91.2	82.3	90.4	-
Joint F-test	50.9	49.3	50.9	49.8	33.4	24.7	24.2	32.8	24.3	23.4	-

Notes: Multipliers follow from Smooth Transition estimates. SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population (see eq. 3.3). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, Recession is the robust F-statistic of the first-stage regression of Pell grants multiplied with $(1 - F(z_{m,t-2}))$. Δ Pell Grants F-test, Expansion is the robust F-statistic of the first-stage regression of Pell grants multiplied with $F(z_{m,t-2})$. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

between 0.7 and 1.5 points in the full sample, the estimated multipliers in expansions are never statistically significant. Similarly, employment growth multipliers 1.7 in expansion and 2.4 in recession, where the difference between the two multipliers across specifications for the full sample range between 0.5 and 0.8 points.

The results for the post-1999 sample, columns (5)-(10), predict much larger gaps in the multipliers between recessions and expansions (about 5 points) with the difference being significant at conventional confidence levels. This result holds for both income and employment growth multipliers. While the standard errors of our estimates suggest that the effect is noisy, the large estimates of the recession multiplier suggest that Pell grants are particularly effective when the local economy is in a recession. These results further the case that Pell grants offer a tool to stimulate economic activity when needed, which means that they can

Figure 3.4: *Percentage of Pell Grants Awarded to For-Profit Schools*

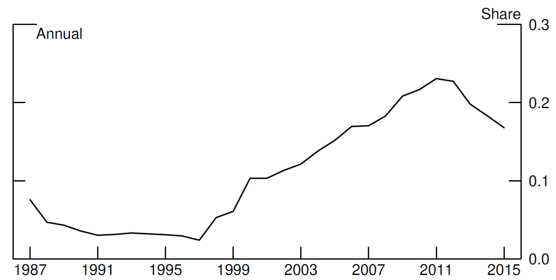

Notes: Figure plots the fraction of national-level Pell grants that is awarded to students who are enrolled at for-profit institutions. Data is obtained from Delta Cost.

potentially be used as part of countercyclical fiscal policy.

3.4.2 Institutions: For-Profit versus Non-Profit Colleges

We next assess whether multipliers depend on the type of institution attended by the beneficiary student. The previous sections have shown that Pell grants have substantial multipliers, especially during recessions. One objection to using Pell grants for countercyclical policy may be, however, that 15–20 percent of grants is spent at for-profit colleges (Figure 3.4).[123] If for-profit colleges have market power, they may be able to charge higher tuition fees in response to higher generosity of Pell grants. Pell grants can therefore operate as an implicit subsidy. As public companies own a large fraction of for-profit colleges, not all of these subsidies will be spent within the college's metropolitan area, reducing local economic effects.[124]

Given these concerns, we explore whether there is indeed a relationship between the multiplier of Pell grants and the for-profit status of institutions at which students study. Because for-profit Pell grants and non-profit Pell grants may be correlated at the MSA level, we estimate both multipliers jointly:

$$\frac{\Delta y_{m,t;t-2}}{y_{m,t-2}} = \beta^{FP} \frac{\Delta e^{FP}_{m,t;t-2}}{y_{m,t-2}} + \beta^{NP} \frac{\Delta e^{NP}_{m,t;t-2}}{y_{m,t-2}} + \phi_m + \psi_t + \gamma' x_{m,t-2} + \mu_{m,t}, \tag{3.7}$$

where $e^{FP}_{m,t}$ denotes the total amount of Pell grants awarded to for-profit schools in metropolitan area m in year t, while $e^{NP}_{m,t}$ denotes the amount awarded to non-profit schools. As instruments we thus use:

$$b^z_{m,t} = \left(\frac{\Delta e^z_{t;t-2}}{y_{t-2}} \right) \cdot s^z_{m,t-2}; \quad z \in \{FP, NP\} \tag{3.8}$$

[123]The reduction after 2013 is the result of "Gainful Employment" regulation. This regulation restricts federal student aid at several for-profit institutions (see, for example, Cellini et al., 2016).

[124]Examples of publicly listed companies that own for-profit colleges are Grand Canyon University (LOPE), Adtalem (ATGE, previously DeVry), American Public University System (APEI), and Bridgepoint Education Inc. (BPI).

Table 3.9: Effect of Pell Grants on Local Income Per Capita: For-Profit versus Non-Profit

	Full Sample				Post 1999						Full Sample
	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS	(9) 2SLS	(10) 2SLS	(11) OLS
Panel A: Income Growth											
Non-Profit Multiplier	4.577**	4.485**	4.935***	4.788**	5.914***	5.956***	5.125**	4.694**	4.728**	5.059**	-1.482
	(1.773)	(1.877)	(1.760)	(1.867)	(2.015)	(2.105)	(2.057)	(2.043)	(2.121)	(2.036)	(0.967)
For-Profit Multiplier	2.047	2.058	2.053	2.071	2.837	2.833	2.908*	3.427*	3.425*	3.462*	-2.459
	(1.596)	(1.593)	(1.633)	(1.628)	(1.783)	(1.783)	(1.730)	(1.806)	(1.806)	(1.789)	(2.657)
Panel B: Employment Growth											
Non-Profit Multiplier	2.240	2.667*	2.550*	2.891*	4.140**	4.429***	3.840**	3.383**	3.674**	3.465**	-1.544*
	(1.472)	(1.487)	(1.467)	(1.500)	(1.608)	(1.616)	(1.604)	(1.586)	(1.594)	(1.586)	(0.865)
For-Profit Multiplier	1.449	1.398	1.507	1.467	2.397	2.370	2.505*	2.585*	2.563*	2.593*	-1.640
	(1.264)	(1.265)	(1.238)	(1.240)	(1.478)	(1.476)	(1.454)	(1.448)	(1.445)	(1.446)	(1.610)
Difference, Income	2.530	2.427	2.882	2.717	3.077	3.123	2.217	1.267	1.303	1.596	0.977
Std. Error, Income	(2.188)	(2.273)	(2.221)	(2.305)	(2.385)	(2.464)	(2.379)	(2.445)	(2.514)	(2.435)	(2.762)
Difference, Employment	0.791	1.269	1.044	1.423	1.743	2.059	1.335	0.798	1.112	0.872	0.097
Std. Error, Employment	(1.859)	(1.882)	(1.837)	(1.873)	(2.026)	(2.041)	(2.008)	(1.982)	(1.996)	(1.984)	(1.799)
Observations	8,432	8,432	8,432	8,432	4,443	4,443	4,443	4,443	4,443	4,443	8,432
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, NP	107.3	97.1	107.1	97.1	85.0	79.9	80.2	85.5	80.3	92.2	-
Δ Pell Grants F-test, FP	28.8	29.8	29.0	29.8	33.4	33.3	33.6	33.0	33.0	34.9	-
Joint F-test	53.7	39.9	53.7	40.3	42.4	20.6	20.6	42.7	20.5	27.3	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at for-profit and non-profit institutions (see eq. 3.8). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization and 30-day mortgage delinquency. Robust F-statistic is for the Pell grants SSIV. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, NP is the robust F-statistic of the first-stage regression of Pell grants at non-profit colleges. Δ Pell Grants F-test, FP is the robust F-statistic of the first-stage regression of Pell grants at for-profit colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

where $e^z_{t;t-2}$ is the per-capita national change in the Pell grants awarded to each group (for-profit, non-profit) and $s^z_{m,t-2}$ is the share of Pell grant recipients within each group (for-profit, non-profit) in an MSA.[125]

Results are presented in Table 3.9. Control variables follow the same sequence as in Table 3.2. The for-profit multiplier of Pell grants estimates the multiplier effects of grants awarded to private for-profit schools, while the non-profit multiplier estimates the effects of Pell grants at other schools. By including both es-

[125] In some cases the recipient data is missing in our databases for a specific MSA-year. For these MSA-years—about 29 percent of all of the MSA-years—, we use MSA-level recipient proxy based on the share of for-profit institutions among all institutions in a specific MSA and the share of undergraduate population in the total population in that MSA. In addition, pre-1999 instruments were computed using the proxy described in Section 3.2.3.

Table 3.10: Effect of Pell Grants on Tuition Fees

Tuition Growth	For-Profit	Non-Profit
	(1)	(2)
Δ Pell Grants (% Tuition)	1.749**	1.715
	(0.822)	(1.678)
Observations	14,697	65,386
Time FE	Y	Y
School FE	Y	Y
Instrument F-test	13.1	13.1

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at for-profit and non-profit institutions. Dependent and independent variables are winsorized at the 1^{th} and 99^{th} percentiles. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

timates in the same specification we control for the correlation between awards at both types of schools. As we can see in Table 3.9, both income growth and employment multipliers are considerably higher when Pell grants are awarded at non-profit schools than at for-profit schools where the difference between the multipliers is between 1.3 and 3.1 points for income growth multipliers and between 0.8 and 2.1 for employment growth multipliers, although it is not statistically significant at conventional levels due to high standard errors. The income multipliers for non-profit schools ranges from 4.5 to 6, while for for-profit schools from 2.1 to 3.6. The employment multipliers range between 2.2 and 4.4 at non-profit schools and 1.4 and 2.5 at for-profit schools. Multipliers initiated from grants to for-profit schools are only significant in columns (7)-(10) in both our panels. These estimates imply that there are notable differences in the policy transmission of the education spending depending on profit orientation of recipient schools. Multipliers in the for-profit education sector are considerably smaller in all specifications.

We next assess the drivers of the smaller multipliers of Pell grant receipts at for-profit schools. Using school-level micro data on enrollment, expenditures, and revenue sources from Delta Cost, we first explore the effect of Pell grant receipts on a school's tuition fees. We define a school's tuition fee as the amount of tuition received directly from students, net of any grants or (institutional) student aid, divided by the number of full-time students. The estimation equation reads:

$$\frac{\Delta\tau_{i,t;t-2}}{\tau_{i,t-2}} = \Gamma\left(\frac{\Delta e_{i,t;t-2}}{e_{i,t-2}}\right) + \phi_i + \psi_t + \mu_{i,t}, \tag{3.9}$$

where τ_{it} is the average tuition rate at school i during academic year t, while $e_{i,t}$ denotes the amount of Pell grants received per full-time equivalent student. We look at biannual changes to match the horizon over which we estimate the multiplier and estimate Γ separately for for-profit and non-profit schools.

The estimation of eq. (3.9) is subject to endogeneity because an increase in demand for schooling may increase both tuition fees and the number of Pell grants a school receives. To address this, we again resort to shift-share instrumental variables. We instrument school-level Pell grants by the interaction of the share of an institution's student body that receives Pell grants—analogous to our city-level shift-share instrument—and interact this with national changes in the Pell grant program along eq. (3.3).

Results are presented in Table 3.10. When Pell grants increase as a percentage of the total tuition revenue, both non-profit and for-profit schools increase their average tuition fees by about 1.7 percent when Pell grants share in the total tuition increase by 1 percentage point. However, only the estimate for the for-profit schools is significant. The standard error for non-profit schools is large and the estimate is not statistically significant, which may suggest greater heterogeneity in non-profit schools' reaction to Pell grant increase. The estimate at for-profit schools suggests that for-profit schools raise tuition fees more than proportionally when the grants increase.

From the results above, it seems that it is likely that both for-profit and non-profit schools raise tuition fees when Pell grant generosity increases. As tuition hikes prevent students from gaining purchasing power when Pell grants disbursements increase, these grants may have a smaller effect on economic activity when tuition fees increase. This may, at least partly, explain why we find lower multipliers at for-profit schools.

While the increase in tuition fees is slightly larger (and significant) at for-profit schools, both types of institutions may to some degree raise their tuition fees in response to Pell grant increases. This is in line with the "Bennett Hypothesis," first proposed by former Secretary of Education William Bennett. The hypothesis roughly yields that colleges expropriate rises in student aid. At for-profit colleges, this may raise their profits. At non-profit colleges, schools may use additional tuition to subsidize and expand their broader activities, such as research. Previous studies supporting the hypothesis are anchored by Cellini & Goldin (2014) which links higher tuition charged by for-profit institutions with their eligibility for federal aid and Lucca et al. (2019) which documents a 60 cents on the dollar pass-through effect on tuition of changes in subsidized loan maximums and about 20 cents on the dollar for unsubsidized federal loans. Among the studies challenging the hypothesis is Rizzo & Ehrenberg (2004), which found no evidence of tuition increases in response to increases in federal or state financial aid, Kelchen (2019) which showed that law schools did not raise tuition prices once federal aid was increased, and Kelchen (2020), which found a similar result for medical and business schools.

As the difference in tuition hikes is relatively small, we examine further what may explain the smaller effect of for-profit Pell grant disbursements on growth. One potential mechanism is that for-profit and non-profit schools differ in how

Table 3.11: Effect of Education Spending on College Expenditures: For-Profit versus Non-Profit

Expenditure Growth	Full Sample				Post 1999						Full Sample
	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS	(9) 2SLS	(10) 2SLS	(11) OLS
Non-Profit Δ Pell Grants	0.763	1.125*	0.777	1.136*	1.115	1.218*	1.226*	1.158	1.261*	1.185	0.976***
	(0.628)	(0.619)	(0.627)	(0.617)	(0.744)	(0.705)	(0.706)	(0.736)	(0.698)	(0.721)	(0.240)
For-Profit Δ Pell Grants	1.507***	1.464***	1.498***	1.457***	1.451***	1.441***	1.430***	1.401***	1.393***	1.403***	1.378***
	(0.199)	(0.194)	(0.201)	(0.196)	(0.222)	(0.219)	(0.221)	(0.227)	(0.224)	(0.228)	(0.190)
Difference	-0.745	-0.339	-0.722	-0.322	-0.336	-0.223	-0.204	-0.243	-0.132	-0.218	-0.402
Std. Error	(0.552)	(0.564)	(0.552)	(0.562)	(0.640)	(0.612)	(0.612)	(0.640)	(0.612)	(0.628)	(0.305)
Observations	8,432	8,432	8,432	8,432	4,443	4,443	4,443	4,443	4,443	4,443	8,432
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, NP	107.3	97.1	107.1	97.1	85.0	79.9	80.2	85.5	80.3	92.3	-
Δ Pell Grants F-test, FP	28.8	29.8	29.0	29.8	33.4	33.3	33.6	33.0	33.0	35.0	-
Joint F-test	53.7	39.9	53.7	40.3	42.4	20.6	20.6	42.7	20.5	27.3	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at for- and non-profit institutions (see eq. 3.8). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, NP is the robust F-statistic of the first-stage regression of Pell grants at non-profit colleges. Δ Pell Grants F-test, FP is the robust F-statistic of the first-stage regression of Pell grants at for-profit colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

they *use* their additional tuition revenue. While tuition fee hikes moderate the effect of Pell grants on students' consumer spending, grants can also positively affect growth when schools spend their tuition revenue productively. To explore this, we estimate the effect of overall expenditures at for-profit and non-profit schools when Pell grants rise.

To see how total expenditures by colleges respond to a change in Pell grants, Table 3.11 estimates the "college spending multiplier" of the grants.[126] The dependent variable is the biannual change in total expenditure as a percentage of aggregate personal income in the MSA, analogous to equation (3.2). Pell grants are instrumented with our shift-share instrument as before.

The results in Table 3.11 show evidence that both non-profit and for-profit colleges raise spending when the Pell grant program increases in generosity. The

[126]We focus on the expenditures directly related to the primary function of universities: education expenditures (instruction, student services) and non-education expenditures (research, public service). Education and non-education expenditures constitute total expenditures in our definition. From our measure of total expenditures we are thus excluding expenses incurred from institutional support (day-to-day operational support, like HR, legal service, etc.), main operating expenses (utilities, insurance, etc.), and grants and scholarships.

point estimate is a touch lower than the one in Table 3.10, suggesting that colleges increase their expenditures at a rate slightly below the one for tuition increases when the Pell Grant Program expands. The point estimate for spending at for-profit schools is about 1.4, while the estimate at non-profit colleges is about 1.1. The results in Table 3.11 are in line with the finding in Dinerstein et al. (2014) that public universities (a subset of our non-profit universities) increased their educational expenditures during the Great Recession as a result of the increase in the maximum Pell grants that occurred during 2009/2010.

As these results do not explain why the effect of Pell grants on economic activity is smaller when grants are disbursed to for-profit schools, we consider two further potential explanations. The first potential explanation is that for-profit and non-profit colleges increase different *types* of college expenditures in response to the increase in generosity of the Pell grant program. The second possibility is that the consumer spending response to Pell grants is different for students at for-profit and non-profit schools. While we attempt to evaluate the first explanation, the evidence in favor of the second explanation is largely by exclusion, as the data to directly address it are not available. Tables C.9–C.14 in the Appendix conduct the estimation separately for education expenditures and non-education expenditures and then further split the education expenditures into instruction and student services, and non-education expenditures into research and public service. Results suggest that for-profit schools only increase education expenditures—where the estimated multiplier is larger for student services than for instruction—while non-profit schools mostly increase non-education expenditures, where only the effect on the research subcomponent is statistically significant. There is also some evidence that these schools increase student services. Ideally, the next step would be to assess whether these different expenditures changes in response to Pell grants' increase in generosity lead to a positive effect of college expenditures on local income. However, this analysis turns out to be challenging due to weak instrument problem when we use our standard instruments.[127] Thus, it is difficult to judge exactly what portion of the Pell grants multiplier on local income may go through the college expenditure, if any, as we cannot reliably estimate the effect of college expenditures on local income. The fact that the "college spending multiplier" is smaller than the overall Pell grant multiplier also suggests that the large effect of Pell grants on growth is unlikely to work primarily through expenditures by the college. Rather, it seems that student loan increases as a result of Pell grants increases and the associated relaxation of students' budget constraint may be the main driver of the grants' economic ef-

[127] We have tried alternative instruments that rely on national college expenditure growth and the lagged share of student population in a city. We report these results in Table C.15. The estimates of the college expenditure multiplier are positive but insignificant: These estimates suffer from the weak instrument problem.

Figure 3.5: *Percentage of Pell Grants Awarded at Two-year Institutions*

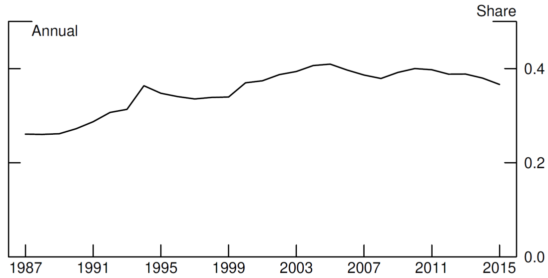

Notes: Figure plots the fraction of national-level Pell grants that is awarded to students who are enrolled at two-year institutions. Data is obtained from Delta Cost.

fects.

3.4.3 Institutions: Two-Year versus Four-Year Colleges

Finally, we assess whether there are differences in the multiplier of Pell grants between four-year and two-year colleges. Two-year colleges are typically community colleges that offer post-secondary education to local students. Four-year colleges are often more broadly engaged in academic activities, including research. The share of Pell grants received by students at two-year colleges has gradually increased over time: while only 25 percent of all Pell grants were disbursed to two-year institutions in 1987, this share has increased in the 80's and 90's and has fluctuated between 35 and 40 percent (see Figure 3.5).

To test whether multipliers between these types of institutions are different, we estimate eq. (3.7) with Pell grant disbursements to two- and four-year institutions rather than for-profit and non-profit institutions. Table 3.12 presents the results. We find that multipliers are larger at four-year institutions compared to two-year institutions. Income multipliers at four-year institutions range from 2.7 to 3.9 and employment multipliers range between 1.4 and 3.5, while income multipliers at two-year institutions range from -0.7 to 2.2 and employment multipliers range between -0.6 and 0.4. Only multipliers at the four-year institution are significantly different from zero. In our preferred specification that uses all available controls for the full sample (column (4)), the income multiplier of four-year schools is 1.1 points higher and employment multiplier is 2.0 points higher. The differences in multipliers between two- and four-year schools are even larger in the post-1999 sample.

We next assess whether the difference in multipliers between two- and four-year colleges is due to differences in the response of spending by these types of institutions. Table 3.13 presents the results, which is analogous to Table 3.11 for for-profit versus non-profit schools. The table shows that at four-year colleges, an increase in Pell grants by 1 percent of local personal income leads to an increase in total college expenditures by 1.4 percent of local personal income. In contrast, two-year institutions do not significantly increase their total expendi-

Table 3.12: Effect of Pell Grants on Local Income Per Capita: Two-Year versus Four-Year Schools

	Full Sample				Post 1999						Full Sample
	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS	(9) 2SLS	(10) 2SLS	(11) OLS
Panel A: Income Growth											
4-year Multiplier	2.814*	2.706*	3.229**	3.079*	4.133**	4.134**	3.709**	3.638**	3.661**	3.923**	-1.443
	(1.481)	(1.545)	(1.539)	(1.600)	(1.760)	(1.813)	(1.759)	(1.700)	(1.751)	(1.723)	(1.163)
2-year Multiplier	2.163	2.223	1.914	1.998	1.099	1.098	0.172	-0.150	-0.164	-0.740	-2.101
	(3.545)	(3.549)	(3.544)	(3.548)	(3.830)	(3.836)	(3.745)	(3.797)	(3.806)	(3.856)	(1.788)
Panel B: Employment Growth											
4-year Multiplier	1.419	1.726	1.759	2.001	3.343**	3.529**	3.261**	2.964**	3.173**	3.030**	-0.919
	(1.362)	(1.383)	(1.370)	(1.406)	(1.510)	(1.535)	(1.508)	(1.461)	(1.483)	(1.463)	(0.822)
2-year Multiplier	0.366	0.194	0.145	0.011	0.279	0.176	-0.523	-0.424	-0.549	-0.561	-2.852*
	(1.936)	(1.925)	(1.954)	(1.948)	(2.110)	(2.104)	(2.042)	(2.019)	(2.016)	(2.027)	(1.607)
Difference, Income	0.651	0.482	1.315	1.081	3.034	3.036	3.536	3.789	3.825	4.663	0.658
Std. Error, Income	(3.657)	(3.711)	(3.708)	(3.763)	(3.713)	(3.770)	(3.668)	(3.668)	(3.727)	(3.767)	(2.241)
Difference, Employment	1.053	1.532	1.613	1.990	3.064	3.353	3.784	3.388	3.721	3.591	1.934
Std. Error, Employment	(2.397)	(2.416)	(2.420)	(2.450)	(2.424)	(2.441)	(2.400)	(2.347)	(2.365)	(2.374)	(1.726)
Observations	8,436	8,436	8,436	8,436	4,447	4,447	4,447	4,447	4,447	4,447	8,436
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes		Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes	Yes	Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, 4-Year	133.8	119.0	133.9	119.3	95.3	86.0	84.7	93.8	84.5	94.9	-
Δ Pell Grants F-test, 2-Year	65.2	65.0	65.2	65.2	54.0	54.1	54.8	54.9	55.2	57.3	-
Joint F-test	32.5	43.1	32.6	43.4	29.5	26.0	25.6	29.9	25.6	20.1	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at four-year and two-year institutions. We estimate the four-year share based on the MSA four-year penetration and the MSA-level number of recipients where missing. SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, 4-Year is the robust F-statistic of the first-stage regression of Pell grants at 4-year colleges. Δ Pell Grants F-test, 2-Year is the robust F-statistic of the first-stage regression of Pell grants at 2-year colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

tures in response to the increase of Pell grants. The difference in response of two- and four-year colleges—that is often statistically significant—may be behind the large standard errors in Table 3.11 for non-profit schools, as roughly one third of students at non-profit schools are enrolled at public two-year institutions (community colleges) that do not raise spending in response to the increase of Pell grants. Thus, the non-profit sector exhibits a significant degree of heterogeneity.

We proceed with a similar analysis to the one we performed for the for-profit and non-profit schools and evaluate the effect of Pell grants on various types of college expenditures. These results are reported in Appendix Tables C.16–C.21. These tables show that the main source of the overall increase in expenditures

Table 3.13: Effect of Pell Grants on College Expenditures: Two-Year versus Four-Year Schools

Expenditure Growth	Full Sample				Post 1999						Full Sample
	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS	(9) 2SLS	(10) OLS	(11) est11
4-year Δ Pell Grants	1.107**	1.392***	1.120**	1.401***	1.362**	1.446**	1.442**	1.366**	1.448**	1.393**	1.152***
	(0.551)	(0.526)	(0.549)	(0.524)	(0.608)	(0.570)	(0.571)	(0.606)	(0.569)	(0.592)	(0.234)
2-year Δ Pell Grants	0.534	0.374	0.523	0.367	0.369	0.322	0.326	0.399	0.350	0.343	0.852*
	(0.335)	(0.323)	(0.336)	(0.323)	(0.344)	(0.337)	(0.338)	(0.345)	(0.337)	(0.343)	(0.438)
Difference	0.573	1.017*	0.598	1.035*	0.993	1.124*	1.117*	0.966	1.097*	1.050*	0.300
Std. Error	(0.547)	(0.556)	(0.547)	(0.554)	(0.607)	(0.577)	(0.579)	(0.606)	(0.576)	(0.594)	(0.501)
Observations	8,436	8,436	8,436	8,436	4,447	4,447	4,447	4,447	4,447	4,447	8,436
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes	Yes	Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, 4-Year	133.8	119.0	133.9	119.3	95.3	86.0	84.7	93.8	84.5	95.0	-
Δ Pell Grants F-test, 2-Year	65.2	65.0	65.2	65.2	54.0	54.1	54.8	54.9	55.2	57.2	-
Joint F-test	32.5	43.1	32.6	43.4	29.5	26.0	25.6	29.9	25.6	20.1	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at four-year and two-year institutions. We estimate the four-year share based on the MSA four-year penetration and the MSA-level number of recipients where missing. SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, 4-Year is the robust F-statistic of the first-stage regression of Pell grants at 4-year colleges. Δ Pell Grants F-test, 2-Year is the robust F-statistic of the first-stage regression of Pell grants at 2-year colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

in Table 3.13 is the increase of education expenditures at four-year institutions—in particular student services—, although in the overall sample non-education expenditures—in particular research expenditures—also display significant effects. Two-year institutions do not increase their expenditures in response to the increase in generosity of the Pell Grant Program. College spending effects are smaller also for two- and four-year colleges than the overall Pell grant multiplier, similarly as in the case for for-profit and non-profit colleges, suggesting that student spending is the main driver of the Pell grant's multiplier on local income.

Our findings offer a broader perspective on the "Bennett Hypothesis". While we confirm that four-year colleges raise expenditures in response to an increase in Pell grants, we also find that the multiplier of Pell grants is largest at these schools. While some part of higher Pell grants gets transferred to schools, this transfer does not seem to impede the beneficial local economic effects of expansions of the Pell Grant Program—the additional spending by colleges may even enhance them.

3.5 Conclusion

This paper estimates the effect of the Federal Pell Grant Program on short-run economic activity. Specifically, we assess how a relative increase in Pell grant disbursements at the metropolitan area raises the area's relative income and relative employment. To do so, we employ a shift-share approach where our identification relies on the variation in Pell grant receipts across metropolitan areas. We deploy a series of validity tests for the shift-share instrument to ensure that the empirical strategy delivers a causal estimate of the effect of the Pell Grant Program on local economic growth.

We find an average income multiplier of 2.8 and an employment multiplier of 1.9 in the main specification. This implies that a 1 percent increase in Pell grants as a fraction of local income raises local income by 2.8 percent and local employment by 1.9 percent. These multipliers are higher than the average estimates of the multipliers from geographical cross-sectional data of other forms of fiscal spending found in the literature, e.g., the multipliers of military spending. Pell grants are fiscal transfers that raise personal income one-by-one and are awarded to students from lower-income households that tend to have higher propensity to consume than wealthier households. Our results suggest that, in part, the Pell grant fiscal multiplier operates through enabling students to attend college and acquire students loans. An increase in generosity of Pell grants increases student loans disbursements. This increase in disbursements further eases students' budget constraint and allows them to spend more. We also find that multipliers are higher when the economy is in recession. Our results imply that besides having beneficial effects in the long run, educational investments can also be used for countercyclical fiscal policy.

Our findings also have implications for education policy. We find that the multiplier of the Pell Grant Program is higher at non-profit colleges. For-profit colleges raise education spending when Pell grants become more generous, but there appears to be no effect of these expenditures on local economic growth or local employment. This result offers some validation for recent restrictions imposed on the eligibility of students at for-profit institutions for Pell grants. Finally, we show that four-year institutions have larger multipliers than two-year institutions. Pell grants are therefore particularly effective as a tool for countercyclical policy if granted to students attending four-year non-profit colleges.

References

Acconcia, A., Corsetti, G., & Simonelli, S. (2014). Mafia and public spending: Evidence on the fiscal multiplier from a quasi-experiment. *American Economic*

Review, *104*(7), 2185–2209.

Auerbach, A., Gorodnichenko, Y., McCrory, P. B., & Murphy, D. (2022). Fiscal multipliers in the covid19 recession. *Journal of International Money and Finance*, *126*, 102669.

Auerbach, A., Gorodnichenko, Y., & Murphy, D. (2020). Local fiscal multipliers and fiscal spillovers in the usa. *IMF Economic Review*, *68*, 195–229.

Auerbach, A. J., & Gorodnichenko, Y. (2012). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, *4*(2), 1-27.

Barnichon, R., Debortoli, D., & Matthes, C. (2021). Understanding the Size of the Government Spending Multiplier: It's in the Sign. *The Review of Economic Studies*, *89*(1), 87-117.

Barro, R. J. (1991). Economic Growth in a Cross Section of Countries. *The Quarterly Journal of Economics*, *106*(2), 407-443. Retrieved from https://ideas.repec.org/a/oup/qjecon/v106y1991i2p407-443..html

Barro, R. J., Sala-i Martin, X., Blanchard, O. J., & Hall, R. E. (1991). Convergence across states and regions. *Brookings papers on economic activity*, 107–182.

Benhabib, J., & Spiegel, M. M. (1994). The role of human capital in economic development evidence from aggregate cross-country data. *Journal of Monetary Economics*, *34*(2), 143 - 173.

Berge, T., De Ridder, M., & Pfajfar, D. (2021). When is the fiscal multiplier high? a comparison of four business cycle phases. *European Economic Review*, *138*, 103852.

Bettinger, E. (2004). How financial aid affects persistence. In *College choices: The economics of where to go, when to go, and how to pay for it* (pp. 207–238). University of Chicago Press.

Bils, M., & Klenow, P. J. (2000). Does schooling cause growth? *American Economic Review*, *90*(5), 1160-1183.

Black, S. E., Denning, J. T., Dettling, L. J., Goodman, S., & Turner, L. J. (2020). *Taking it to the limit: Effects of increased student loan availability on attainment, earnings, and financial well-being* (Working Paper No. 27658). National Bureau of Economic Research.

Blanchard, O. J., & Leigh, D. (2013). Growth forecast errors and fiscal multipliers. *The American Economic Review*, *103*(3), 117–120.

Borusyak, K., Hull, P., & Jaravel, X. (2022). Quasi-experimental shift-share research designs. *The Review of Economic Studies*, *89*(1), 181–213.

Castleman, B. L., & Long, B. T. (2016). Looking beyond enrollment: The causal effect of need-based grants on college access, persistence, and graduation. *Journal of Labor Economics*, *34*(4), 1023–1073.

Cellini, S. R., Darolia, R., & Turner, L. J. (2016, December). *Where Do Students Go when For-Profit Colleges Lose Federal Aid?* (NBER Working Papers No. 22967).

National Bureau of Economic Research, Inc.

Cellini, S. R., & Goldin, C. (2014). Does federal student aid raise tuition? new evidence on for-profit colleges. *American Economic Journal: Economic Policy*, *6*(4), 174-206.

Chao, J. C., Hausman, J. A., Newey, W. K., Swanson, N. R., & Woutersen, T. (2014). Testing overidentifying restrictions with many instruments and heteroskedasticity. *Journal of Econometrics*, *178*, 15-21. (Annals Issue: Misspecification Test Methods in Econometrics)

Chhabra, Y., Levenstein, M. C., & Owen-Smith, J. (2019). *The local economic impact of science spending: Evidence from the american recovery and reinvestment act* (Working Paper No. 1383). Ross School of Business.

Chodorow-Reich, G. (2019). Geographic cross-sectional fiscal spending multipliers: What have we learned? *American Economic Journal: Economic Policy*, *11*(2), 1-34.

Chodorow-Reich, G., Feiveson, L., Liscow, Z., & Woolston, W. G. (2012). Does state fiscal relief during recessions increase employment? evidence from the american recovery and reinvestment act. *American Economic Journal: Economic Policy*, *4*(3), 118-45.

Conley, T. G., & Dupor, B. (2013). The american recovery and reinvestment act: Solely a government jobs program? *Journal of Monetary Economics*, *60*(5), 535 - 549. (Aggregate Implications of Local Public Finance)

Corsetti, G., Meier, A., & Müller, G. J. (2012). What determines government spending multipliers? *Economic Policy*, *27*(72), 521–565.

Delisle, J. (2017). The pell grant proxy: A ubiquitous but flawed measure of low-income student enrollment. *Evidence Speaks Reports*, *2*(26), 1–12.

Denning, J. T., Marx, B. M., & Turner, L. J. (2019). Propelled: The effects of grants on graduation, earnings, and welfare. *American Economic Journal: Applied Economics*, *11*(3), 193-224.

De Ridder, M., & Pfajfar, D. (2017, November). Policy shocks and wage rigidities: Empirical evidence from regional effects of national shocks. *Working Paper*.

Dinerstein, M. F., Hoxby, C. M., Meer, J., & Villanueva, P. (2014). Did the fiscal stimulus work for universities? [Book]. In *How the financial crisis and great recession affected higher education* (p. 263-320). University of Chicago Press.

Dunn, A. C., & Gholizadeh, M. (2023, January). *The Geography of Consumption and Local Economic Shocks: The Case of the Great Recession* (BEA Working Papers No. 0205). Bureau of Economic Analysis. Retrieved from https://ideas.repec.org/p/bea/wpaper/0205.html

Dupor, B., & McCrory, P. B. (2018, June). A Cup Runneth Over: Fiscal Policy Spillovers from the 2009 Recovery Act. *Economic Journal*, *128*(611), 1476-1508. Retrieved from https://ideas.repec.org/a/wly/

econjl/v128y2018i611p1476-1508.html doi: 10.1111/ecoj.12475

Dupor, B., & Mehkari, M. S. (2016). The 2009 recovery act: Stimulus at the extensive and intensive labor margins. *European Economic Review*, *85*, 208 - 228.

Dynarski, S. M. (2003). Does aid matter? measuring the effect of student aid on college attendance and completion. *American Economic Review*(93), 179-88.

Fack, G., & Grenet, J. (2015). Improving college access and success for low-income students: Evidence from a large need-based grant program. *American Economic Journal: Applied Economics*, *7*(2), 1-34.

Feyrer, J., & Sacerdote, B. (2011). *Did the stimulus stimulate? real time estimates of the effects of the american recovery and reinvestment act* (Working Paper No. 16759). National Bureau of Economic Research.

Fuller, W. A. (1977). Some properties of a modification of the limited information estimator. *Econometrica: Journal of the Econometric Society*, 939–953.

Ghassibe, M., & Zanetti, F. (2022). State dependence of fiscal multipliers: the source of fluctuations matters. *Journal of Monetary Economics*.

Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, *110*(8), 2586–2624.

Goodman, S., & Volz, A. H. (2019). Attendance spillovers between public and for-profit colleges: Evidence from statewide variation in appropriations for higher education. *Education Finance and Policy*, *0*(ja), 1-67.

Hasna, Z. (2021). *The grass is actually greener on the other side: Evidence on green multipliers from the united states.*

Hausman, J. A., Newey, W. K., Woutersen, T., Chao, J. C., & Swanson, N. R. (2012). Instrumental variable estimation with heteroskedasticity and many instruments. *Quantitative Economics*, *3*(2), 211-255.

Ilzetzki, E., Mendoza, E. G., & Végh, C. A. (2013). How big (small?) are fiscal multipliers? *Journal of monetary economics*, *60*(2), 239–254.

Jaeger, D. A., Ruist, J., & Stuhler, J. (2018). *Shift-share instruments and the impact of immigration* (Working Paper No. 24285). National Bureau of Economic Research.

Jappelli, T., & Pistaferri, L. (2010). The consumption response to income changes. *Annual Review of Economics*, *2*(1), 479-506.

Jappelli, T., & Pistaferri, L. (2014). Fiscal policy and mpc heterogeneity. *American Economic Journal: Macroeconomics*, *6*(4), 107-36.

Johnson, D. S., Parker, J. A., & Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, *96*(5), 1589-1610.

Jordà, Ò., & Taylor, A. M. (2016). The time for austerity: estimating the average treatment effect of fiscal policy. *The Economic Journal*, *126*(590), 219–255.

Kelchen, R. (2019). An empirical examination of the bennett hypothesis in law school prices. *Economics of Education Review*, *73*, 101915.

Kelchen, R. (2020, May). Does the Bennett Hypothesis Hold in Professional Education? An Empirical Analysis. *Research in Higher Education*, *61*(3), 357-382. Retrieved from https://ideas.repec.org/a/spr/reihed/v61y2020i3d10.1007_s11162-019-09557-9.html doi: 10.1007/s11162-019-09557-

Lucca, D., Nadauld, T., & Shen, K. (2019). Credit supply and the rise in college tuition: Evidence from the expansion in federal student aid programs. *Review of Financial Studies*, *32*(2), 423-466.

Manuelli, R. E., & Seshadri, A. (2014). Human capital and the wealth of nations. *American Economic Review*, *104*(9), 2736-62.

Marx, B. M., & Turner, L. J. (2018). Borrowing trouble? human capital investment with opt-in costs and implications for the effectiveness of grant aid. *American Economic Journal: Applied Economics*, *10*(2), 163-201.

McCrory, P. B. (2020, July). *Tradable Spillovers of Fiscal Policy: Evidence from the 2009 Recovery Act* (Institute for Research on Labor and Employment, Working Paper Series No. qt04n482qf). Institute of Industrial Relations, UC Berkeley. Retrieved from https://ideas.repec.org/p/cdl/indrel/qt04n482qf.html

Murphy, D. (2017). Excess capacity in a fixed-cost economy. *European Economic Review*, *91*, 245–260.

Nakamura, E., & Steinsson, J. (2014, March). Fiscal Stimulus in a Monetary Union: Evidence from US Regions. *American Economic Review*, *104*(3), 753-92.

Olea, J. L. M., & Pflueger, C. (2013). A robust test for weak instruments. *Journal of Business & Economic Statistics*, *31*(3), 358–369.

Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, *37*(2), 187–204.

Parker, J. A., Souleles, N. S., Johnson, D. S., & McClelland, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, *103*(6), 2530-53.

Ramey, V. A., & Zubairy, S. (2018). Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data. *Journal of Political Economy*, *126*(2).

Rizzo, M., & Ehrenberg, R. G. (2004). Resident and Nonresident Tuition and Enrollment at Flagship State Universities. In *College Choices: The Economics of Where to Go, When to Go, and How to Pay For It* (p. 303-354). National Bureau of Economic Research, Inc.

Sargan, J. D. (1958). The estimation of economic relationships using instrumental variables. *Econometrica*, *26*(3), 393–415.

Shoag, D. (2013). Using state pension shocks to estimate fiscal multipliers since the great recession. *American Economic Review*, *103*(3), 121-24.

Suárez Serrato, J. C., & Wingender, P. (2016). *Estimating local fiscal multipliers* (Working Paper No. 22425). National Bureau of Economic Research.

Tenreyro, S., & Thwaites, G. (2016). Pushing on a string: Us monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics*, *8*(4), 43-74.

Turner, L. J. (2017). *The Economic Incidence of Federal Student Grant Aid* (mimeo). University of Maryland.

Webber, D. A. (2017). State divestment and tuition at public institutions. *Economics of Education Review*. doi: 10.1016/j.econedurev.2017

Appendix A

Appendix to Chapter 1: Human Capital and Growth: The Role of high-skilled labor Concentration

A.1 Romer-based Growth

In this section, I derive the Romer-based growth rate from an increase in human capital supply. Since the goal here is not a full derivation of the model, I stick to the aspects that matter to us, primarily the R&D production function. I will use a similar notation to Jones (1995), meaning that the parameters here are not related to the ones used in the main text.

In Romer (1990), the production function of the final good takes the form:

$$Y = L_Y^{1-\alpha} \sum_{j=1}^{A} x_j^{\alpha} \tag{A.1}$$

where Y is output, L_Y is labor used in the production of the final good, x_j are intermediate goods, α is a constant, and A is the number of intermediate goods. The latter can be thought of as the number of product ideas in the economy.

The total number of intermediate goods evolves according to:

$$\dot{A} = \gamma L_A^{\lambda} A^{\phi} \tag{A.2}$$

where γ, λ and ϕ are constants, and L_A is the number of workers engaged in innovation. The labor market clearing condition is, then, $L = L_A + L_Y$, where L is total labor supply.

In steady state, the share of labor employed in R&D is constant, i.e. $L_A = s_A L$. From Equation A.2 we can write the growth rate of product variety in this model as:

$$g_{Romer} = \frac{\dot{A}}{A} = \gamma (s_a L)^{\lambda} A^{\phi-1} \tag{A.3}$$

In the original model, this growth rate is constant implying that the time derivative of the right-hand side of Equation A.3 is zero. That is:

$$\lambda \frac{\dot{L}}{L} + (\phi - 1)\frac{\dot{A}}{A} = 0 \tag{A.4}$$

In our case, we assume that s_A grows constantly for a period of time as the economy moves toward a new steady state where the labor share employed in innovation is s'_A, $s'_A > s_A$. During this transition, we can write the change in the growth rate as:

$$\frac{\dot{g}_{Romer}}{g_{Romer}} = \lambda\left(\frac{\dot{s}_A}{s_A} + \frac{\dot{L}}{L}\right) + (\phi - 1)\frac{\dot{A}}{A} \tag{A.5}$$

where we now take into account that the share s_A is changing.

If we make the assumption that, throughout the transition, the economy moves between steady states, we can substitute Equation A.4 into Equation A.5 to get that:

$$\frac{\dot{g}_{Romer}}{g_{Romer}} = \lambda\left(\frac{\dot{s}_A}{s_A}\right) \tag{A.6}$$

Hence, the change in the growth rate is only due to changes in the high-skill labor share. We then use Equation A.6 to calculate the Romer-based expected growth rate from a change in high-skill supply. This is done in Figure A20 for $\lambda = 0.435$ (Pessoa, 2005).

We assume here that the only change relative to the original model is the increase in the high-skill labor share. Naturally, this increase can be countered by changes to parameters γ, λ, and ϕ. For example, a reduction in innovation productivity (i.e. lower γ) or an increase in concavity (i.e. lower λ) could lead to a constant growth rate. As such, the values in Figure A20 should be interpreted as a measure of *potential growth rate*.

A.2 Monte Carlo Simulations on Identification with Endogenous Controls

In the baseline estimation in Section 3.3, we control for non-high-skill hiring (and a proxy for capital formation) to make sure that the exclusion restriction on the SSIV holds, i.e. that any low-skill hiring induced by the SSIV has no significant effect on local GDP growth. This, in turn, might introduce bias if non-high-skill hiring is happening as a result of additional high-skill hiring. This is known as a "bad control" problem. In this section, I show that we can still identify the coefficients of interest when we add endogenous controls and instrument them with the extra IVs.

We start by defining the data-generating process. There are four IVs (Z_{1-4}), three endogenous variables (H_1, H_2, and L_1), and a dependent variable Y, all of which relate as:

$$
\begin{aligned}
H_1 &= \beta_0 + \beta_1 Z_1 + \beta_2 Z_2 + \beta_3 Z_3 + \beta_4 Z_4 + v \\
H_2 &= \gamma_0 + \gamma_1 Z_1 + \gamma_2 Z_2 + \gamma_3 Z_3 + \gamma_4 Z_4 + w \\
L_1 &= \alpha_0 + \alpha_1 Z_1 + \alpha_2 Z_2 + \alpha_3 Z_3 + \alpha_4 Z_4 + u \\
L_2 &= \delta_0 + \delta_1 H_1 + \delta_2 H_2 + \nu \\
Y &= \xi_0 + \xi_1 L_1 + \xi_2 H_1 + \xi_3 H_2 + +\xi_4 L_2 + \epsilon
\end{aligned}
\tag{A.7}
$$

where (u, v, w, ν, ϵ) are error terms. Equation A.7 can be understood as follows. Instruments Z_{1-4} generate variation in both low-skill hiring (L_1) and high-skill concentration, the latter split between low (H_1) and high (H_2) skill concentration places. Low-skill hiring (L_2) can also happen due to substitutability or complementarity with high-skill hiring. Finally, both low- and high-skill workers contribute to output (Y). We are interested in identifying ξ_2 and ξ_3.

We then match all parameter moments to their estimated values in Section 3.3 and we draw 100,000 joint observations of the IV and error term sets matching their empirical distributions, in particular their in-sample covariance structure. Finally, we estimate ξ_2 and ξ_3 using 2SLS in three different scenarios: one where we fix δ_{0-2} and we vary ξ_4/ξ_1, one where we set $\xi_1 = \xi_4$ and we vary δ_1, and finally one where we vary both δ_1 and the ξ_4/ξ_1 ratio. The idea is to assess identification as we vary the intensity of the "bad control" channel with respect to both the effect of high-skill concentration on low-skill hiring (δ_{0-2}) and the effect of the change in low-skill hiring on growth (ξ_1 and ξ_4). If identification fails, it is important to determine the range of parameter values for which it happens. Ideally, results will show that we can identify the parameters of interest for any realistic range of the intensity of the unobserved "bad control" channel.

I report results in Figure A32 for $\hat{\xi}_2$.[128] On the top-left plot, we fix the intensity of the effect of high-skill concentration on low-skill hiring though we increase the effect of the additional low-skill hiring on growth. On the top-right one, we set $\xi_1 = \xi_4$ and we increase the effect of high-skill concentration on low-skill hiring. Finally, in the bottom plot we increase both channels simultaneously. Importantly, there are two takeaways from this exercise. First, identification starts to weaken as the "bad control" channel becomes more significant in magnitude. We see this in the increasing distance between the horizontal lines representing the parameter's true value and the point estimates. However, unless both the change in low-skill hiring due to high-skill concentration and its effect on growth relative to other low-skilled workers are quite large the parameter of interest is identifiable. A value above 8 for both δ_1 and ξ_4/ξ_1 would imply an implausibly large high-to-low-skill elasticity and that low-skilled workers hired through this

[128]Results are similar for the channel through $\hat{\xi}_3$.

channel are more productive than other low-skilled workers by almost an order of magnitude.

As such, the approach taken in Section 3.3 of controlling for non-high-skill hiring (and analogously for the capital proxy) does not seem to introduce a significant bias due to a "bad control" problem under reasonable values for its intensity while allowing us to identify the effect of high-skill concentration on growth.

A.3 DiD Falsification Tests and Robustness Checks

We now assess both the identification assumption and the robustness of our difference-in-differences results on the effect of an increase in skill supply on local economic growth.

We start with the identification assumption. The estimation relies on the choice of municipality for a new college being as-good-as-random with respect to local growth.[129] While it is reassuring to find no significant pre-trends in Figure 1.2, we can check the data for further evidence. First, I show in Figure A33 that we also do not observe significant pre-trends in the local stock of formal employees.[130] Second, Figure A34 shows that we also fail to reject the parallel trends hypothesis on both the population share of the graduating college cohort and the difference between the number of new high-school graduates and the incoming college cohort. Both represent different ways to gauge local demand for college education. The lack of pre-trends corroborates our intuition that local demand was not a major factor in determining where to open a new college as supply was severely constrained prior to 1996 and had yet to catch up by the end of my sample.

We can also check the robustness of the identification assumption to changes in the sample. The intuition behind this exercise is that if results remain robust in settings where threats to identification are lower, then we have evidence that such threats are not driving our results. I start by showing in Figure A35 the estimation results using the matched subsample where treated and control observations are matched on observables. Results on the difference $\beta_{2,k} - \beta_{1,k}$ are similar to the baseline, which is evidence that the differences in covariates reported in Table A.2 are uncorrelated with local GDP growth. In absolute levels, however, results using the matched sample are overall higher than the ones in Figure 1.2 as we now do not observe a significant decline in growth in places with higher skill concentration. However, this is explained in the model I introduce in Section 2.5.

[129] Note that for results on the difference between high and low skill concentration places, confounders would need to correlate not only with local growth but also with local skill concentration to be able to affect results, to the extent that if both coefficients are equally biased, the difference cancels out the bias. As such, the identification strategy for the difference in coefficients is stronger.

[130] While we do not observe negative and significant estimates on employment, this is likely due to shifts between formal and informal sectors.

The reason for a lack of decline in growth is that the matched sample has a lower average high-skill concentration level. Hence, the increase in skill concentration from more human capital is not large enough, in this sample, to induce a decline in growth.[131]

Finally, we can assess our identification strategy using different control groups. As in the previous exercise, if results are robust in a setting where the identification assumption is slightly different, we have evidence that our assumption is valid. I do this in two ways by switching the untreated control group with either the last-treated or the not-yet-treated observations, both of which have been shown to provide valid comparison groups (Sun & Abraham, 2021, Callaway & Sant'Anna, 2021). In both cases, treated and control municipalities receive a new college at some point in my sample. Hence, the identification assumption is now on the timing of receiving a new college such that treatment assignment between early- and later-treated municipalities looks as-good-as-random.[132] This is likely since groups look similar on observables as shown in Table A.8 for last treated and not-yet treated. Anecdotal evidence also corroborates the assumption on the timing as new colleges take, on average, many years to be created as founders need government approval, appropriate facilities and staff, and a procedure to formally select students (e.g. exams). All of these steps can take different amounts of time for reasons that are unrelated to local growth.

Results using last-treated and not-yet-treated groups are similar to baseline estimates. Starting with the former, I show in Figures A36 and A37 results for the individual coefficients and the difference in effect between places with high and low skill concentration. Although noisier, last-treated estimates are in line with baseline ones. I then show in Figure A38 results using the not-yet treated as control. As the estimation is noisier and requires calculating estimates between different municipality cohorts, I make two important modifications. First, I set the threshold between low and high to the 20^{th} percentile to reduce standard errors. Second, I set the treatment period $t = 0$ to the period when the first student cohort is expected to graduate. This increases the number of not-yet-treated observations, though both changes render estimate-by-estimate comparisons with the baseline estimation difficult. Nonetheless, results look qualitatively similar, and we can visually infer that the difference between coefficients in high and low skill concentration places is significant and negative. Hence, both last-treated and not-yet-treated estimations show evidence in support of the validity of the identification assumption.

Moreover, I show robustness of results to changes in the sample and in the specification. In the former, estimates are robust to restricting the sample to municipalities that do not have colleges, either in all periods (control) or in the pre-

[131] C.f. Section 1.5 for a detailed explanation of the mechanism behind these results.

[132] Similar to the assumption made in Nimier-David (2023).

treatment period (treated), as shown in Figure A39. This is reassuring as we might worry that including places with pre-existing colleges could bias results if these colleges expand their student intake in response to college creation elsewhere. In the latter, I first show that results remain the same if we increase or decrease the threshold p that defines a high or low high-skill concentration municipality, as shown in Figures A40 and A41 for $p = 12\%$ and $p = 17\%$, respectively. I then show in Figure A42 that results remain unchanged if we add to the specification dummies for the leads and lags of municipalities that reported receiving new colleges twice or three times. Finally, results are robust to running a weighted specification where we weight by the log of local population, as I show in Figure A43. This is evidence that results are not being driven by the direct economic effect of new colleges on local GDP, mainly because new colleges represent little of the local economic activity. Importantly, in all cases we find evidence of no significant pre-trends.[133]

Results are also unchanged if we use estimators robust to heterogeneous treatment effects and non-binary treatments. The literature on difference-in-differences estimators has shown that estimates can be biased in the presence of heterogeneity in treatment effects (c.f. Roth et al., 2023 for a summary). Moreover, in our context it is possible that more than one college is created within a single municipality over time, which ultimately constitutes multiple treatments. We can, then, assess whether alternative estimators that take into account such cases give different results. For the case of heterogeneous effects, I show in Figure A45 results using the estimator proposed in Sun & Abraham (2021) that restrict the control group to never-treated units, avoiding the issue of "forbidden comparisons" which may bias estimates. Reassuringly, results remain indistinguishable from baseline ones. As for non-binary treatments, I show in Figure A46 estimates using the estimator proposed in de Chaisemartin & D'Haultfœuille (2024) which aggregates the treatment effect of municipalities experiencing different treatment paths. Once again, estimates are quite similar to baseline ones. Finally, results remain robust to adding local-level controls and using a robust estimator as shown in Figure A47, which is evidence that potential imbalances between treated and control groups are not driving results.

Finally, results on high-skill concentration are robust to different specifications, changes to the sample, and alternative estimators. First, I show that estimates remain similar if we use the matched placebo group as our control group, as shown in Figure A48, which is evidence that any imbalance between treated and control groups is not affecting estimates. Results are also robust to restricting the sample to municipalities that do not have colleges, at all (control) or prior

[133]Results are also robust to removing exports from local GDP as I show in Figure A44. This is evidence that results are not being driven by inter-municipality firm competition, in line with the local labor market mechanism I propose in Section 2.5.

to treatment, as shown in Figure A49. Figure A50 shows results are robust to adding controls for the leads and lags of municipalities that reported receiving new colleges twice or three times, while Figure A51 shows robustness to using a HHI-based measure of local high-skill concentration. Finally, results remain unchanged if we use instead the robust estimator proposed in Sun & Abraham (2021), with or without local-level controls, as shown in Figures A52 and A53, or the estimator proposed in de Chaisemartin & D'Haultfœuille (2024), as shown in Figure A54.

A.4 Shift-Share Design: Skill Premium

We can leverage the SSIV design in Section 1.3.2 to study the effect of high-skill concentration at large firms on the skill premium. While the latter is not part of my key results, it is useful as additional validation for the model in Section 2.5 as the increase in skill concentration can have secondary effects other than on growth.

We proceed by using the same SSIV design. In particular, I use a specification similar to Equation 1.3 where I replace per-capita GDP growth with local skill premium, here defined as the ratio between wages for high-skill and non-high-skill workers within a municipality. I then show estimation results in Table A.9 which follows the same framework as Table 1.1. I also instrument high-skill concentration with the public loans SSIV described in Section 1.3.2. However, an issue with this estimation is that the SSIV calculated using loans to small firms gets weaker when high-skill concentration is high. This is expected as small firms play a less significant role in the local economy when concentration at large firms is high. Because of this, when estimating the effect on skill premium this leads to a failed overidentification test due to spurious coefficients from $B_{i,t-2,small}$. As such, I remove $B_{i,t-2,small}\mathbb{1}\{\{HSConc_{i,t-1} > p\} = 1\}$ from the set of IVs. I show evidence in Section A.5 that this is caused by observations with high levels of high-skill concentration and that results are robust to instrumenting with $B_{i,t-2,small}$ without an interaction term.

Results for the skill premium are similar to those for GDP growth as we see a non-monotonic relationship between skill concentration and the skill premium. While an increase in skill concentration leads to a higher skill premium at low levels of skill concentration, further increases to skill concentration reduce the skill premium. Coefficients are significant in all specifications implying results are robust to potential threats to identification. Joint F-statistics are above the usual weak-IV threshold except for Column (7), and the effective F-statistics calculated individually for each regressor of interest are above the critical values. Finally, we do not reject the null for the J-test of overidentification when using more IVs than regressors, providing support for identification.[134]

As in Section 1.3.2, we can gauge the importance of the high-skill concentration channel to movements in aggregate skill premium. As I show in Figure A27, there has been an important decline in the skill premium in Brazil since the early 2000's. We can, then, use our estimates in Column (2) of Table A.9 to assess the relative importance of the high-skill concentration channel. Doing a similar back-of-the-envelope calculation as the one in Section 1.3.2 yields a 0.29 drop in the skill premium over 11 years from the increase in aggregate skill concentration. This decline is quite significant as it can potentially explain the entire aggregate decline in skill premium between 1999 and 2010 when aggregate skill concentration plateaued. As such, the large increase in high-skill concentration has led to a significant decline in skill premium.

This novel non-monotonic result between the skill premium and skill concentration extends the existing literature on monopsony power in the labor market. As shown in Azar et al. (2022) and Schubert et al. (2024), as firms gain more labor market power they push wages down. While this is reasonable for unskilled labor, my results show that high-skill wages react differently to labor market power given R&D competition dynamics.

A.5 SSIV Falsification Tests and Robustness Checks

We now assess both the identification assumption of the SSIV design and the robustness of our non-monotonic results to changes in the baseline specification. Importantly, what interests us is not exactly the stability of point-estimates but whether the coefficient signs are robust to changes in the estimation, i.e. whether our non-monotonicity result is robust.

I first analyze the validity of the shift-approach through falsification tests. While the overidentification tests in Table 1.1 are encouraging, I also test the shock exogeneity assumption using the tests proposed in Borusyak et al. (2021). Though these tests cannot prove instrument validity, passing results strengthen the identification assumption.

I start by showing how much variation we have at the shock level. This is important as the validity assumption requires enough shock-level variation for consistency. Table A.10 shows summary statistics for both shocks $g_{n,t-2}$ and shares $s_{in,t-3}$, split between large and small firms. For the shocks, statistics are weighted by the shares and I also report statistics after residualizing with year fixed-effects. That is, I regress shocks on year fixed-effects while weighting with shares, which allows us to gauge within-period variation. In addition, I report the effective

[134]Differently from Table 1.1 I do not run a LIML specification when using all endogenous regressors since, in this case, we are dealing with a just-identified specification.

sample size measured as the inverse of the share HHI, i.e. $1/\sum_{n,t} s_{n,t-3}$ where $s_{n,t-3} = \sum_i s_{in,t-3}$ are the sector-level shares. This indicates how concentrated sector exposure is and, hence, measures whether we have enough sector-level variation for asymptotic validity. Borusyak et al. (2021) show that an effective sample of at least 20 provides enough variation for large-sample approximations.

Table A.10 shows that we have significant variation at the sector level. Largest shares for both small and large firms are 1.1% and 5.1% respectively, indicating that no single sector-period has an overweight on the distribution. Shock distributions for both large and small firms look regular and have standard deviations that are larger than their means. Residualizing shock distributions for large and small firms with year fixed-effects has a significant effect on the former as the standard deviation drops by around 50%. However, the effective sample for both large (26) and small (190) local-level shares is above the threshold of 20 which suggests high enough variation.

I then implement falsification tests at the shock level. These consist of regressing sector-level controls and the dependent variable, both taken prior to the realization of shocks, on shocks directly weighting by the shares. Formally, let q_{it} be a control variable used in Equation 1.3. We then run:

$$\bar{q}_{n,t-3} = \beta g_{n,t-2} + \gamma V_{n,t-2} + \epsilon_{n,t} \tag{A.8}$$

where $\bar{q}_{n,t-3} = \frac{\sum_i s_{int} q_{i,t-3}}{\sum_i s_{int}}$ is the exposure-weighted average of $q_{i,t-3}$ and $V_{n,t-2}$ is the set of all controls used in Equation 1.3, including time fixed-effects, except $q_{i,t-2}$.[135] When using the lagged dependent variable on the left-hand side, I replace $q_{i,t-3}$ with $y_{i,t-3}$ (i.e. per-capita GDP growth or skill premium). Finally, I also use data from sector-level surveys at the national level to check for balance between sectors on supply-side parameters, although I can only run this specification on the combined shock to both small and large firms (vs. running it separately for small- and large-firm shocks).

The intuition behind these tests is two-fold. First, in assessing whether there are any significant correlations between shocks and prior observables at the sector level we can look for significant differences between industries exposed to large and small shocks. If we find any, we may potentially worry that our results in Table 1.1 and Table A.9 are biased due to correlations with unobservables even though I control for the observables being tested (Oster, 2019). Second, regressing lagged dependent variables on shocks provides us with a pre-trend test similar to difference-in-differences specifications. A significant shock coefficient could indicate that high-shock sectors were on a different trend relative to low-shock sectors prior to the realization of the shock, posing a threat to identification.

[135] I run the regression at the sector level in order to avoid the clustering issue shown in Adão et al. (2019).

I show results for the falsification and pre-trend tests in Table A.11 for both large and small firm shocks using local-level variables, and in Table A.12 using sector survey data. In the former, Columns (1)-(4) show the pre-trend tests, both for local growth rates and the skill premium, while Columns (5)-(14) show the balance test. All variables have been demeaned and normalized to have unit variance so that coefficients are easier to interpret, and standard errors are clustered at the sector level. All but one coefficient pass the balance test of non-significant results and we observe no pre-trends with respect to GDP growth and skill premium. While the coefficient for large firms when regressing the population share of workers receiving minimum wage or less is statistically significant, the magnitude is small: a one-standard-deviation increase in the shock is associated with a decline in the minimum wage share of around -4% of its standard deviation. This difference between high- and low-shock sectors does not seem large enough to drive results.

As for the sector survey data, I find no significant shock imbalance with respect to supply-side variables. These consist of the growth in net revenues and value added, the ratio of wages, intermediate inputs costs, and fuel and electricity costs to value added, and the share of production workers over total sector employment, measured either at the end of the year or as an yearly average. While sector survey coverage is lower than the one in the RAIS database, I manage to cover most sectors. All coefficients between variables and shocks are not significant, a result we would expect if shocks are indeed as-good-as-randomly assigned to industries each year.[136]

With enough shock-level variation and having passed the falsification tests, the *a priori* assumption of shock exogeneity for my SSIV seems plausible. Although local demand for loans from the BNDES depends on local conditions, changes to the national amount disbursed to different sectors and firm sizes seem exogenous to local- and sector-level parameters. The evidence, then, points to the validity of the SSIV identification strategy.

I now move to the robustness checks of the SSIV. We first assess whether having the IVs interact with $\mathbb{1}\{HSConc_{i,t-1} > p\}$ leads to biased estimates. As $\mathbb{1}\{HSConc_{i,t-1} > p\}$ is a function of the endogenous variable, albeit one with little variation, we may worry that we might be reintroducing endogenous variation into the IV set. I report in Table A.13 results where I use polynomial terms for both high-skill concentration and the SSIVs instead. Although IV relevance is significantly lower, we observe a similar non-monotonic (and concave) relationship between skill concentration and GDP growth. I also report the point-estimates of

[136]Another concern when using SSIVs is that a strong serial correlation of shocks, combined with latent dynamic adjustments of the dependent variable, may bias our results (Jaeger et al., 2018). In our case, $g_{n,t-2,large}$ and $g_{n,t-2,small}$ have low serial correlation: -0.047 and -0.083, respectively. As such, any dynamic bias would not affect results significantly.

the point where the slope changes sign to confirm that they are within the [0,1] domain. As a final check, I report in Table A.14 results from a specification that interacts high-skill concentration with $\mathbb{1}\{HSConc_{i,t-1} > p\}$ but instruments with polynomial terms of the SSIV. Results remain robust though we lose significance of the coefficient at high levels of skill concentration when the IV weakens. Evidence, then, suggests that results are not affected by having the IVs interact with $\mathbb{1}\{HSConc_{i,t-1} > p\}$.

Next, we analyze whether results remain unchanged when we use a narrower definition of high skill. Up to this point, we have considered all workers with some college education as high-skill employees. However, we can narrow down this definition to include only those who actually work in occupations that require high critical thinking. Although this procedure may exclude workers who can potentially do innovation yet are underemployed relative to their capabilities, this narrower classification reinforces the link between our empirical results and innovation dynamics. I show in Tables A.15 and A.16 results on growth and skill premium, respectively, using high critical-thinking workers as defined in Section 2.2. We find similar non-monotonic relationships between high-skill concentration, growth, and skill premium as in the baseline case.

On skill premium, I show evidence that the estimated negative slope is due to a large negative coefficient at an intermediate level of high-skill concentration. As discussed in Section A.4, the estimation using skill premium suffers from a weak IV problem regarding the SSIV for small firms when skill concentration is high. To assess this issue, I run in Table A.17 a specification where instead of splitting the regressor between low and high levels of skill concentration, I do it between low and intermediate levels ($\mathbb{1}\{HSConc_{i,t-1} < p_1\}$ and $\mathbb{1}\{p_1 < HSConc_{i,t-1} < p_2\}$, respectively). I then fix p_1 at the 10^{th} percentile of the skill concentration distribution and vary p_2 for different estimations. Results show that the negative slope at high levels of concentration is due to a more negative coefficient at mid-levels (between -5.0 and -2.5 vs. -1.4 in the baseline). This explains why the overidentification test fails: as we increase the p_2 threshold, the SSIV for small firms becomes weaker, which can be seen in the joint F-statistic jumping from 38.3 to 55.1 between Columns 4 and 5 once we remove $B_{i,t-2,small}(\mathbb{1}\{p_1 < HSConc_{i,t-1} < p_2\} = 1)$ from the IV set. The IVs, then, seem to capture the heterogeneity in the slope and the failure of the test does not seem to be due to a violation of IV validity. I show further evidence of this in Table A.18 where I use $B_{i,t-2,small}$ as an IV without interaction terms. The non-monotonic result is robust and we cannot reject the null hypothesis of the overidentification test.

We also obtain similar results when we run the specification on a subsample restricted to the non-tradable sector.[137] This is shown in Tables A.19 and A.20 for GDP growth and skill premium, respectively. The fact that we observe similar results is reassuring as tradable firms can engage in product competition with companies outside their municipality. As I show in Section 2.5, the mechanism I propose to explain the causal results involves the strategic interaction between a leading and a follower firm competing through innovation within the same labor market. Since I run my baseline specification at the municipality level, I also capture tradable firms competing out-of-municipality. By restricting the analysis to the non-tradable sector, particularly in a context where the data is at the establishment level, results can be more directly linked to my mechanism.

Finally, results are robust to other changes in the specification. We still observe non-monotonicity and significant results between skill concentration, growth, and skill premium if we run a weighted regression weighting by the log of local population (Tables A.21 and A.22). Results are also robust to lagging the SSIV exposure shares one additional period (Tables A.23 and A.24, the latter uses the low and mid-level thresholds to increase IV relevance). Finally, I show in Table A.25 that while point-estimates are sensitive to the choice of threshold p, our finding on the non-monotonic relationship does not depend on a particular value of p as long as the cut-off is near the point where the slope changes in the relationship between high-skill concentration and growth (or skill premium). In our case, results show that the change in slope occurs around the range of 15% and 30%.

A.6 low-skilled labor Supply Assumption, Proofs, and Derivations

In this section, I discuss the assumption of perfect elastic low-skilled labor supply and provide the necessary proofs and derivations for the model derived in Section 2.5.

Perfect elastic low-skilled labor supply: As in Aghion et al. (2001), I assume in the model that low-skilled labor supply is perfectly elastic. Along with the simplification, this assumption, in fixing the low-skill wage, leads to a straightforward link between variations in the skill premium and what is happening in the high-skilled labor market. As such, the non-monotonic result on skill premium is only being driven by changes in high-skill wages.

I assess this assumption empirically in Table A.26 using the 2SLS specification in Section A.4. We see that changes in skill concentration do not lead to significant changes in low-skill wages. Importantly, the large standard errors in Columns

[137] I define the non-tradable sector as any sector outside agriculture and manufacturing. Although some service subsectors can be deemed tradable, there is no local-level GDP data by subsector.

(1) and (2) are due to the low in-sample variance of low-skill wages. Moreover, Columns (3) and (4) show a positive relationship between high-skill concentration and low-skill hiring when the former is low. These results are in-line with assuming that low-skill wages are fixed while low-skilled labor supply adjusts.

In reality, low-skilled labor supply is elastic though not infinitely so. If we had instead assumed a finite labor elasticity, some of the conclusions from the infinitely inelastic case studied in Aghion et al. (1997) would apply. We can expect low-skill wages to follow changes in aggregate demand, which in turn are related to profits, high-skill hiring by the non-innovative sector, and low-skill hiring. Demand, however, varies little across different gap levels as lower low-skill hiring due to higher relative productivity is compensated by higher profits. This implies that the non-monotonic shape of the skill premium curve as a function of high-skill concentration would not change significantly. I assess this point further in Section A.7 by running the model under different labor market assumptions.

Optimal R&D investment and high-skilled labor demand: The firm's optimal R&D choice can be derived from the maximization problem in Equation 1.15. Starting with λ_s, the first-order condition for the leader's problem results in (analogously for λ_{-s} and λ_0):

$$\lambda_s = \frac{A_\lambda (J_{s+1} - J_s)}{\rho} \tag{A.9}$$

Similarly for labor demand $l_{s,H}$, which requires solving:

$$w_{s,HS} = A_l \alpha l_{s,HS}^{\alpha-1} (J_{s+1} - J_s) - \kappa l_{s,HS} \left(\frac{\delta}{B\theta_s^{-\varphi} u_s} \right)^2 \tag{A.10}$$

where we used Equation 1.16 to replace for v_s as a function of $l_{s,HS}$. To solve for labor demand, we substitute for the wage rate using Equation 1.19.

High-skill wage: To get Equation 1.19, we multiply Equation 1.18 by r and replace W_s and U_s with their definitions along with the Nash bargaining solution:

$$r\xi S_s = rW_s - rU_s = w_{s,HS} - \delta(\xi S_s) - b - B\theta_s^{1-\varphi} \xi v_s S_s - B\theta_s^{1-\varphi} \xi v_{-s} S_{-s} \tag{A.11}$$

where we used the fact that the match surplus for the non-innovative firm is zero. We can then rearrange terms to get Equation 1.19.

Growth rate: The derivation follows Acemoglu & Akcigit (2012). Start with a single sector at gap s. Since $y_s = \gamma_s l_{s,HS}$ and $l_{s,HS}$ is constant in steady state, y_s grows at the same rate as γ_s, i.e.:

$$g_s = \lim_{\Delta t \to 0} \frac{\ln \gamma_s (t + \Delta t) - \ln \gamma_s (t)}{\Delta t} \tag{A.12}$$

Given Bertrand competition, we only need to look at the leader's production for $s > 1$ and the neck-and-neck case, though we will link it back to the follower's case at the end. Note that at any interval Δt, in expectation, the leader innovates at a rate $\eta_s \Delta t + o(\Delta t)$ while neck-and-neck firms innovate at a rate $2\eta_0 \Delta t + o(\Delta t)$. Each innovative step increases γ_s by γ. Then:

$$\ln\gamma_s(t+\Delta t) = \ln\gamma_s(t) + \ln\gamma\left[\mathbb{1}_{s=0} 2\eta_0 \Delta t + \mathbb{1}_{s>0}\eta_s \Delta t\right] \tag{A.13}$$

Replacing Equation A.13 into A.12 results in Equation 1.23.

The final step is to notice that aggregate growth is the weighted average of all g_s by the sector share μ_s. Notice also that, in steady state, the technological frontier (i.e. leaders and neck-and-neck firms) and followers must grow at the same rate, implying that:

$$g = \ln(\gamma)\left(\sum_{s=1}^{\infty} \mu_s \eta_s + 2\mu_0 \eta_0\right) = \ln(\gamma) \sum_{s=1}^{\infty} \mu_s [\eta_{-s} + h_l l_{-s,HS}^{\alpha} + h_c] \tag{A.14}$$

A.7 high-skilled labor Supply Assumption

In the model presented in Section 2.5, high-skilled labor is hired through search. In this section, I show how results change if we remove labor frictions, providing intuition for their importance. I then show the importance of labor frictions in matching the skill premium.

We first consider a version of the model with high-skilled labor which is closest to Aghion et al. (2001). That is, we start with a similar set-up to the one in Section 2.5, i.e. a step-by-step growth model with two firms, and we allow for two types of labor: high and low skilled, each being paid at wage w_k, $k = \{HS, LS\}$. As in Section 2.5, high-skilled labor is used in R&D production. Importantly, there are no search frictions, hence no unemployment. This implies that the labor market for high-skill workers has to clear as follows:

$$L_{HS} = l_{s,HS} + l_{-s,HS} \tag{A.15}$$

Notice how Equation A.15 implies that total high-skill hiring (L_{HS}) is invariant with respect to the gap s. This aspect of this simple model affects results significantly. To see that, first realize that in this model the leading firm still has relatively higher incentives to hire skilled workers than the follower as s increases. This is because we have not made any changes to firm competition or how innovation works. As such, starting at $s = 0$, as s increases high-skill concentration at the leader still goes up. However, once incentives to innovate decline as the leader is too far ahead, the leading firm cannot shed labor as the skilled labor market has to clear and the follower is even less willing than the leader to hire. As such, $l_{s,HS}$ is necessarily a monotonically increasing function of the gap s. I show this in Figure A55 where I plot the firms' value functions and input decisions as a function of the gap s.[138] We observe that the leading firm's high-skilled labor hiring increases monotonically in s and stays near total supply L_{HS} for s large enough.

This change in the high-skilled labor market clearing also affects results on growth. I show this in Figure A56. Since the leading firm cannot lower its high-skill hiring, growth does not decline at a high enough level of the gap. Key to this is that labor cannot adjust downward as it does in the baseline model. This highlights the importance of allowing for unemployment in the model, which I achieve with search frictions.[139]

Labor frictions are important not only for results on growth but also to match the skill premium. To show this clearly, we can take results from Section 2.5 and counterfactually change the assumption on high-skilled labor.

We start with a similar setup to the one in Section 2.5, i.e. a step-by-step growth model with two types of labor. However, we assume now that high-skilled labor can be hired without frictions and is supplied inelastically. All firms pay a single wage rate, conditional on the gap s, which is set so that the labor market clears. Taking the first-order condition with respect to labor demand in Equations 1.15 and 1.17 results in:

$$\begin{aligned} l_{s,HS} &= \left(\frac{A_l\alpha(J_{s+1}-J_s)}{w_{s,HS}}\right)^{\frac{1}{1-\alpha}}, \; l_{-s,HS} = \left(\frac{A_l\alpha(J_{-s+1}-J_{-s})}{w_{s,HS}}\right)^{\frac{1}{1-\alpha}} \\ l_{0,HS} &= \left(\frac{A_l\alpha(J_1-J_0)}{w_{0,HS}}\right)^{\frac{1}{1-\alpha}}, \; l_{o,HS,s} = (1-\nu)\frac{D_s}{w_{s,H}} \end{aligned} \tag{A.16}$$

Intuitively, Equation A.16 says that high-skilled labor is paid at its marginal product which, for the innovative sector, is the marginal benefit, in expectation, from a successful R&D innovation that increases the gap by 1. Importantly, in this case total labor demand does not vary with s as there are no restrictions in the labor market. We can then solve Equation A.16 using, where needed, estimates from the baseline estimation, including the estimates for the firms' value functions.[140] I plot results for the wage premium in Figure A57. Comparing the wage premium curve between the inelastic and labor search ("baseline") cases, we observe that we achieve non-monotonicity in both cases as firms move from the region of intense competition to the one where the lazy monopolist effect kicks

[138] To solve the model, we normalize low-skilled labor supply to 1. We only have 5 parameters to estimate: $\{\gamma, \rho, A_l, A_\lambda, h_c\}$. I estimate those using 5 moments: average real GDP per capita growth rate, average high-skill concentration at large firms, average firm profitability, R&D share of sales, and share of markets where high-skill concentration is below or equal to 50%.

[139] Results are similar if we go a step forward and add an outside sector as in Section 2.5. Though this sector can in theory absorb high-skilled labor once the leader faces lower incentives to hire, this is limited by the outside sector's demand. Results for this version of the model are available upon request.

[140] To make results comparable, I set total labor supply L_{HS} to average employment in the baseline model.

in. This is because we are using the results from the model with search frictions, which allow high-skilled labor hiring to adjust. However, even when we use the baseline value function values there are two shortcomings of the inelastic case. First, wage premium can be (and is) below one which does not make empirical sense. Second, since at large s both firms have low incentives to invest in R&D effort, the high-skill wage approaches zero.[141] As such, the wage premium also goes to zero as s grows large which also does not reflect reality.

Another approach would be to make high-skilled labor supply elastic by adding labor disutility in the utility function of workers. The shape of the wage premium curve in this case would depend on the exact functional form of the utility function. In cases where income and substitution effects cancel out, the end result is a constant L_{HS} and a simple level shift from the inelastic case. However, if labor supply is a monotonic function of the wage the shape of the wage premium curve changes slightly though it is still dictated by the change in the marginal benefit of R&D. I show one parametrization of this case in Figure A57 where I set the Frisch elasticity ζ to 0.5 and the disutility scalar is set to match the average unemployment rate in the baseline case.[142] In this scenario, as the disutility from working at different companies is the same we have to impose a geographical restriction to the labor market where firms hire from within different areas of a municipality.[143] The resulting wage premium is also below 1 and tends to zero as s grows.

It is clear from Figures A56 and A57, then, that the assumption of search frictions is helpful, both qualitatively and quantitatively. Qualitatively, having unemployment allows firms to adjust high-skilled labor downward, in tandem with changes in incentives to hire as a function of the gap s. Quantitatively, the baseline scenario can capture empirical trends, particularly when it comes to the skill premium, and does not rely on a particular shape of the disutility of labor. Moreover, the elastic case still requires an assumption on labor mobility as firms pay different wages. Finally, we also require some restriction on the labor market to capture unemployment. Search frictions are, then, a natural choice.

A.8 GMM Estimation Moments

In this section, I go over each empirical moment and the model mapping of the GMM estimation in detail:

[141] In the limit, for s reaching infinity high-skill wage is effectively zero as the value function becomes flat.

[142] An example of a case where $l_{s,HS}$ does not depend on the wage is $U(x,L) = \ln(x) - K\frac{l_{s,HS}^{1+\frac{1}{\zeta}}}{1+\frac{1}{\zeta}}$ which delivers $l_{s,HS} = (\frac{1}{K})^{\frac{\zeta}{1+\zeta}}$. In the case where $l_{s,H}$ depends on the wage, I abstract from the exact functional form and set $l_{s,HS} = (\frac{w_{s,HS}}{K})^{\zeta}$ where $K = 3.62$.

[143] Otherwise, all labor would work at the leading firm, who pays a higher wage.

i) *Real GDP growth*: data comes from IBGE for the 1999-2017 period. I show how to calculate the aggregate growth rate in the model in Section 2.5;

ii) *Skill Premium*: average calculated using in-sample data where I weight observations by the number of workers. In the model, skill premium is the labor-weighted average of high-skilled labor in firms i, $-i$, and o;

iii) *Labor Market Tightness*: data comes from the Catho-Fipe series which provides an indexed time series. The average level between 2004-2017 is calculated using a nominal value reported in October 2013, allowing me to de-index the data;

iv) *High-Skill Wage at Non-Large Firms*: calculated using in-sample data weighted by the number of workers. Wage data is provided as a multiple of the yearly minimum wage. To recover annual wage rates, I multiply the wage multiple by the monthly minimum wage rate from IBGE. I then multiply it by 13 to get annual wages, taking into account the mandatory end-of-the-year bonus. To avoid an empirical moment with a large order of magnitude, I divide the annual wage by 100,000;

v) *High-Skill Concentration*: average calculated using in-sample data. I show how to calculate high-skill concentration in the model in Section 2.5;

vi) *Firm Profitability*: calculated as the sum of returns on riskless assets and the equity risk premium (ERP). For the former, I use $r = 8\%$ (Section 2.5). For the ERP, I use the value provided in Damodaran (2023) (July/23 edition) for Brazil, i.e. $ERP = 9.57\%$. This value, however, is post-tax. I convert it to the pre-tax level using an effective corporate tax rate of 18.08% (Pires et al., 2023). In the model, firm profitability is defined in Equation 2.12, which must be averaged using the sector shares μ_s;

vii) *R&D Investment-to-Sales Ratio*: data is from the Survey for Technological Innovation (PINTEC) which is conducted over a period of three years since 1998. I use total spending in internal R&D as my measure of R&D investment though it requires two adjustments. First, I remove government subsidies to R&D which account for around 11% of private spending (Betarelli Junior et al., 2020). Second, the survey-measured spending includes wages to people employed in R&D activities which should not be taken into account here as we separate investments from labor costs in the model. While the wage share is not measured by the survey, it is known to be substantial as innovation relies heavily on the knowledge capital of high-skill workers. I assume this share to be around two-thirds (67%), in line with the literature and US data (Hall & Lerner, 2010). In the model, I fit this ratio with the share of total R&D investment over aggregate demand for the innovative sector;

viii) *Cost of Hiring per Job*: I estimate this using data from the US where the average cost of hiring per job was \$4,683 in 2021.[144] I then calculate the cost of hiring as a share of the annual average wage (\$67,610 in 2021 according to the Bureau of Labor Statistics). Finally, I estimate the cost for Brazil as proportionate to the number of days that it takes to hire someone (39.6 days in Brazil vs. 23.8 in the US).[145]. In the model, I calculate this as the average vacancy costs share of high-skill wages;

ix) *High-Skill Unemployment*: data comes from IBGE for the period between 2012 and 2017 for people who have attended college though might not have graduated from it. I adjust average unemployment to take into account that some high-skill workers are in the informal market, which is something the model does not account for. According to Veloso et al. (2022), around 25% of workers with 16 or more years of study (equivalent to having a college degree) were informal workers between 1999 and 2017. I assume counterfactually that, in the absence of an informal market, half of currently informal workers would become unemployed (vs. formally employed or becoming inactive). The targeted moment is, then, average high-skill unemployment (6.07%) plus half of those who are informal workers (12.5%);

x) *Share of Markets with Concentration Below 50%*: calculated in-sample after removing markets where concentration is either below 10% or above 90% as those are not properly captured in the model.

Finally, the outside moment ("R&D Worker Share") is calculated using data from PINTEC which reports the number of full-time workers employed in R&D activities and total number of workers for a sample of manufacturing and high-technology firms. In the model, I calculate the R&D worker share as the share of workers employed in R&D at both leader and follower firms.

A.9 Transition Dynamics

In this section, I go over the transition dynamics of an increase in high-skilled labor supply as the one we observe in Brazil between the early 2000's and the late 2010's.

First, notice that I show in Section 3.3 causal evidence of the timing of the effect of an increase in local human capital on both GDP growth and skill concentration. Relative to the time we would expect the first cohort of a new college to graduate, results show a steady increase in skill concentration, which almost doubles after 12 years, and a long-term decline in local growth across municipalities. This is direct evidence that the mechanism I study in this paper fits within the aggregate

[144] Average is from surveys conducted by the Society for Human Resource Management (SHRM).

[145] Data on length of hiring process is from a Glassdoor survey in 2017 available at `https://www.glassdoor.com/research/time-to-hire-in-25-countries`.

growth trend in Brazil which I show in Figure A20: the increase in skill supply since the early 2000s was followed by a growth boom, which subsided around 13 years later and declined further after that point. Although I do not target this timing in my model estimation, it is important to understand whether the model is able to capture it.

To solve for the transition dynamics, I follow Liu et al. (2022) and Chikis et al. (2023). From the production function in Equation 1.7, we have that the logarithm of aggregate output $Y(t)$ can be written as:

$$\ln Y(t) = \sum_s \mu_s(t) \ln y_s(t) = \sum_s \mu_s(t) \ln(\gamma_s(t) l_l(t)) \tag{A.17}$$

where, given Bertrand competition, $y_s(t)$ is the leader's output for $s \geq 1$ and the combined output of both firms when $s = 0$. As such, we can rewrite Equation A.17 as:

$$\ln Y(t) = \sum_s \mu_s(t) \ln(\gamma_F(t)) + \sum_s \mu_s(t) \ln(\gamma^s) + \sum_s \mu_s(t) \ln(l_L(t)) \tag{A.18}$$

where $\gamma_F(t)$ is the follower's productivity. Finally, I assume throughout the transition that each shock to high-skilled labor supply is unanticipated, implying that what takes time to adjust is the distribution μ_s. We then use Equation A.18 to get the variation in time of growth:

$$g(t) = \frac{\dot{Y}(t)}{Y(t)} = \sum_s \mu_s(t) \frac{\dot{\gamma}_F(t)}{\gamma_F(t)} + \sum_s \dot{\mu}_s(t) \ln(\gamma^s l_L(t)) \tag{A.19}$$

where the first term, i.e. the evolution of the follower's technology frontier, is derived in Section A.6 and leads to the steady-state growth rate once μ_s is stable. The second term refers to the change in output composition between the different states s due to temporary changes in the firm distribution.

We can then use Equation A.19 to calculate the transition path of $g(t)$. I assume the economy is in steady state at $L_{HS} = 1$ and I set parameters to the 1999 to 2004 period. I then assume the economy is hit by successive and unanticipated shocks to L_H which raise it linearly from 1 to 2.5 in 18 years.

There are, however, a couple necessary adjustments. As aforementioned, Section 3.3 shows evidence of how long it takes for effects from an increase in high-skilled labor supply to show up. In particular, we see that high-skill concentration increases by around 12% after 12 years since the first cohort graduates from a new college. The same increase in the model is significantly slower. To see this, we have to re-scale the increase in high-skill concentration as the variation in the model is much larger than the one we see empirically. Aggregate skill concentration in the model varies from around 55% to almost 95%, yet Figure A1 shows that, in the data, aggregate skill concentration does not go over 70%.[146] I show in Figure A58 the transition path for skill concentration for both the not-scaled and scaled cases. The latter, which provides a more "apples-to-apples" comparison with the data, shows that the model reaches the 12% increase in skill concentration at around 94 quarters (24 years). This is twice as slow as what we see in the data, which calls for an adjustment of the transition dynamics.

Hence, I propose two adjustments to the transition path of high-skilled labor. First, I assume that the follower takes longer to adjust its optimal hiring, relative to the leader, after a supply shock to aggregate high-skilled labor. Second, that the leader experiences an initial overshooting of its hiring which dissipates with time.[147] These adjustments speed up the transition path for the high-skill concentration as shown in Figure A59, where it now takes around 14.5 years for the model to produce a 12% increase in skill concentration, a value much closer to what we observe empirically.

We can finally assess the transition path of $g(t)$. I show results in Figure A60 for two time frames. A couple of things are noteworthy. First, as noted in Chikis et al. (2023), the full transition path of step-by-step models can be quite long as the growth rate is still slightly adjusting after almost 200 years. However, most of the action happens in the first 40 years. Second, we observe an initial boom in growth as firms adjust their hiring to the new high-skilled labor supply while the distribution of firms μ_s still has a heavier mass at low levels of s. As the leader innovates relatively more, however, a higher share of sectors moves to higher levels of the gap, bringing growth down. This initial boom period lasts around 15.6 years, in line with what we observe in the data as shown in Figure A20: the increase in high-skill supply is followed by a period of high growth between the early 2000s and the mid-2010s. Growth, then, declines relative to the initial period, which the transition path is also able to capture.

[146] This is expected given that the model lacks certain mechanisms that we have in the data, for instance firm creation or minimum employment of high-skill workers — firm founders are usually high-skilled people. I also adjust for the fact that in the difference-and-difference setting the increase in local skill supply is around 100% (vs. 150% here, which is in line with the aggregate increase during the period of my sample).

[147] Effectively, I assume both the slow hiring by the follower and the overshooting by the leader dissipate over a period of 3 years.

A.10 Figures and Tables

Figure A1: *Evolution of high-skill concentration and the high to non-high-skill concentration ratio*

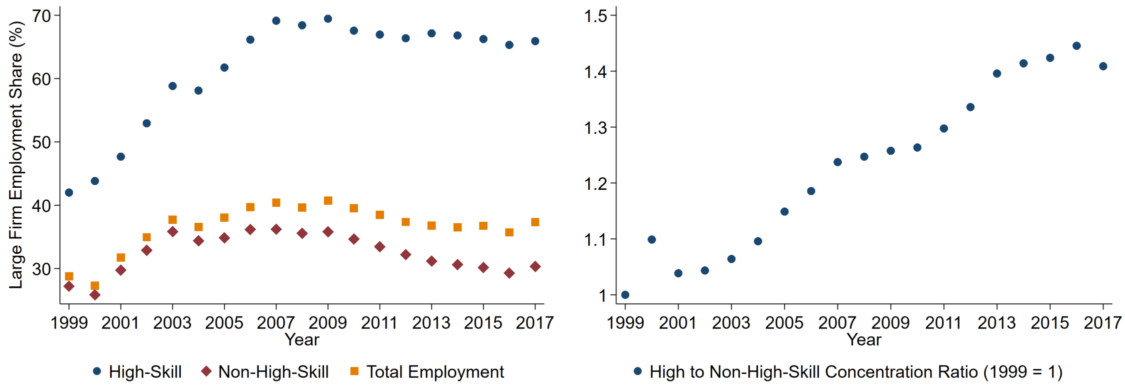

Note: High-skill (non-high-skill) concentration is the median across municipalities of the local share of high-skilled (non-high-skilled) people working at large firms over total local supply. High-skilled workers are those with at least some college education, though they might not have finished their degree.

Figure A2: *Evolution of the HHI-style concentration for different types of labor*

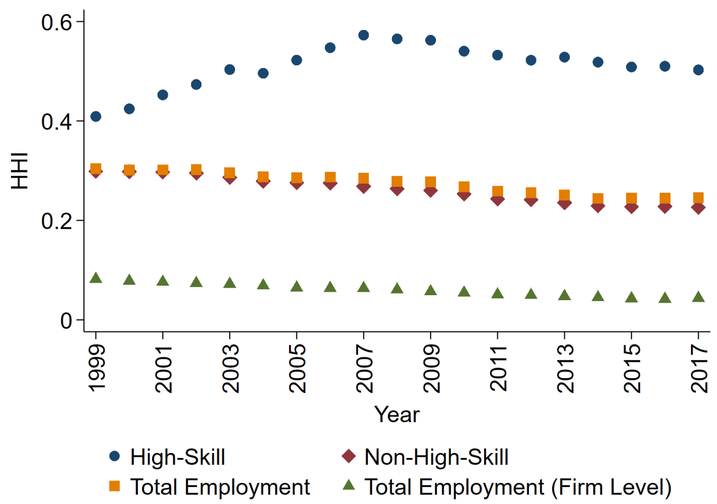

Note: The HHI-based measure of concentration for high-skill, non-high-skill, and total employment is calculated by splitting firms by size bins and calculating the employment HHI between those bins, i.e. by using bin employment shares. To avoid cases where bins of smaller firms have more employees than those of larger firms, I drop localities where that happens. This guarantees that a high value indicates labor concentration at large firms. Total Employment (Firm Level) shows an HHI measure calculated using firm-level employment shares.

Figure A3: *Evolution in the number of colleges and the share of college graduates in the population*

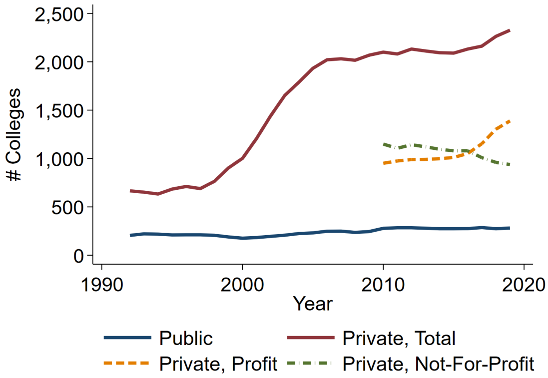

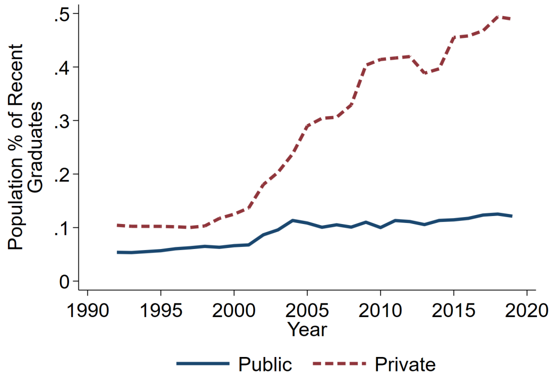

Figure A4: *Effect of college creation on local share of high-skilled people*

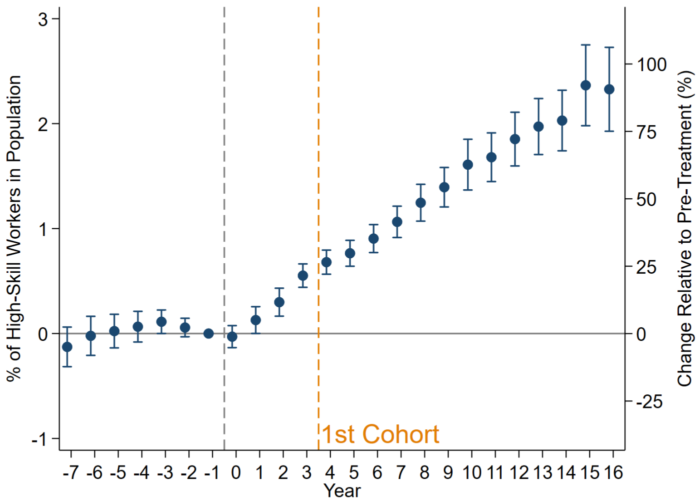

Note: Change relative to pre-treatment period shows the ratio between the coefficient estimates and the average share of high-skilled people across untreated municipalities. Year is relative to the arrival of a new college and dashed orange line represents the period when the first student cohort is expected to graduate. Vertical bars represent the 95% confidence interval.

Figure A5: *Effect of college creation on local share of high-skilled people by treated municipality group*

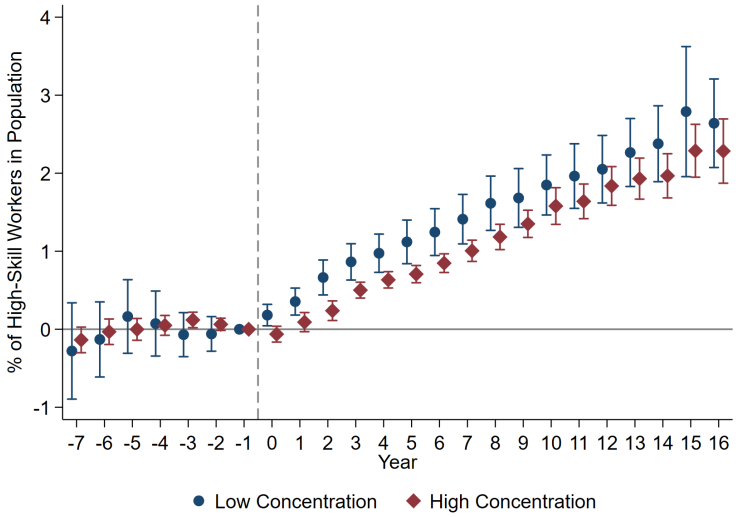

Note: Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A6: *Estimates of the effect of college creation on local growth*

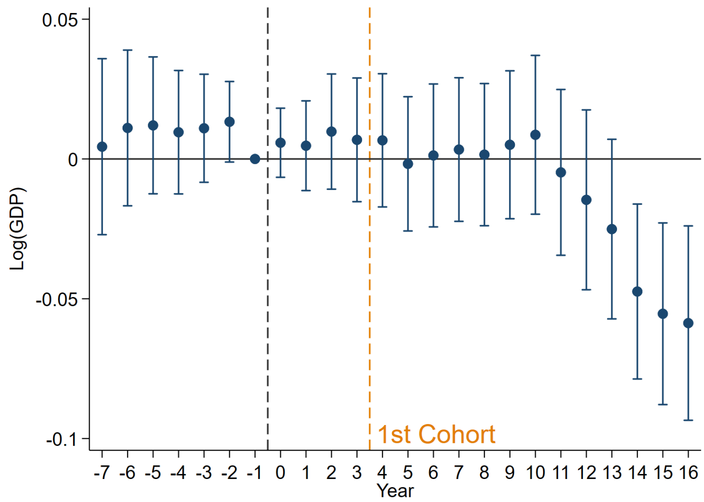

$Log(GDP)$ is the log of local real GDP per capita. Year is relative to the arrival of a new college and dashed orange line represents the period when the first student cohort is expected to graduate. Vertical bars represent the 95% confidence interval.

Figure A7: *Evolution of high-skill share in the population and local skill concentration*

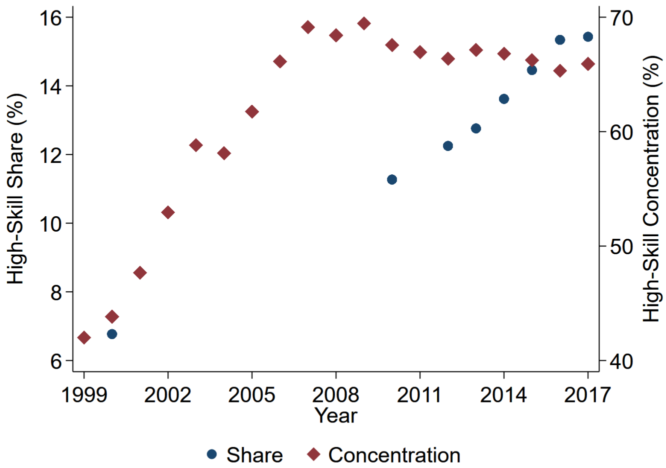

Note: High-skill share corresponds to the ratio between the number of people with a college degree and the total population who is at least 25 years old (UNDP et al., 2024). High-skill concentration is the median across municipalities of the local share of high-skilled people working at large firms over total local supply.

Figure A8: *Binned scatter plot between local GDP per-capita growth and the 1^{st} stage predicted values, both unconditional and conditional on being below or above p*

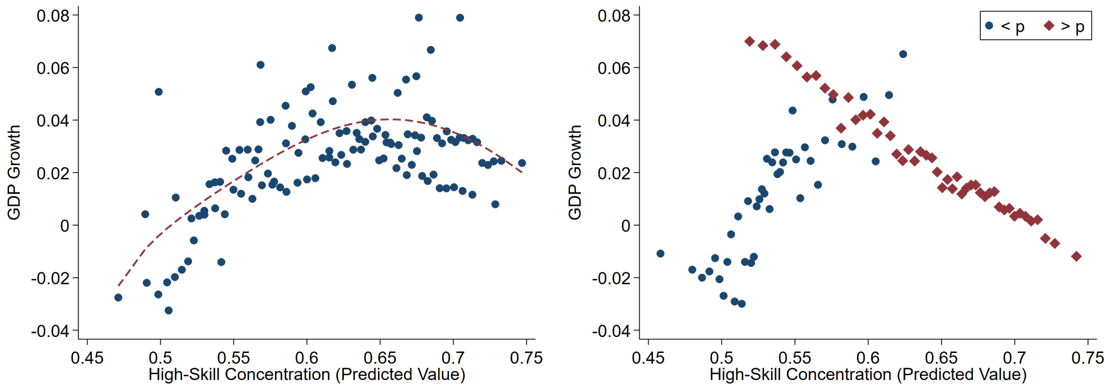

Note: Plots show the predicted value of the 1^{st} stage of the 2SLS ($\widehat{HSConc_{i,t-1}}$) on the x-axis. As p is defined over high-skill concentration ($HSConc_{i,t-1}$), right-hand side figure shows separate binned scatter plots for observations below or above the threshold p. GDP Growth is real per-capita local GDP growth winsorized at the 1^{st} and 99^{th} percentiles. Plots were made using the procedure in Cattaneo et al. (2024), controlling for the local variables used in Column (2) of Table 1.1 along with local and time fixed-effects.

Figure A9: *Left: Value function curves; Right: leader's and laggard's high-skilled labor and investment choices, all as a function of the gap s*

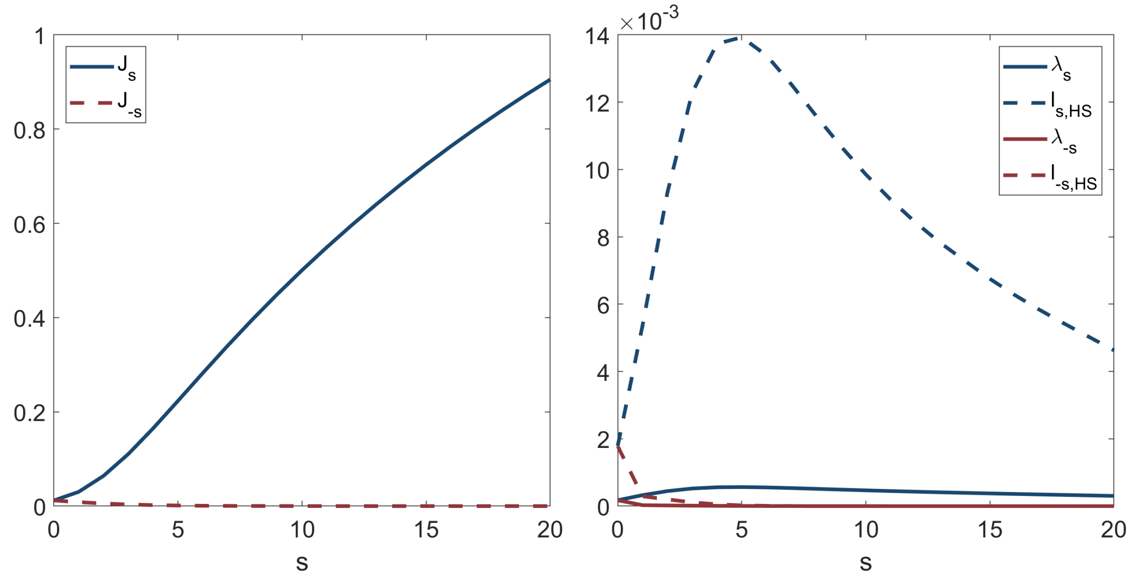

Note: J_s (J_{-s}) refers to the value function of the leader (follower). λ_s (λ_{-s}) refers to R&D investment by the leader (follower). $l_{s,HS}$ ($l_{-s,HS}$) refers to high-skilled labor hired by the leader (follower).

Figure A10: *Growth and the skill premium as a function of high-skilled labor concentration*

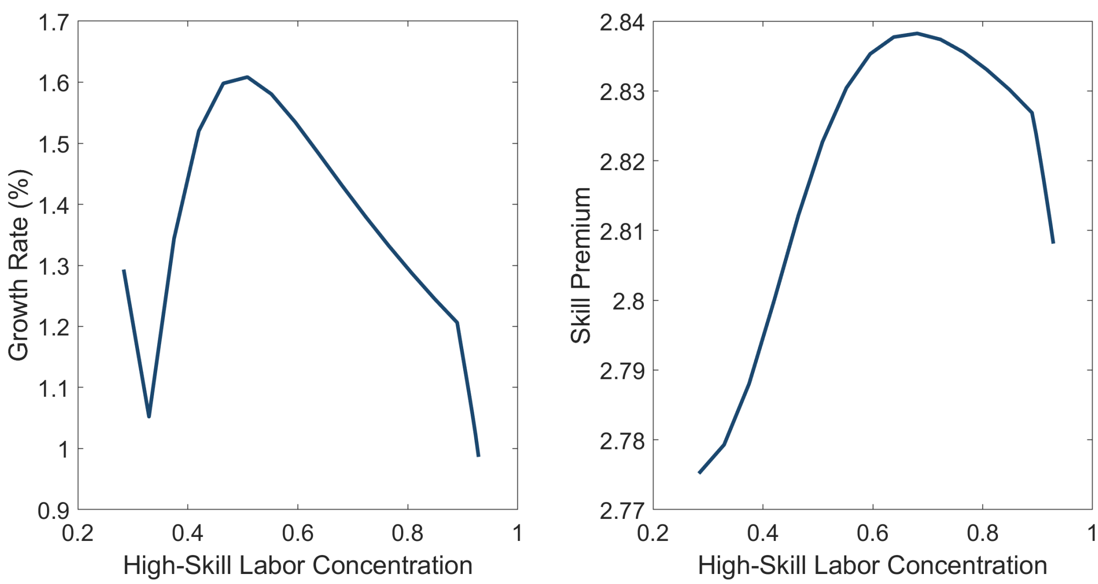

Figure A11: *Share of active R&D catch-up and skill concentration as a function of the gap s*

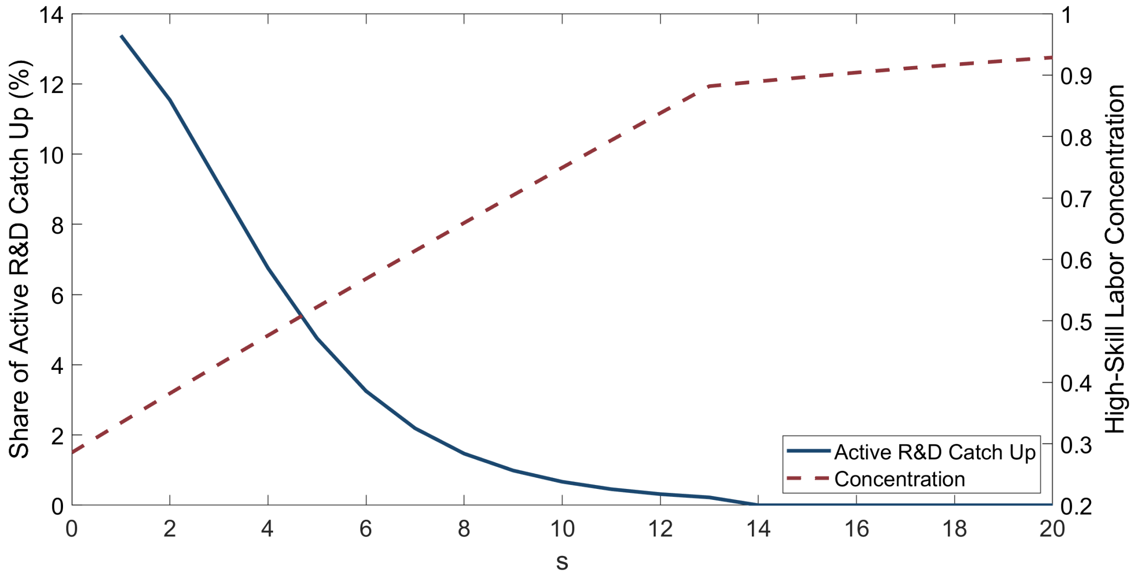

Figure A12: *Growth rate, wage premium, and high-skill concentration as a function of the gap (without the non-innovative, outside firm)*

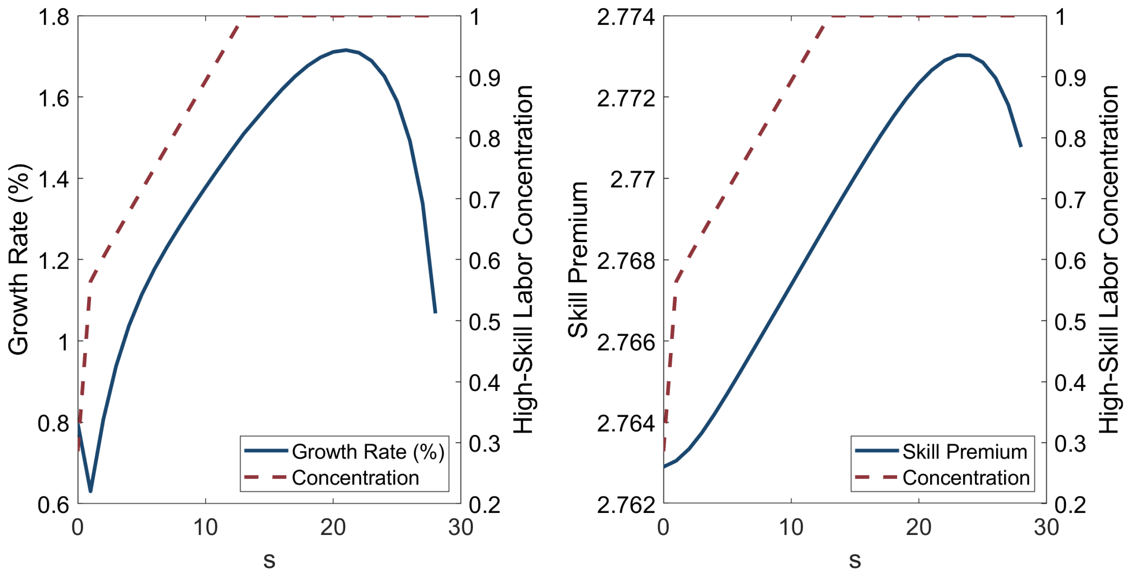

Figure A13: *Growth rate, wage premium, and high-skill concentration as a function of the gap for different values of the R&D labor elasticity*

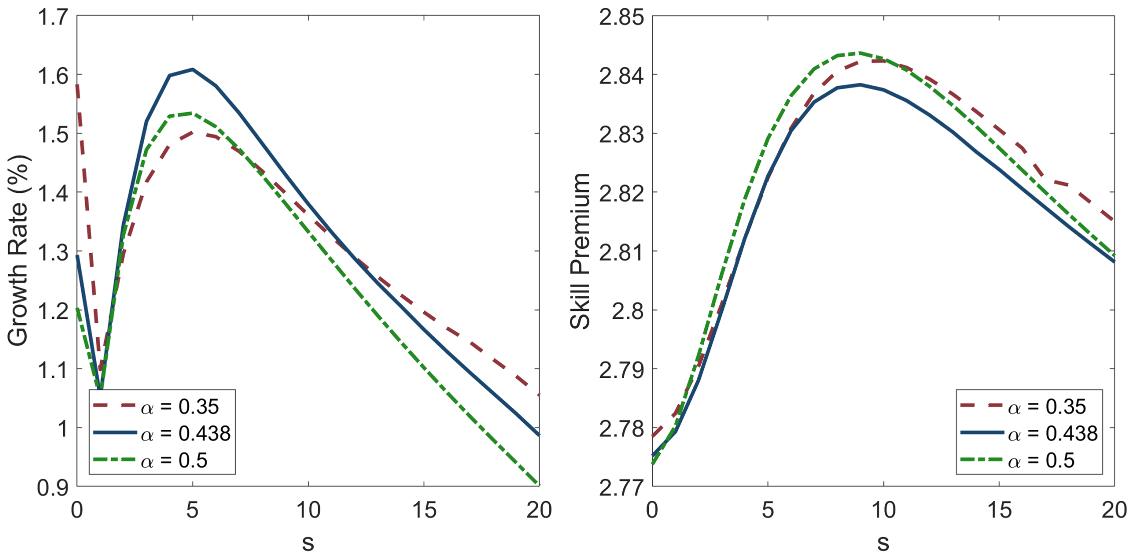

Figure A14: *Growth rate, wage premium, and high-skill concentration as a function of the gap for different convexity values of the R&D cost function*

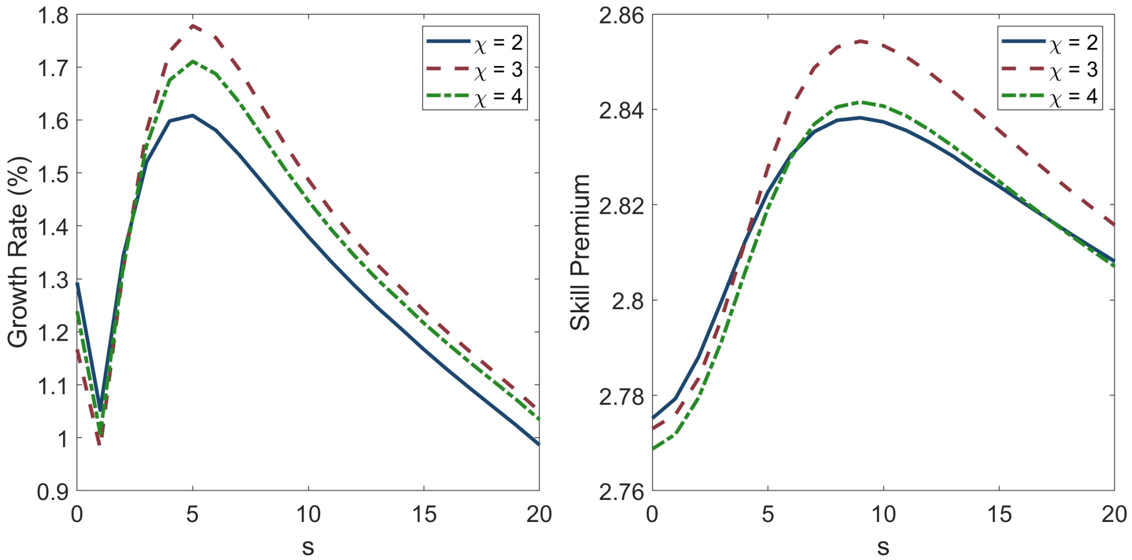

Note: For each curve, we consider the following R&D cost function: $C(\lambda_s) = \rho \frac{\lambda_s^\chi}{\chi}$.

Figure A15: *Growth rate, wage premium, and high-skill concentration as a function of the gap using a smaller numerical adjustment*

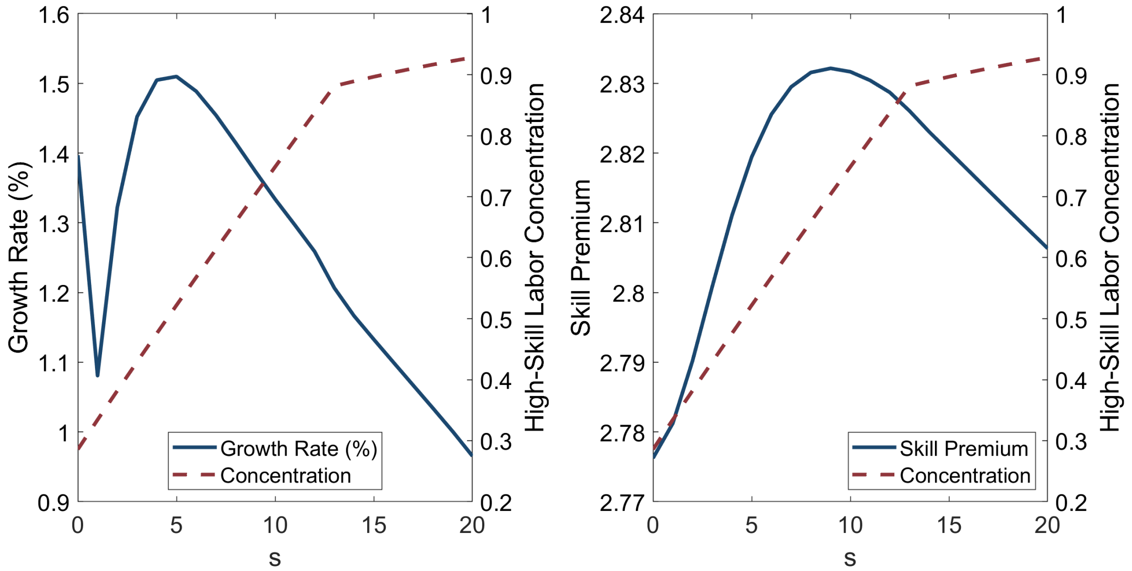

Note: While in the baseline estimation I adjust aggregate demand by adding 0.015 to it, i.e. $D(t) = w_{LS}(t)l_{i,j,L}(t) + w_{o,j}(t)l_{o,j}(t) + \Pi_j(t) + 0.015$, in this plot I use 0.005 instead.

Figure A16: *Growth rate, wage premium, and high-skill concentration calculated without the non-innovative sector, as a function of the gap*

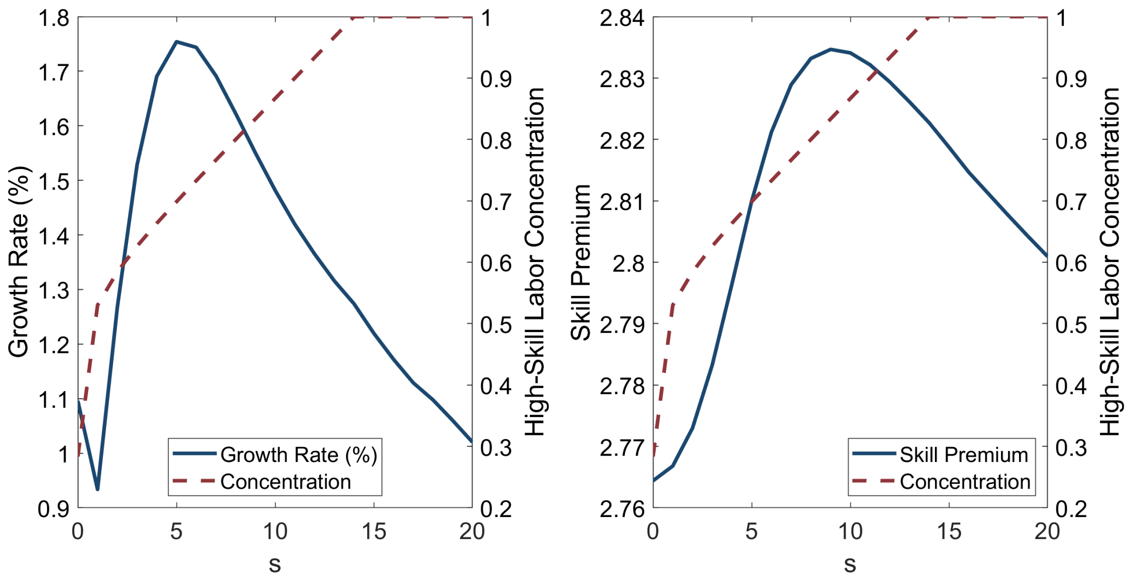

Figure A17: *Growth rate, wage premium, and high-skill concentration, calculated without the LC_2 adjustment, as a function of the gap*

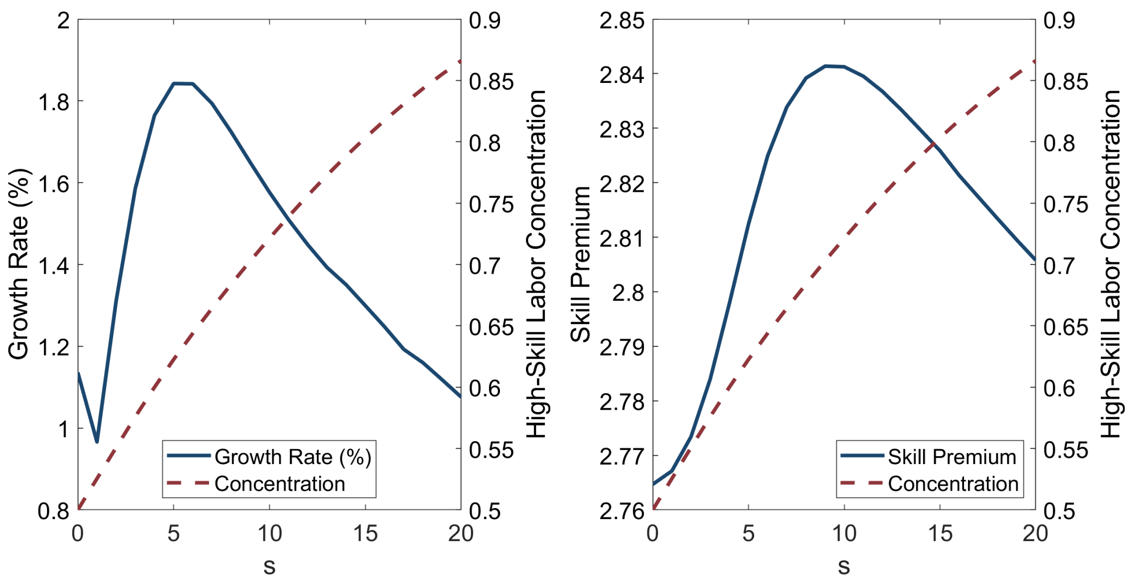

Figure A18: *Evolution of the government-run National Student Performance Exam in Brazil (ENADE) and the Preliminary Course Score (CPC)*

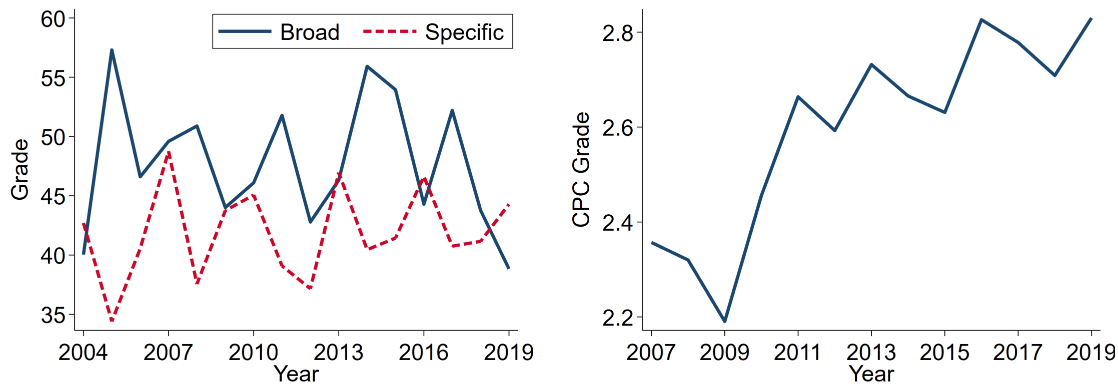

Note: Broad refers to the part of the exam that is common to all degrees. Specific refers to the part of the exam that is specific to a degree. CPC is a composite indicator of quality which takes into account the ENADE grade, teaching staff quality, student feedback, and an indicator of learning value added.

Figure A19: *Evolution of college graduates composition between areas of study*

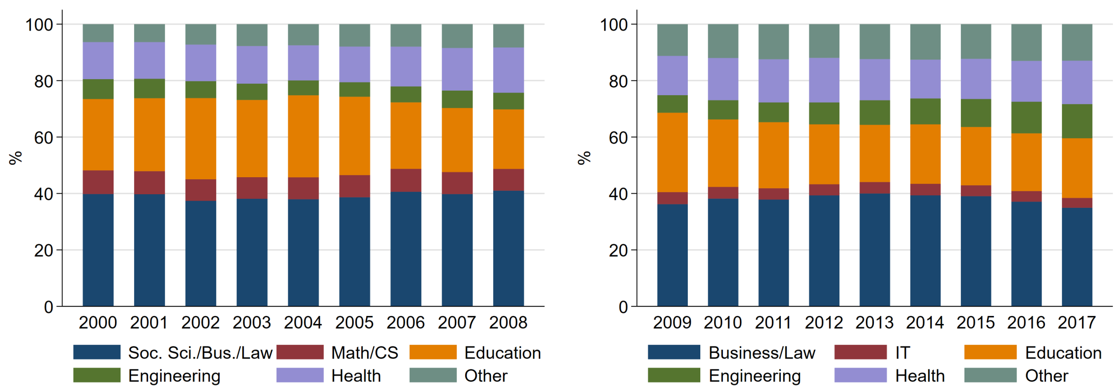

Note: A new area-of-study classification from 2009 onward leads to a breakdown in the series. Soc. Sci. refers to Social Sciences. Bus. refers to Business. CS refers to Computer Science.

Figure A20: *high-skilled population share, GDP per-capita growth trend, and Romer-model-based expected growth in Brazil between 1991 and 2019*

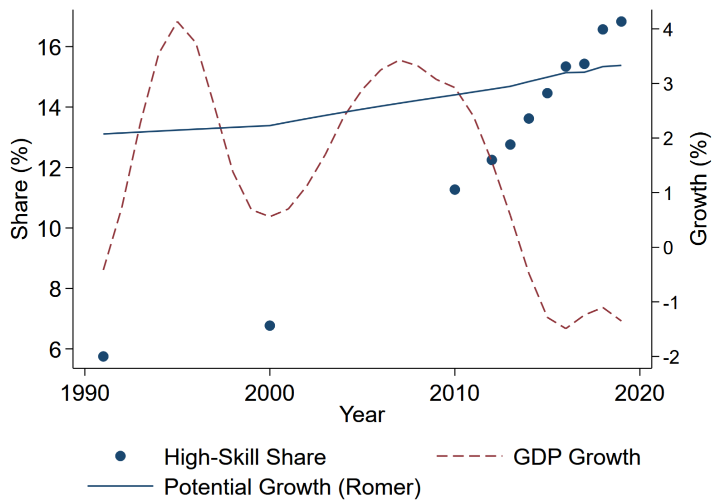

Note: High-skill share data is from the Atlas of Human Development in Brazil (UNDP et al., 2024). High-skill share corresponds to the ratio between the number of people with a college degree and the total population who is at least 25 years old. GDP growth is the real GDP per-capita growth trend after filtering data since 1970 using a Hodrick-Prescott filter. Expected Growth is the expected growth from the increase in high-skill labor share in a Romer-based model (c.f. Section A.1).

Figure A21: *Growth rate as a function of the gap s and the distribution of gaps in the economy at different L_{HS}*

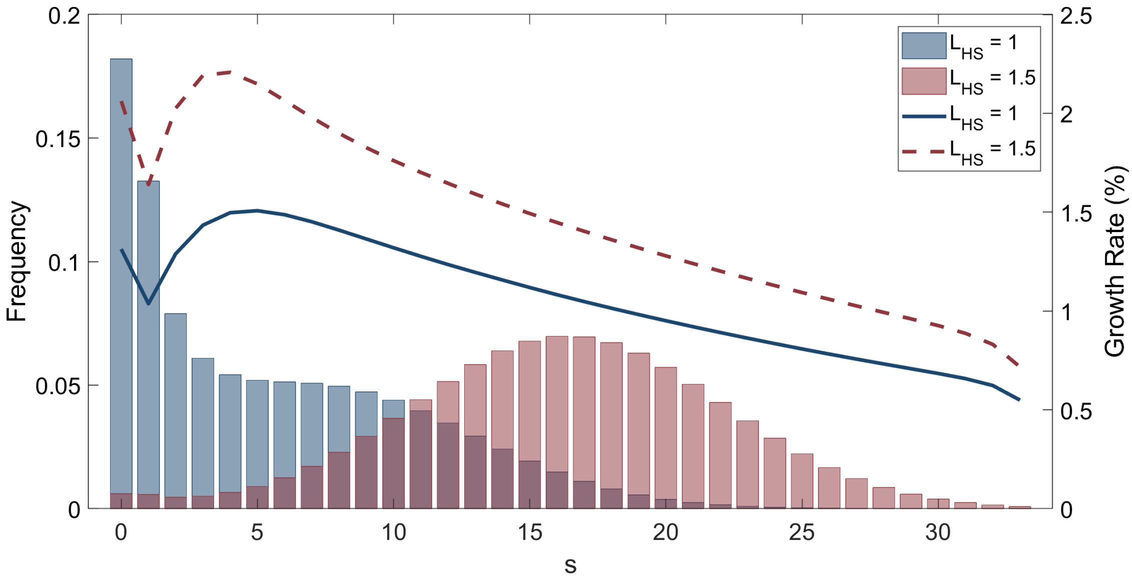

Figure A22: *R&D effort growth and breakdown of the change in the value function as high-skill supply increases from 1 to 1.5*

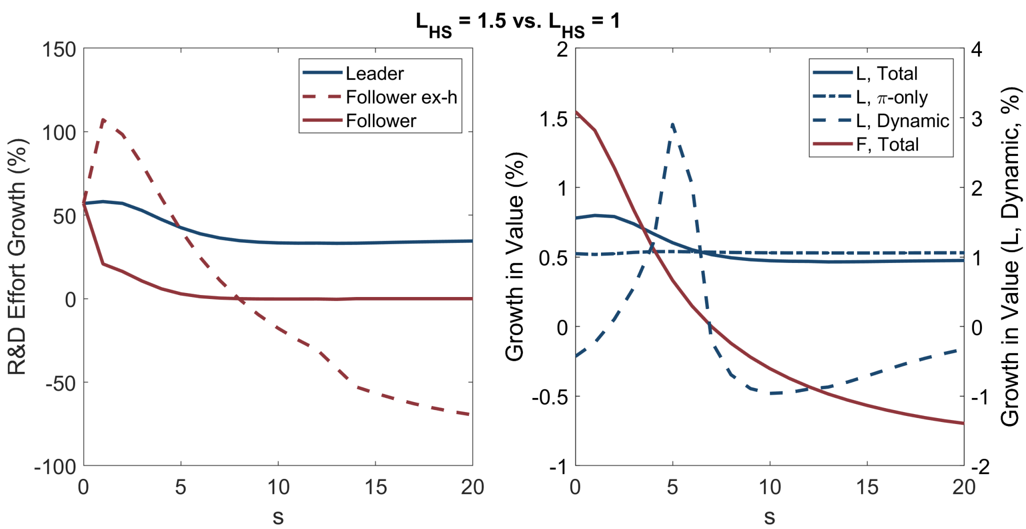

Note: Ex-h refers to the follower's R&D effort without the catch-up term. L, Total, π-only, and Dynamic refer to the change in the leader's total, profit-only, and dynamic parts of its value function, respectively. F, Total refers to the change in the follower's total value function.

Figure A23: *Breakdown of the change in the leader's value function as high-skill supply increases from 1 to 1.5*

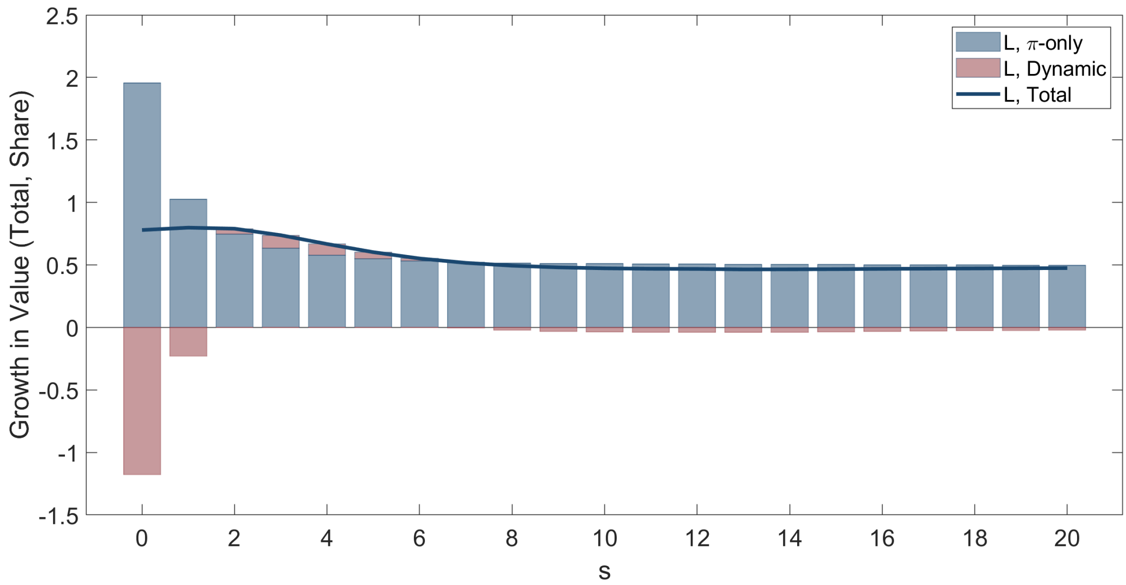

Note: L, Total, π-only, and Dynamic refer to the change in the leader's total, profit-only, and dynamic parts of its value function, respectively.

Figure A24: *Growth and skill concentration as a function of human capital supply for different levels of initial skill concentration*

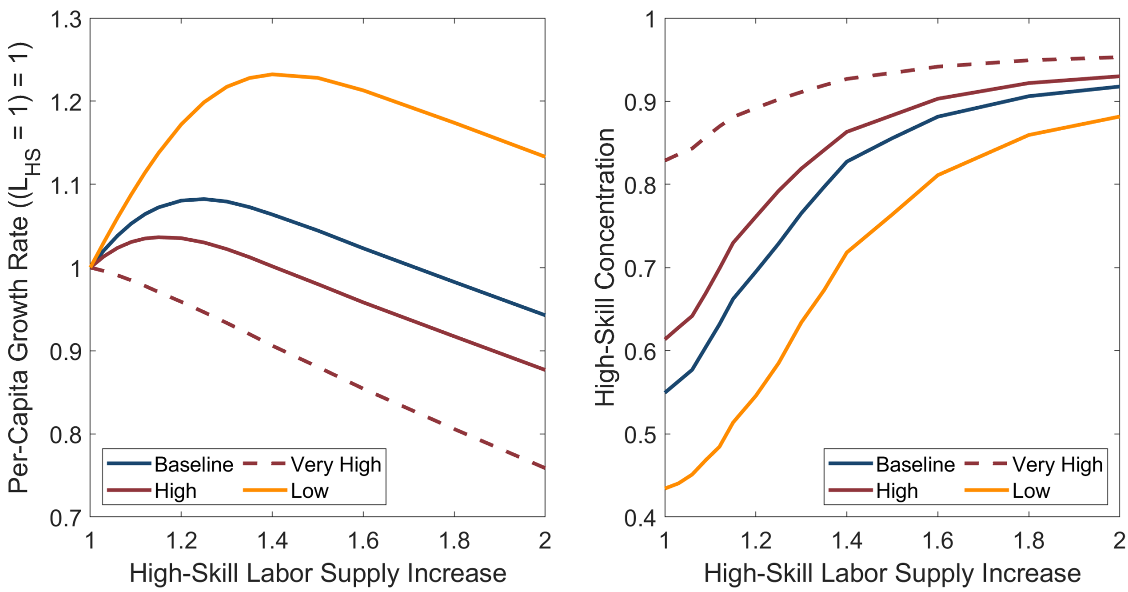

Note: Baseline refers to values using parameter estimates from Table A.7. High, Very High, and Low (Skill Concentration) use the same set of parameters as the baseline scenario except for h_c whose value is $h_{c,baseline}/1.05$, $h_{c,baseline}/1.25$, and $1.1h_{c,baseline}$, respectively.

Figure A25: *Effect of increasing high-skilled labor supply on the skill premium and high-skill unemployment for different increases in aggregate labor supply*

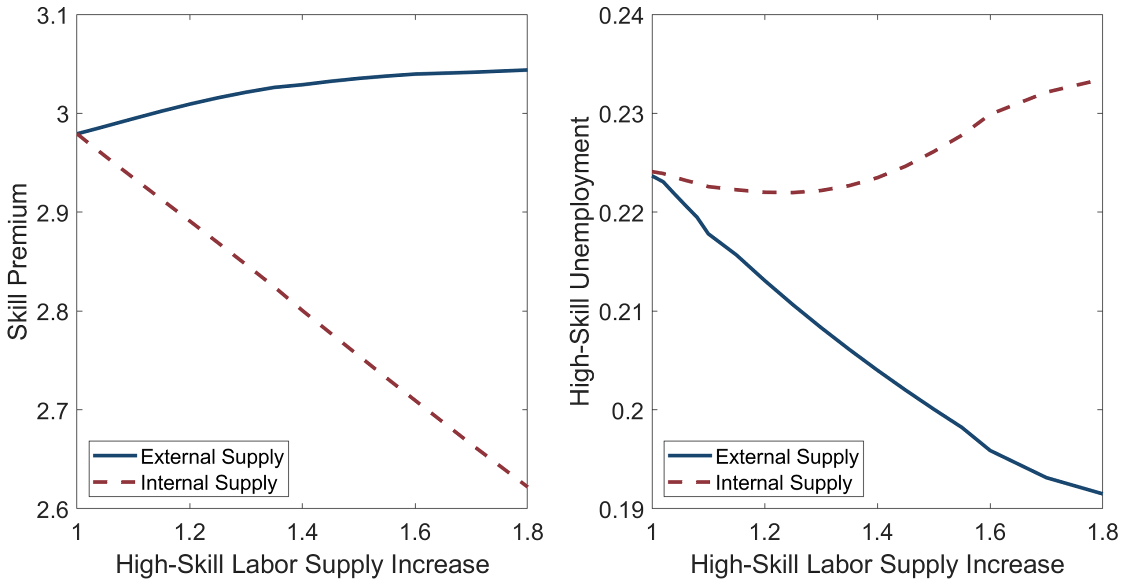

Figure A26: *Breakdown of the effect of increasing high-skilled labor supply on the skill premium by channel*

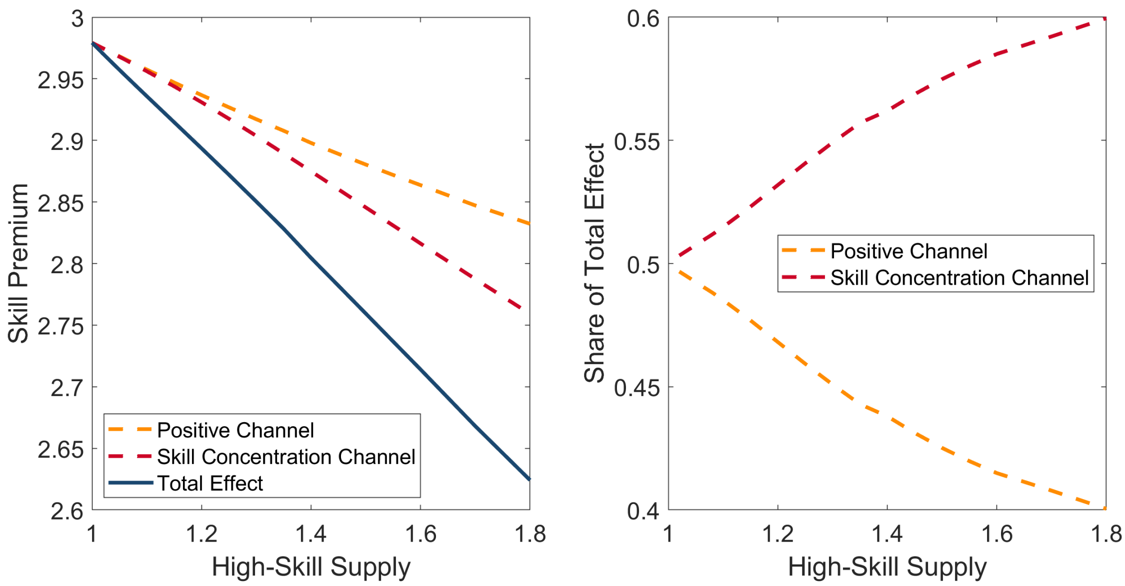

Note: Positive Channel refers to the partial effect of higher R&D effort from an increase in the supply of high-skilled labor. Skill Concentration Channel refers to the partial effect of the shift in firm gap distribution. Total Effect is the total effect of the increase in skill supply on the skill premium.

Figure A27: *Evolution of the skill premium in Brazil*

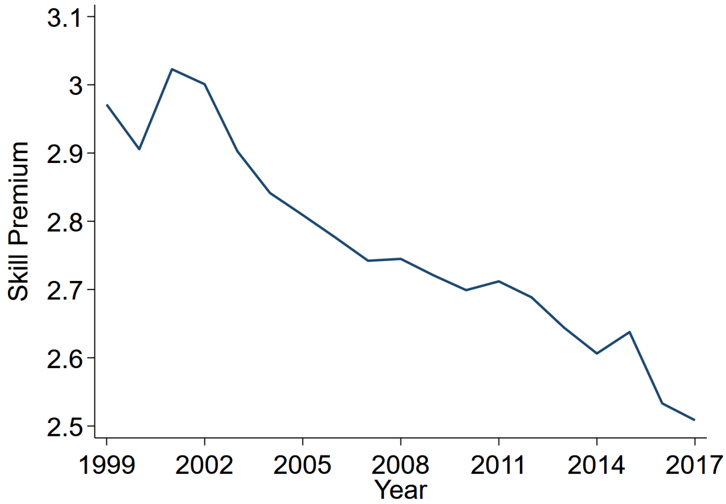

Note: Skill premium consists of the weighted average of the municipality-level ratio between high and low-skill wages, weighted by the total number of local workers.

Figure A28: *Evolution of high-skill underemployment and unemployment-to-population shares ratio in Brazil*

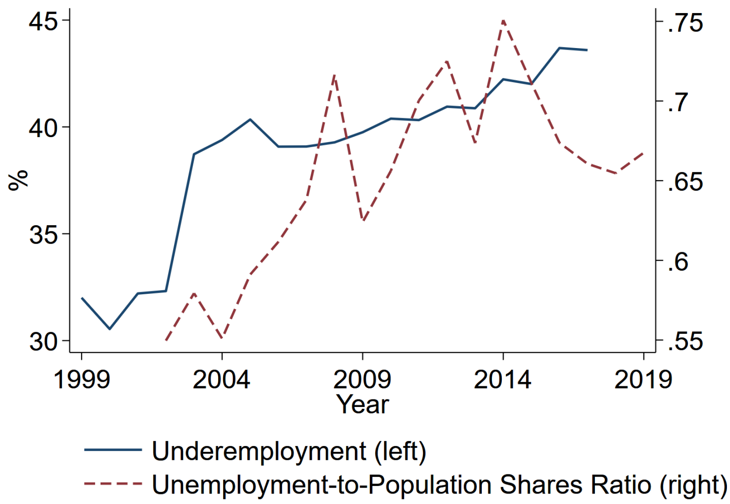

Note: High-skill denotes those with 15+ years of study which corresponds to at least a college degree. Underemployment data comes from RAIS and is defined as an employee with a college degree working in Groups 3-9 in Brazil's occupational classification system (CBO). Unemployment and population shares come from the National Household Sample Survey (PNAD). Values for 2010 are interpolated as the household survey is not run during Census years.

Figure A29: *Active R&D catch-up as high-skilled labor supply increases*

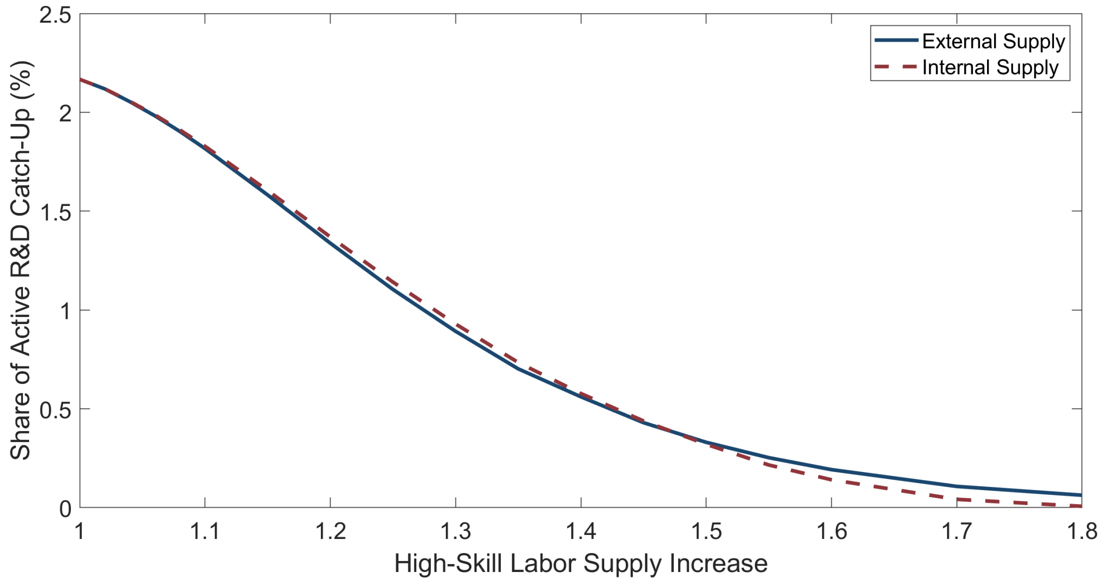

Figure A30: *Firms' R&D effort at* $L_{HS} = 1.5$ *for the baseline and subsidy cases*

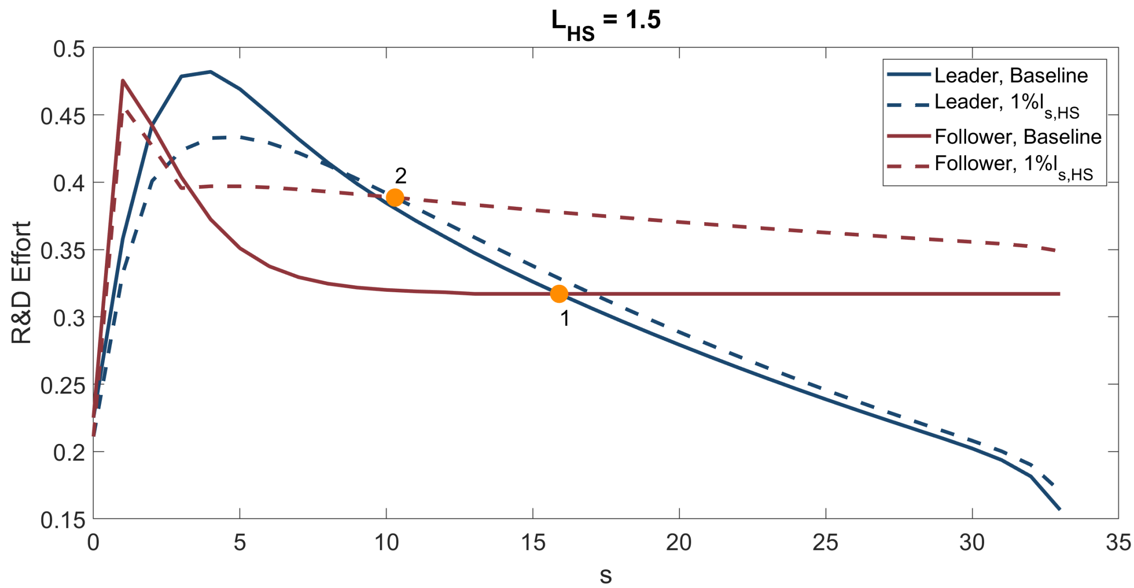

Figure A31: *Effect of increasing high-skilled labor supply for different levels of h_c*

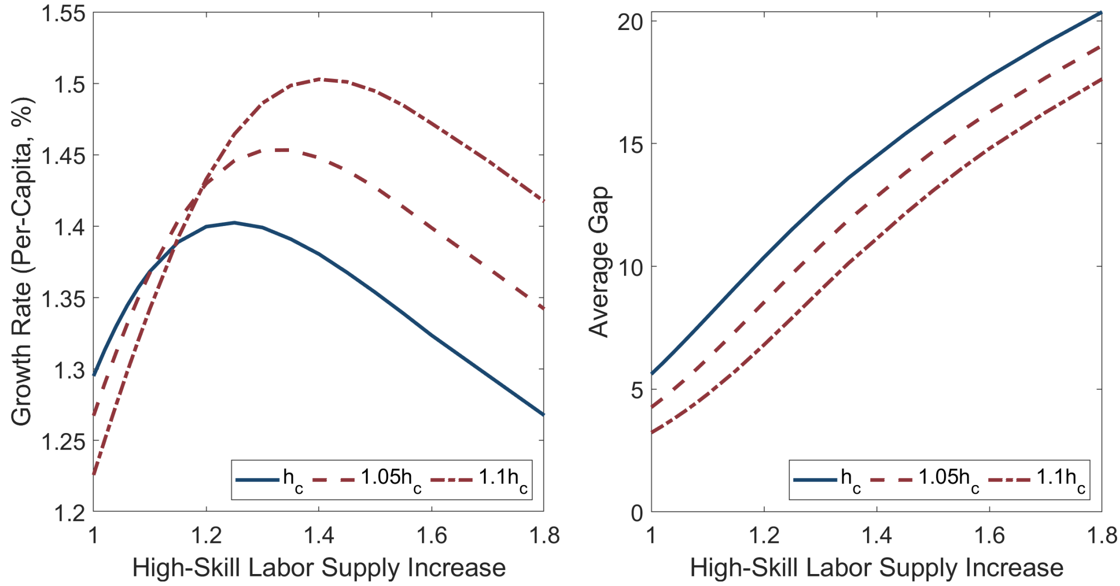

Figure A32: *Estimated $\hat{\xi}_2$ for different "bad control" conditions*

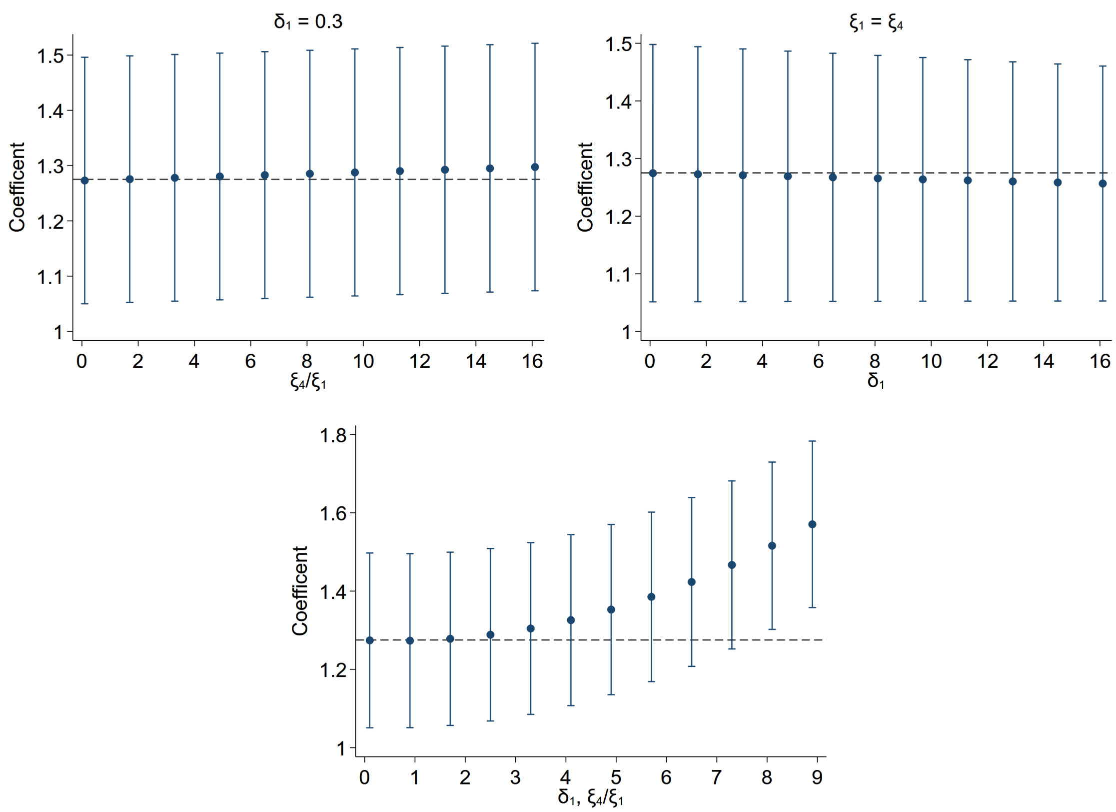

Note: Horizontal line corresponds to the true value of ξ_2. Vertical bars show the 95% confidence interval.

Figure A33: *Estimates of the effect of college creation on local formal employment*

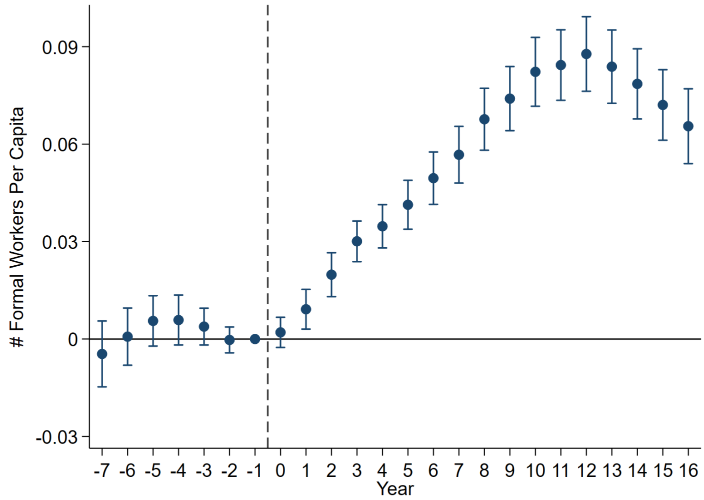

Employment data uses employer-employee links to calculate the local stock of formal workers. Year is relative to the arrival of a new college. Vertical bars represent the 95% confidence interval.

Figure A34: *Trends on college supply competition and college demand relative to the arrival of a new college*

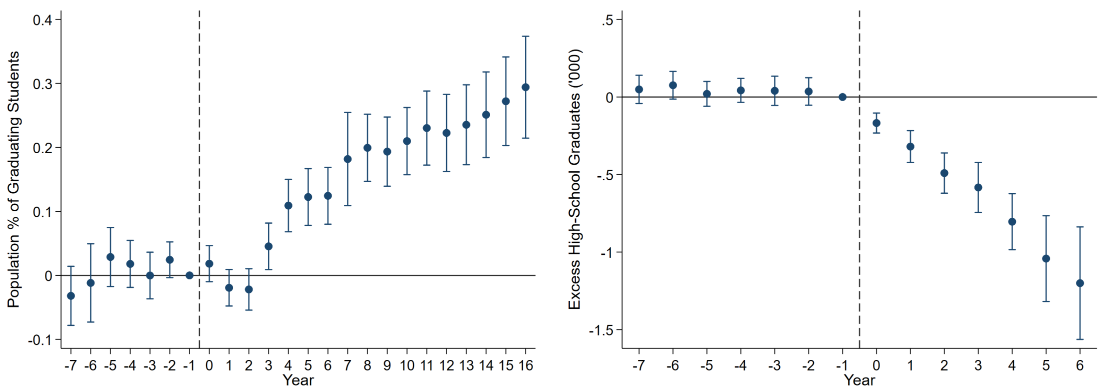

Population % of graduating students refers to the population share of college students who graduated in each year. Excess high-school graduates refers to the difference between the number of new high-school graduates in each year and the size of the incoming first-year college cohort, in thousands. High-school data comes from INEP and is restricted to the 1999-2006 period. Year is relative to the arrival of a new college. Vertical bars represent the 95% confidence interval.

Figure A35: *Estimates of the effect of college creation on local growth at municipalities with high and low high-skill concentration using the placebo-to-treated matched sample*

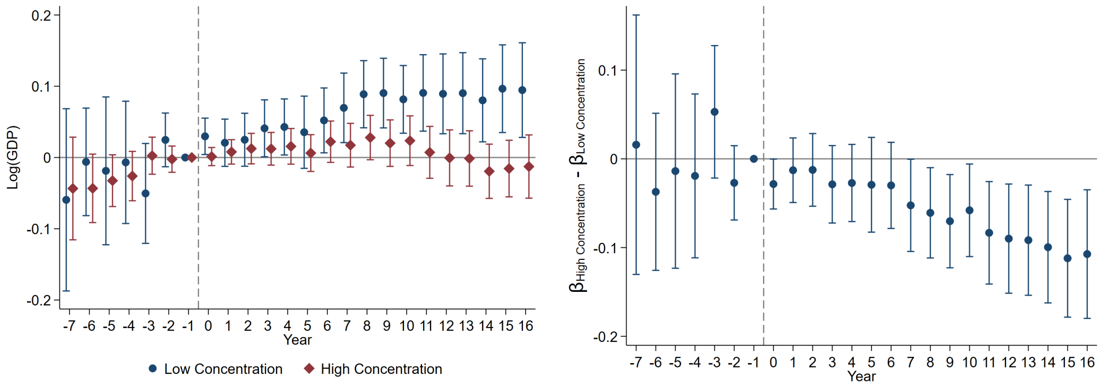

Note: $Log(GDP)$ is the log of local real GDP per capita. High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Sample match is on population level, share earning minimum wage or less, share who only completed the 5^{th} grade, unemployment rate, and illiteracy rate, all in 2000. Treated observations are matched to those in control using the coarsened matching method in Iacus et al. (2012). Low (high) concentration municipalities are defined as those below (above) the 14^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A36: *Estimates of the effect of college creation on local growth at municipalities with high and low high-skill concentration using the last-treated control group*

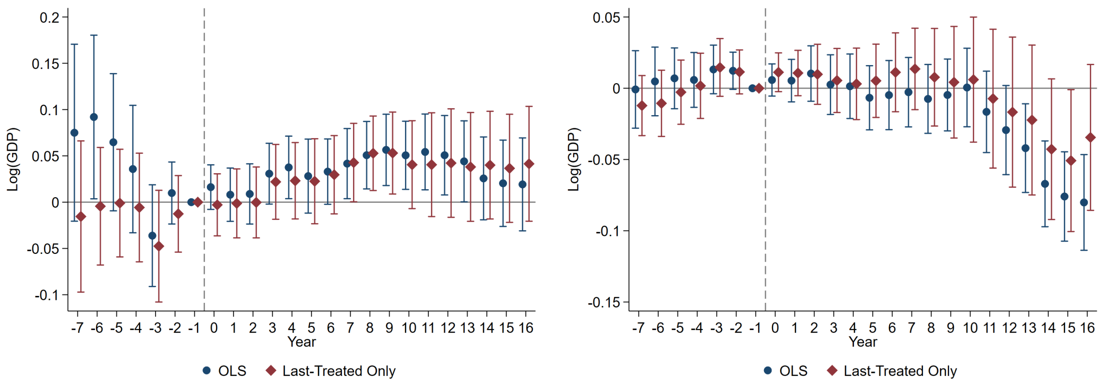

Note: $Log(GDP)$ is the log of local real GDP per capita. High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 14^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. Estimates restrict the control group to last-treated units and use the estimator proposed in Sun & Abraham (2021). Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A37: *Estimates of the difference in the effect of college creation on local growth between municipalities with high and low high-skill concentration using the last-treated control group*

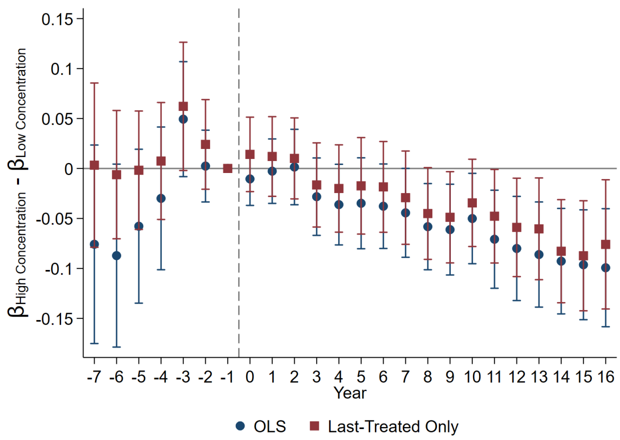

Note: High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 14^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. Estimates restrict the control group to last-treated units and use the estimator proposed in Sun & Abraham (2021). Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A38: *Estimates of the effect of college creation on local growth at municipalities with high and low high-skill concentration using the not-yet-treated control group*

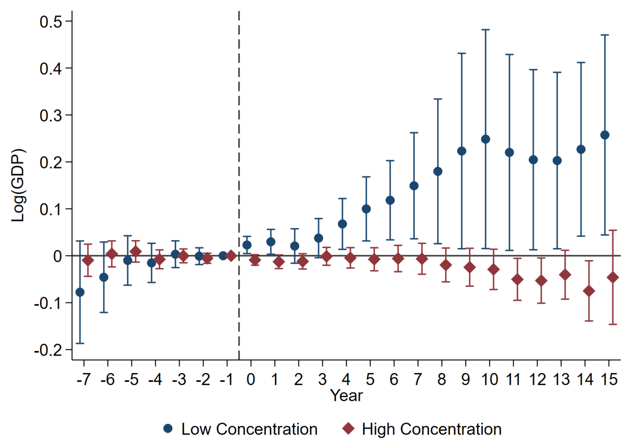

Note: $Log(GDP)$ is the log of local real GDP per capita. High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 20^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. Year refers to the time relative to when the first cohort is expected to graduate. Estimates restrict the control group to not-yet-treated units and use the estimator proposed in de Chaisemartin & D'Haultfœuille (2024). Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A39: *Estimates of the effect of college creation on local growth at municipalities with high and low high-skill concentration for the no-college sample*

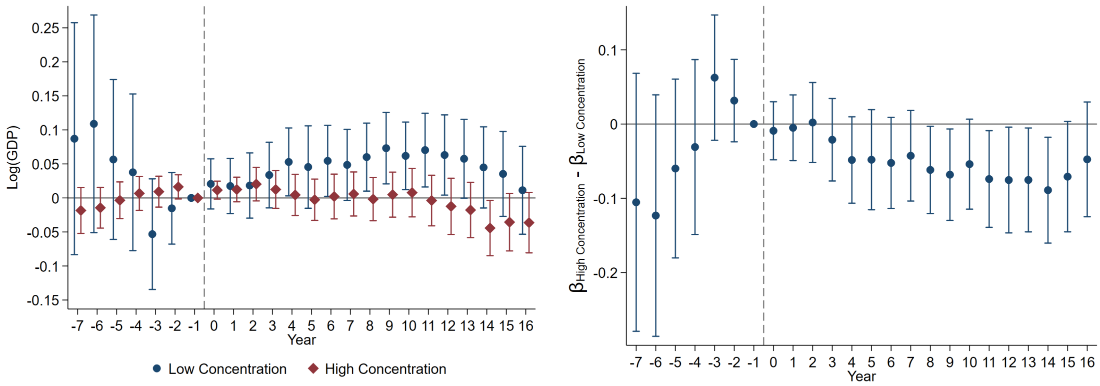

Note: $Log(GDP)$ is the log of local real GDP per capita. No-college sample only includes observations that do not have a college in all periods (control) or in the pre-treatment period (treated). High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 14^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases and between the no-college sample and the baseline one. Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A40: *Estimates of the effect of college creation on local growth at municipalities with high and low high-skill concentration for* $p = 12\%$

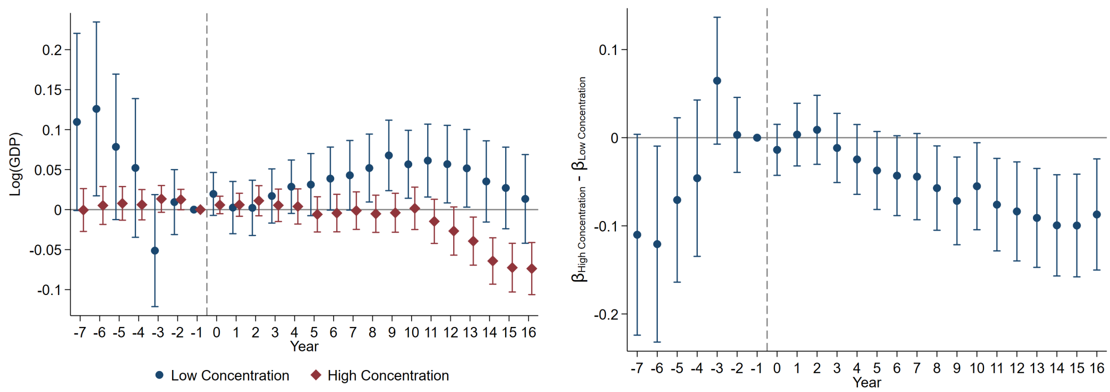

Note: $Log(GDP)$ is the log of local real GDP per capita. High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 12^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A41: *Estimates of the effect of college creation on local growth at municipalities with high and low high-skill concentration for* $p = 17\%$

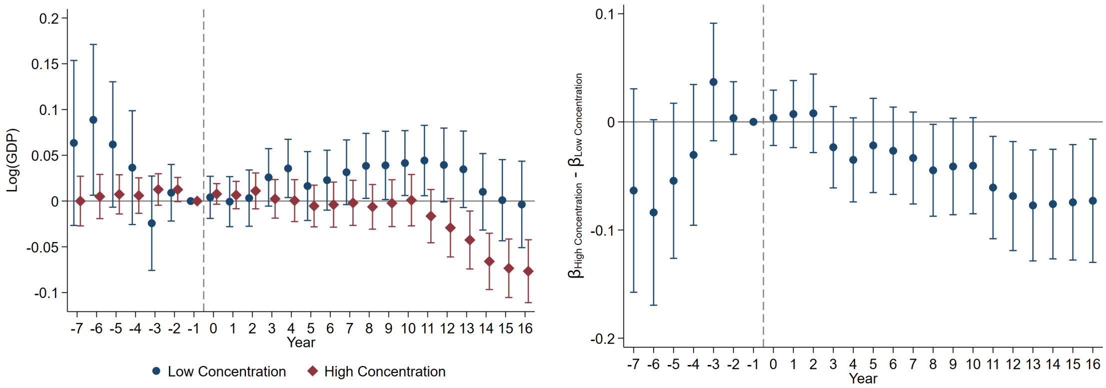

Note: $Log(GDP)$ is the log of local real GDP per capita. High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 17^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A42: *Estimates of the effect of college creation on local growth at municipalities with high and low high-skill concentration controlling for multiple treatments*

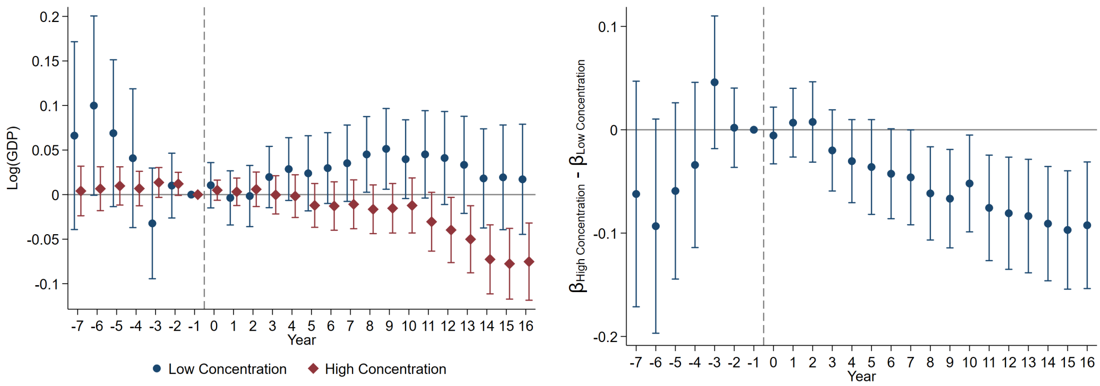

Note: $Log(GDP)$ is the log of local real GDP per capita. Controls include the leads and lags of places treated twice and/or three times. High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 14^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A43: *Estimates of the effect of college creation on local growth at municipalities with high and low high-skill concentration (weighted by log(population))*

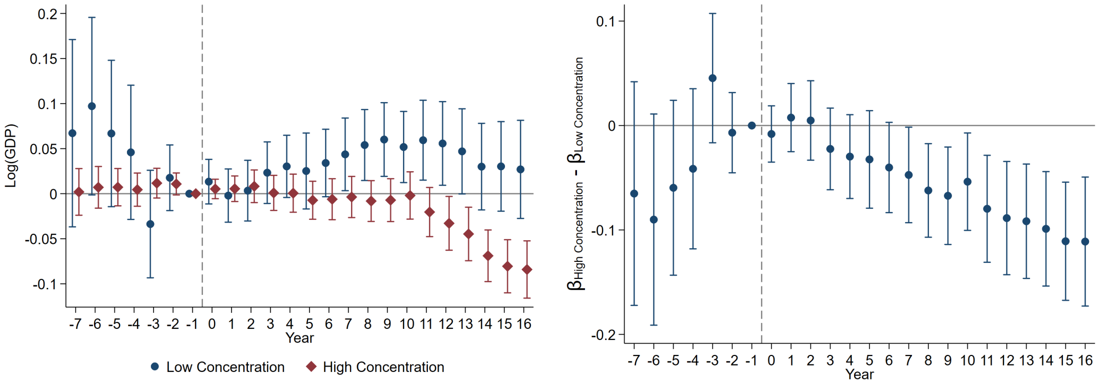

Note: $Log(GDP)$ is the log of local real GDP per capita. High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 14^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. Specification runs a weighted regression using the logarithm of local population as weights. Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A44: *Estimates of the effect of college creation on local growth (removing exports) at municipalities with high and low high-skill concentration*

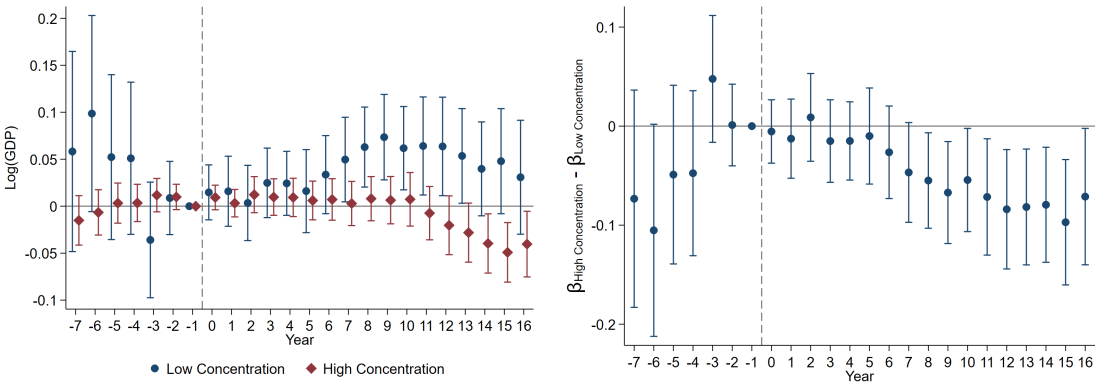

Note: $Log(GDP)$ is the log of local real GDP per capita. High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 14^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. Local GDP excludes exports. Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A45: *Estimates of the effect of college creation on local growth using robust estimators*

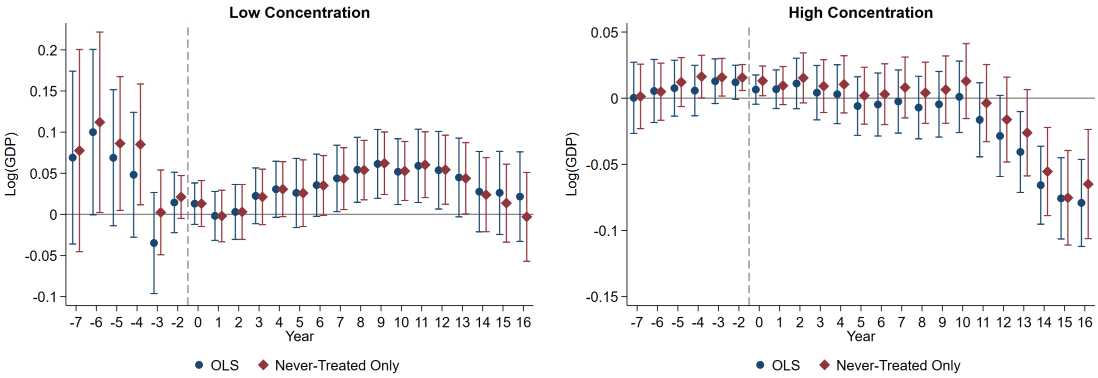

$Log(GDP)$ is the log of local real GDP per capita. High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 14^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. Never-Treated Only estimates restrict the control group to never-treated units and use the estimator proposed in Sun & Abraham (2021). Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A46: *Estimates of the effect of college creation on local growth (baseline vs. non-binary treatment estimator)*

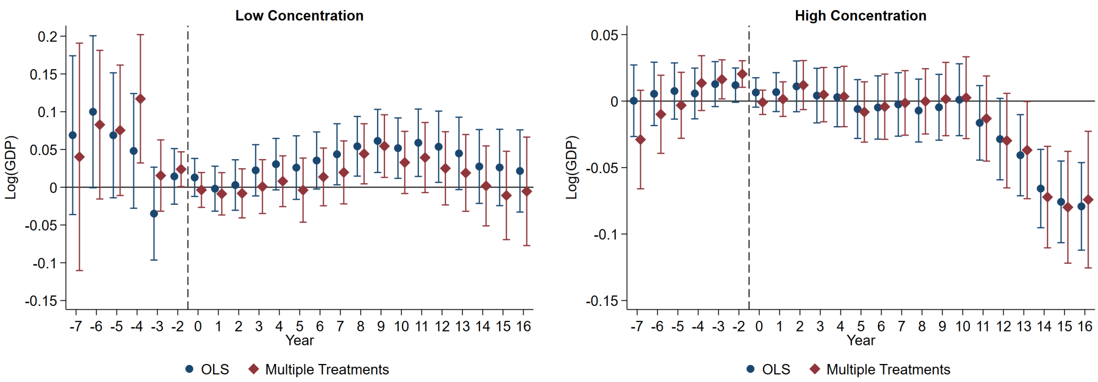

$Log(GDP)$ is the log of local real GDP per capita. High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 14^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. Multiple Treatments estimates (de Chaisemartin & D'Haultfœuille, 2024) restrict the control group to never-treated units and the treated group to places that saw an increase in the number of colleges. Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A47: *Estimates of the effect of college creation on local growth using robust estimators and controls*

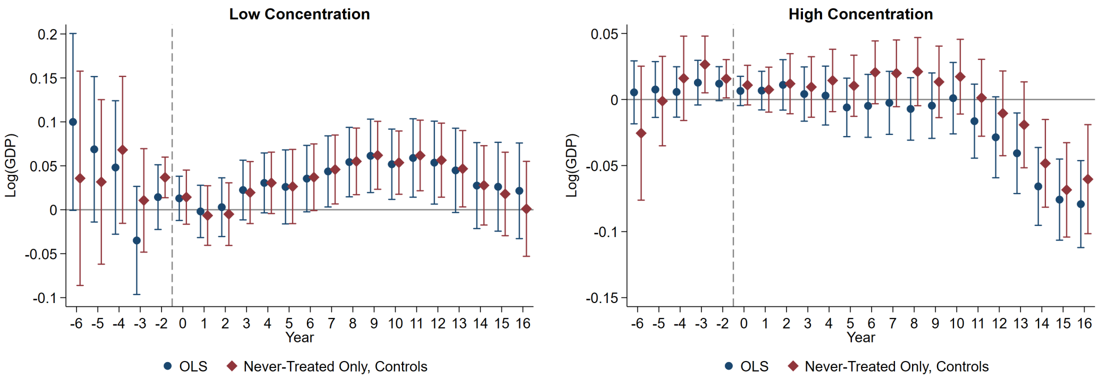

Log(*GDP*) is the log of local real GDP per capita. High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 14^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. Never-Treated Only, Controls estimates (Sun & Abraham, 2021) restrict the control group to never-treated units and control for the log of population, average real wage, the population share of workers receiving minimum wage or less, employment shares by 1-digit sectors, and the share of workers with different levels of education. Year is relative to the arrival of a new college. Vertical bars represent the 90% confidence interval.

Figure A48: *Estimates of the effect of college creation on local high-skill concentration using the placebo group as control*

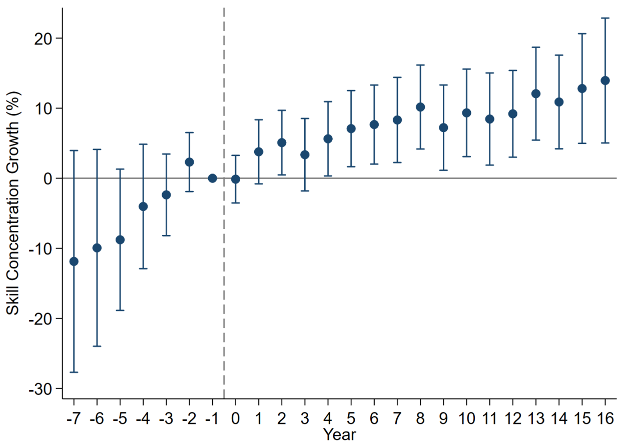

Note: High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Sample excludes observations with no workers at large firms. Sample match is on population level, share earning minimum wage or less, share who only completed the 5^{th} grade, unemployment rate, and illiteracy rate, all in 2000. Treated observations are matched to those in control using the coarsened matching method in Iacus et al. (2012). Year is relative to the arrival of a new college. Vertical bars represent the 95% confidence interval.

Figure A49: *Estimates of the effect of college creation on local high-skill concentration for the no-college sample*

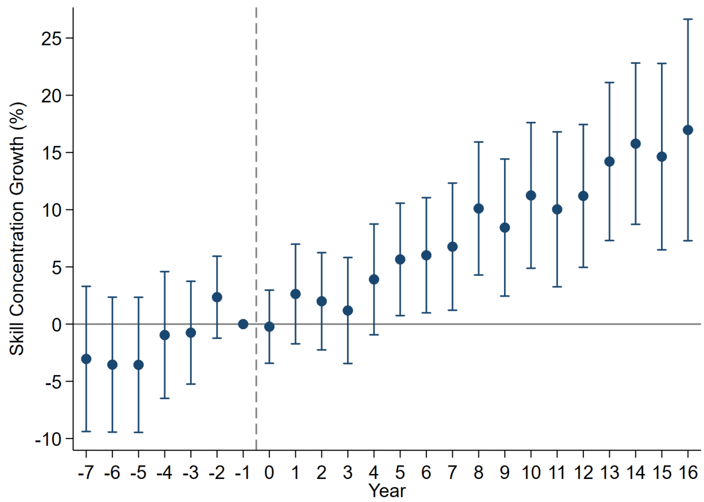

Note: High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Sample excludes observations with no workers at large firms. No-college sample only includes observations that do not have a college in all periods (control) or in the pre-treatment period (treated). Year is relative to the arrival of a new college. Vertical bars represent the 95% confidence interval.

Figure A50: *Estimates of the effect of college creation on local high-skill concentration controlling for multiple treatments*

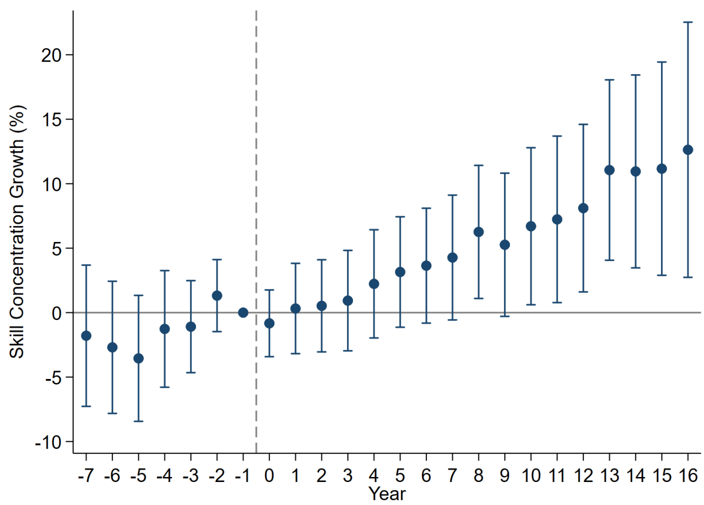

Note: High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Sample excludes observations with no workers at large firms. Controls include the leads and lags of places treated twice and/or three times. Year is relative to the arrival of a new college. Vertical bars represent the 95% confidence interval.

Figure A51: *Estimates of the effect of college creation on local high-skill concentration (HHI-based)*

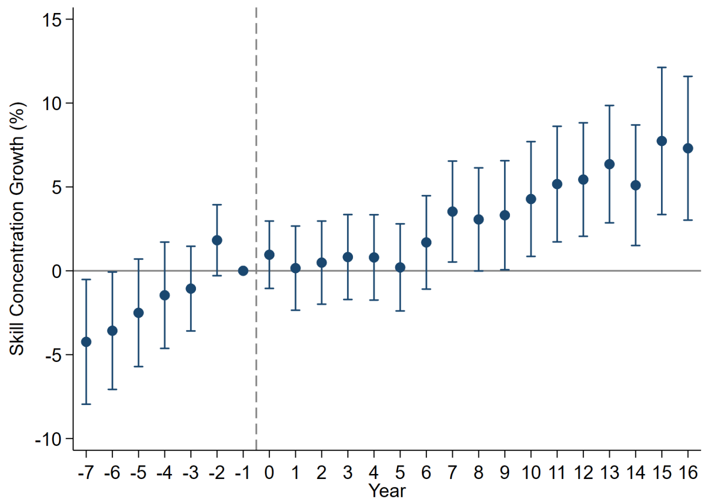

Note: High-skill concentration is the HHI-based measure calculated using firm size bins, i.e. for each firm-size range in the RAIS dataset I calculate the sum of the square of the corresponding local employment share. Sample excludes observations with no workers at large firms. Year is relative to the arrival of a new college. Vertical bars represent the 95% confidence interval.

Figure A52: *Comparison between baseline and robust difference-in-difference estimates*

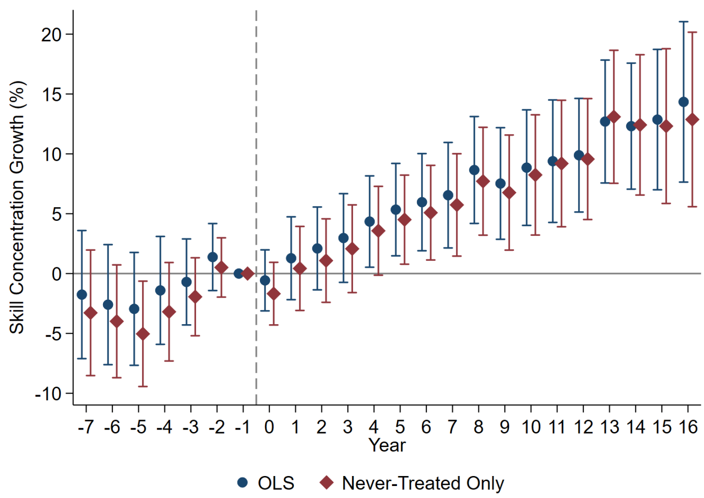

Note: High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Sample excludes observations with no workers at large firms. Never-Treated Only estimates restrict the control group to never-treated units and use the estimator proposed in Sun & Abraham (2021). Year is relative to the arrival of a new college. Vertical bars represent the 95% confidence interval.

Figure A53: *Comparison between baseline estimates and robust difference-in-difference estimates controlling for local-level observables*

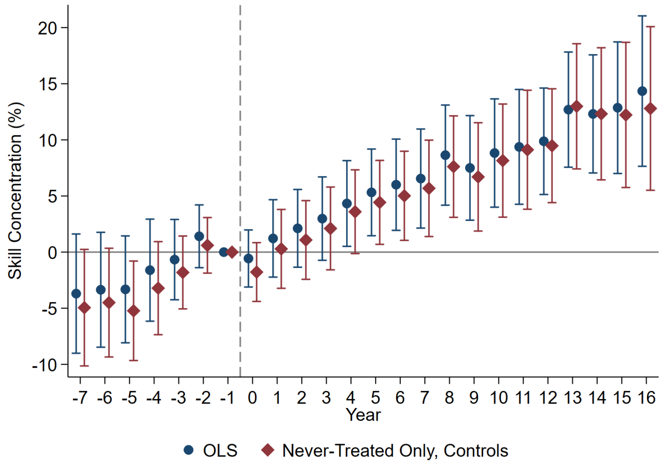

Note: High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Sample excludes observations with no workers at large firms. Never-Treated Only, Controls estimates (Sun & Abraham, 2021) restrict the control group to never-treated units and control for the log of population, average real wage, the population share of workers receiving minimum wage or less, employment shares by 1-digit sectors, and the share of workers with different levels of education. Year is relative to the arrival of a new college. Vertical bars represent the 95% confidence interval.

Figure A54: *Comparison between baseline results and estimates using non-binary treatment*

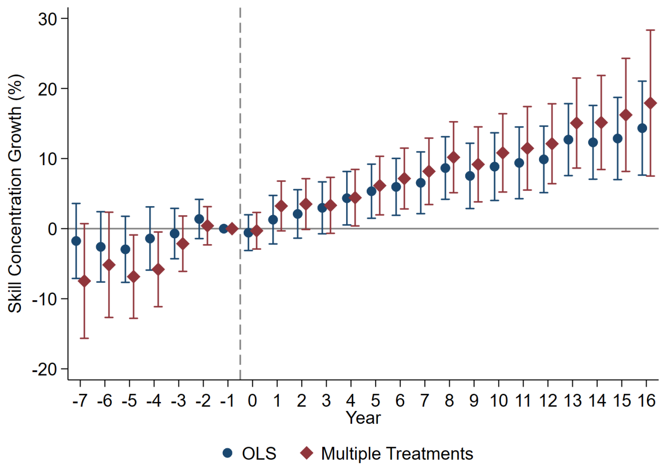

Note: High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Sample excludes observations with no workers at large firms. Multiple Treatments estimates (de Chaisemartin & D'Haultfœuille, 2024) restrict the control group to never-treated units and the treated group to places that saw an increase in the number of colleges. Year is relative to the arrival of a new college. Vertical bars represent the 95% confidence interval.

Figure A55: *Model results without labor search and an outside sector*

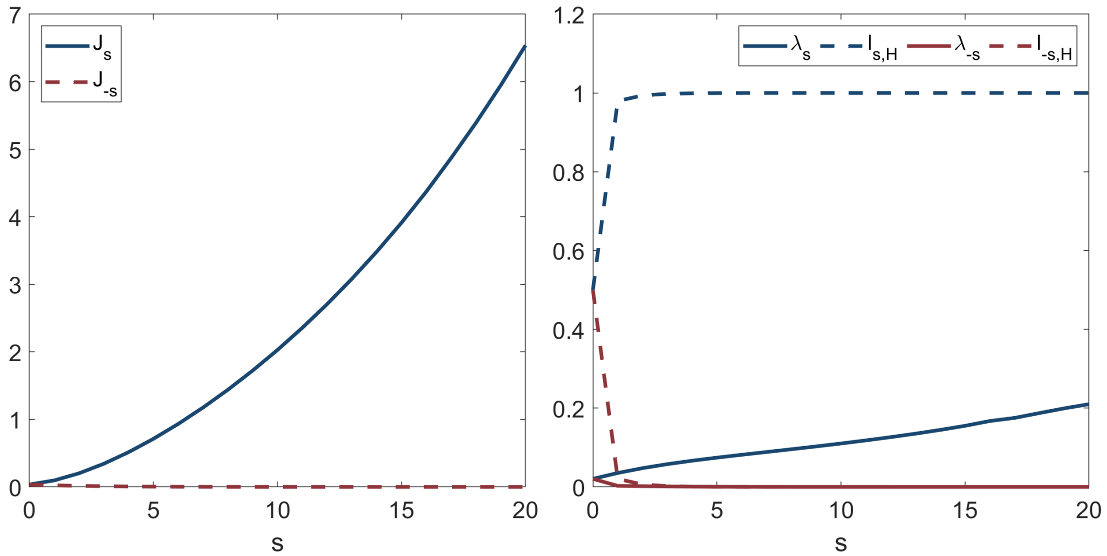

Note: J_s (J_{-s}) refers to the value function of the leader (follower). λ_s (λ_{-s}) refers to R&D investment by the leader (follower). $l_{s,HS}$ ($l_{-s,HS}$) refers to high-skilled labor hired by the leader (follower).

Figure A56: *Growth and high-skilled labor concentration in the model without labor search and an outside sector*

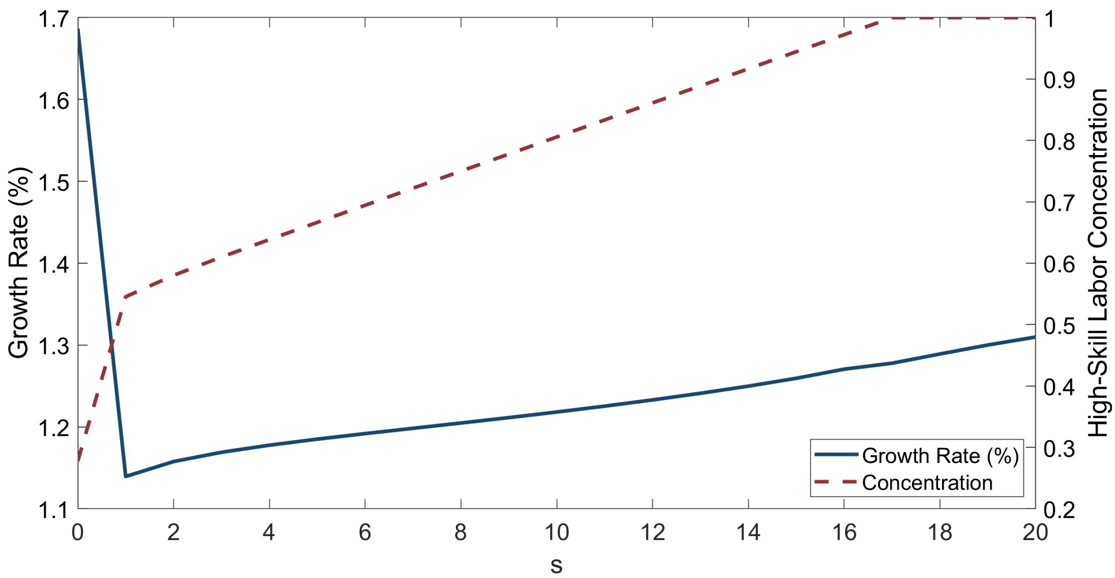

Figure A57: *Wage premium as a function of the gap s for different assumptions on high-skilled labor supply*

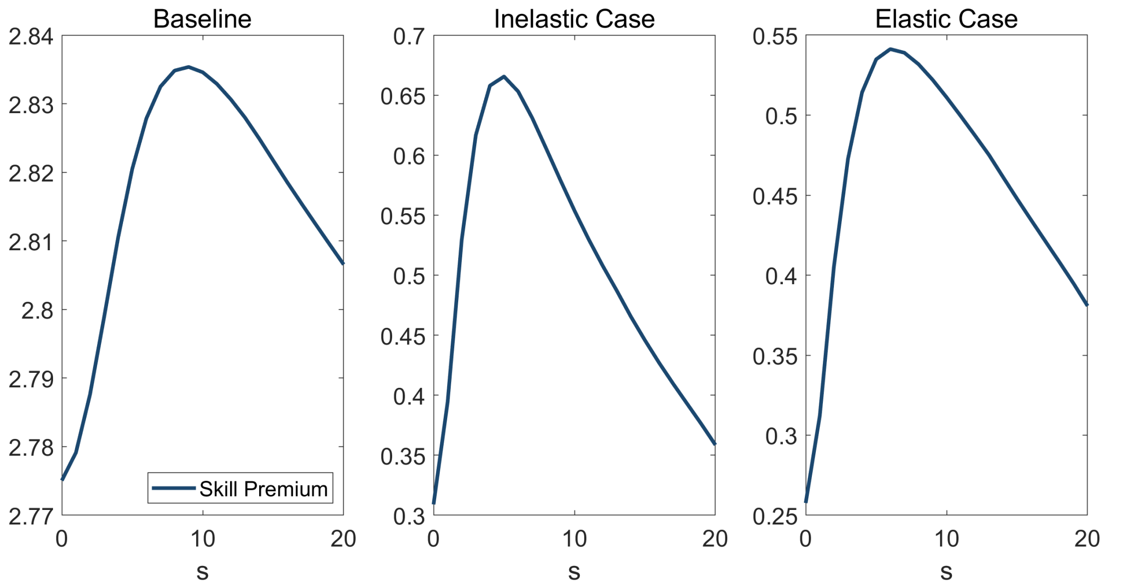

Note: Baseline refers to the baseline model estimation as described in Section 2.5. Inelastic Case refers to the case where high-skilled labor is paid at its marginal product, L_{HS} is set to the average employment rate of the baseline case, and we use baseline parameter estimates including for the firms' value function. Elastic Case refers to the case where high-skilled labor supply is elastic due to labor disutility. We assume $l_{s,HS} = (\frac{w_{s,HS}}{K})^{\zeta}$ where the Frisch elasticity ζ is set to 0.5 and $K = 3.62$ to match the average unemployment rate in the baseline case.

Figure A58: *Transition path for high-skill concentration given successive increases in high-skilled labor supply*

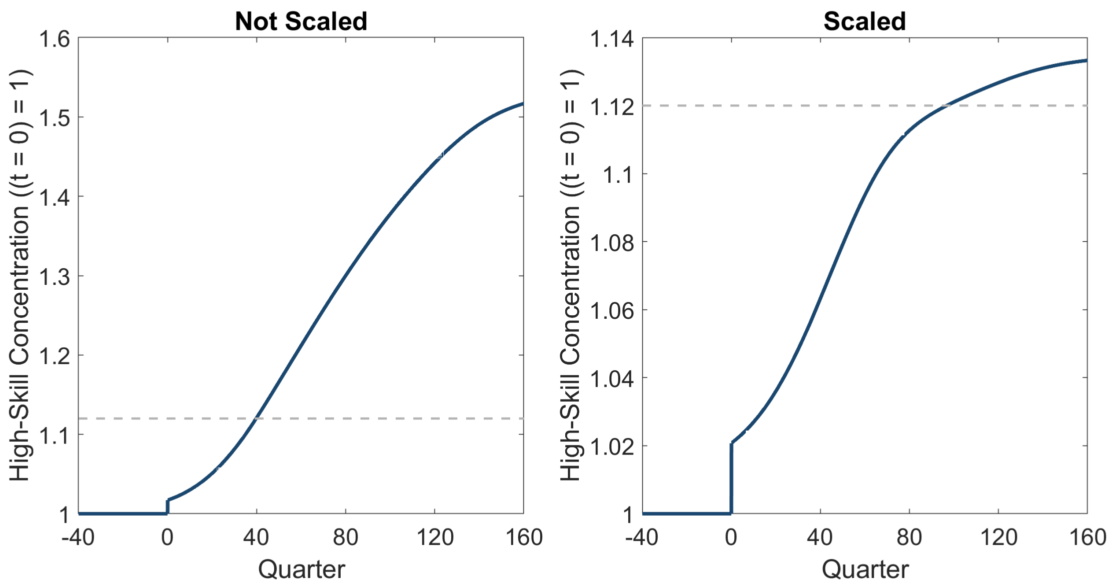

Note: Not Scaled refers to the unaltered skill concentration measure. Scaled refers to the adjusted measure which matches the variation range observed in the data along with a linear adjustment to match the 2.5× increase in L_{HS} with the 2× observed in the difference-in-differences setting.

Figure A59: *Transition path for high-skill concentration given successive increases in high-skilled labor supply (fully adjusted)*

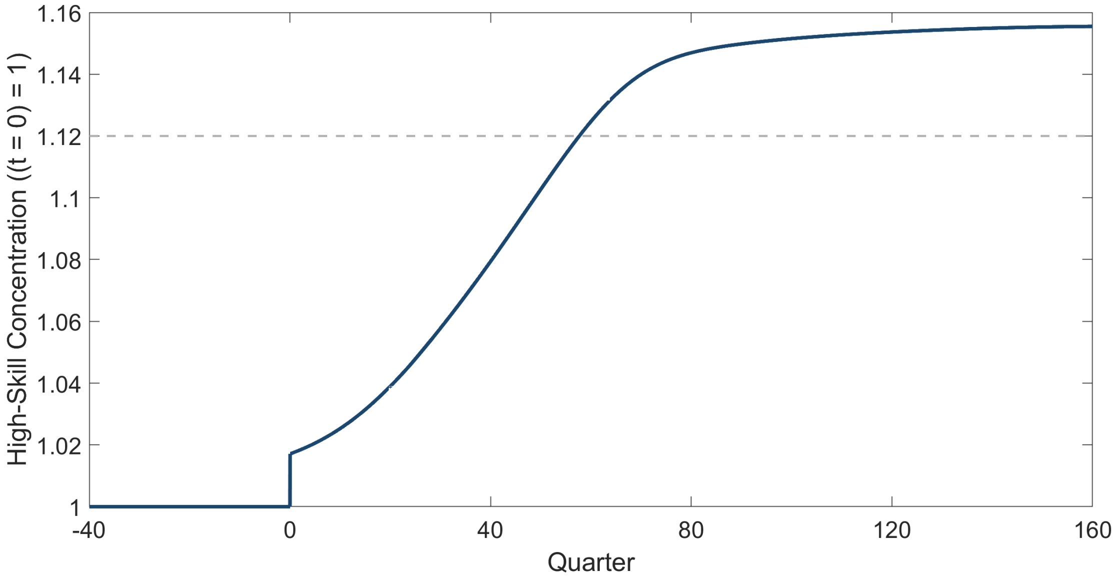

Note: High-skill concentration is the scaled measure which matches the variation range observed in the data along with a linear adjustment to match the 2.5× increase in L_{HS} with the 2× observed in the difference-in-differences setting. Model is adjusted so that the follower firm adjusts its hiring within 3 years after each shock while the leading firms overshoots its hiring, also within a period of 3 years.

Figure A60: *Transition path for the growth rate given successive increases in high-skilled labor supply*

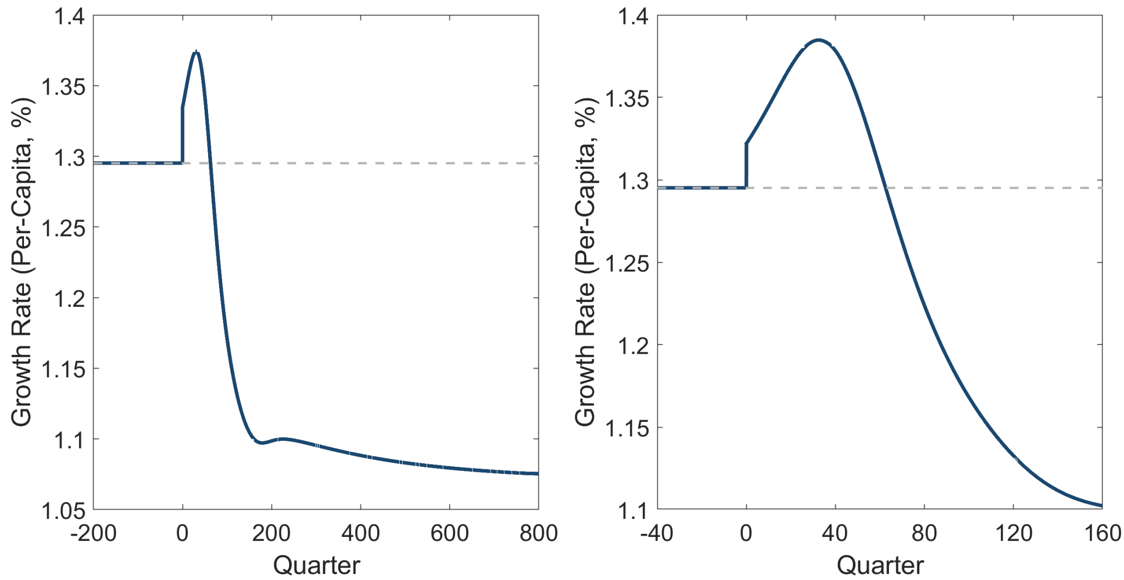

Note: Left-hand side plot shows the transition path over 200 years. Right-hand side plot zooms in on the first 40 years. Model is adjusted so that the follower firm adjusts its hiring within 3 years after each shock while the leading firms overshoots its hiring, also within a period of 3 years.

Table A.1: Summary statistics

	Mean	St. Dev.	Obs	Min	Max
GDP Per Capita Growth	0.031	0.181	74,209	-0.840	12.7
Skill Premium	2.114	0.601	68,691	0.293	11.7
Skill Premium - CT Workers	2.086	0.682	68,171	0.355	11.8
High-Skill Concentration	0.624	0.266	74,209	0.000	1.0
CT Worker Concentration	0.608	0.287	73,733	0.000	1.0
High-Skill Workers (th.)	2.153	31.338	74,209	0.010	2319.5
CT Workers (th.)	1.188	15.337	74,209	0.000	1076.3
Non-High-Skill Workers (th.)	9.453	86.051	74,209	0.001	5768.9
Electricity Consumption Growth	-0.013	0.243	74,209	-0.942	7.4
Real Wages (th.)	1.543	0.478	74,209	0.232	9.4
Population (mm.)	0.036	0.209	74,209	0.001	12.0
Total Workers (th.)	11.299	114.655	74,209	0.003	8042.9
Net New Workers Per Capita	0.002	0.017	74,209	-0.867	1.1
Minimum Wage Population Share	0.010	0.013	74,209	0.000	0.6
High-Skill Population Share	0.024	0.024	74,209	0.000	1.8
CT Workers Population Share	0.014	0.015	74,209	0.000	1.6
Informality Share (2000)	0.512	0.165	74,209	0.073	1.0

High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Skill premium is defined as the ratio between the average high-skill wage over the average non-high-skill wage within a municipality. CT (critical-thinking) workers are those with at least some college education who are also employed in occupations at the top skill quartile for one of the following: Math, Science, Critical Thinking, Active Learning, and Complex Problem Solving.

Table A.2: Summary statistics on municipalities

	Full-Sample		Matched	
	Control	Treatment	Control	Treatment
N	4,775	665	2,404	623
Population (th.)	15.6 (91.3)	117.0 (435.6)	34.9 (80.6)	81.8 (99.7)
Share Earning Min Wage (%)	11.9 (19.2)	4.2 (8.1)	3.4 (5.9)	3.4 (5.9)
Real Wage	1,213.2 (528.7)	1,565.9 (565.6)	1,544.3 (617.8)	1,564.6 (552.3)
Unemployment Rate (%)	9.9 (5.9)	13.2 (5.0)	12.8 (4.9)	13.0 (4.8)
Share Earning $<$ 0.25 x Min Wage (%)	41.2 (22.3)	25.8 (17.4)	26.9 (17.7)	24.9 (16.8)
Share in Agriculture (%)	4.2 (12.3)	1.8 (4.0)	2.8 (8.1)	1.8 (4.1)
Share in Mining (%)	0.7 (5.4)	0.6 (3.6)	0.8 (4.8)	0.6 (3.7)
Share in Manufacturing (%)	6.0 (12.6)	11.9 (13.3)	10.3 (14.9)	12.3 (13.5)
Share in Utilities/Transportation (%)	2.9 (10.2)	3.3 (4.4)	2.6 (6.4)	3.3 (4.5)
Share in Construction (%)	1.4 (7.9)	1.6 (4.1)	1.5 (6.5)	1.6 (3.6)
Share in Retail/Wholesale (%)	9.0 (14.2)	11.6 (8.3)	9.5 (10.7)	11.7 (8.2)
Share in FIRE (%)	8.8 (15.0)	12.8 (9.9)	10.6 (13.0)	12.7 (9.5)
Share in Public Sector (%)	57.2 (32.0)	32.1 (20.8)	46.8 (28.9)	31.6 (20.5)
Share in Other Services (%)	9.9 (16.0)	24.2 (14.7)	15.0 (16.7)	24.4 (14.7)
Illiterate Share (%)	4.0 (6.7)	2.6 (2.9)	2.5 (3.8)	2.4 (2.5)
$< 5^{\text{th}}$ Grade Share (%)	16.1 (12.7)	11.2 (8.4)	11.0 (8.2)	11.0 (8.0)
$= 5^{\text{th}}$ Grade Share (%)	16.2 (11.0)	14.4 (7.4)	17.0 (10.0)	14.5 (7.2)
6^{th} to $< 9^{\text{th}}$ Grade Share (%)	14.6 (9.0)	16.5 (6.2)	17.1 (7.3)	16.8 (6.2)
$= 9^{\text{th}}$ Grade Share (%)	13.4 (9.7)	16.6 (6.3)	16.3 (8.0)	16.7 (6.3)
Incomplete High-School Share (%)	6.7 (5.7)	9.2 (3.7)	8.0 (4.6)	9.3 (3.6)
High-School Share (%)	22.0 (14.1)	22.0 (9.2)	20.5 (9.9)	21.7 (9.0)
Incomplete College Share (%)	2.1 (3.7)	2.2 (1.7)	2.2 (2.0)	2.2 (1.6)
College+ Share (%)	4.0 (4.1)	5.4 (3.8)	5.3 (3.7)	5.4 (3.7)

Table reports sample means and standard errors, the latter in parenthesis. Statistics are for the year 2000. Sample match is on population level, share earning minimum wage or less, share who only completed the 5^{th} grade, unemployment rate, and illiteracy rate, all in 2000. Treated observations are matched to those in control using the coarsened matching method in Iacus et al. (2012). FIRE refers to Finance, Insurance, and Real Estate.

Table A.3: Balance Test on First Year of Treatment

	First Year of New College
	(1)
Log(Population)	0.015
	(0.053)
Share Earning Min Wage (%)	0.295**
	(0.096)
Real Wage	-0.093
	(0.059)
Illiterate Share (%)	-0.139
	(0.106)
$<$ 5$^{\text{th}}$ Grade Share (%)	0.118
	(0.142)
= 5$^{\text{th}}$ Grade Share (%)	0.019
	(0.121)
6$^{\text{th}}$ to $<$ 9$^{\text{th}}$ Grade Share (%)	0.100
	(0.103)
= 9$^{\text{th}}$ Grade Share (%)	-0.060
	(0.113)
Incomplete High-School Share (%)	-0.139
	(0.091)
High-School Share (%)	0.167
	(0.150)
Incomplete College Share (%)	-0.000
	(0.096)
Unemployment Rate (%)	0.023
	(0.058)
Share in Agriculture (%)	-0.031
	(0.111)
Share in Mining (%)	0.042
	(0.052)
Share in Manufacturing (%)	0.054
	(0.051)
Share in Utilities/Transportation (%)	0.043
	(0.086)
Share in Construction (%)	0.045
	(0.079)
Share in Retail/Wholesale (%)	0.032
	(0.082)
Share in FIRE (%)	0.056
	(0.069)
Share in Public Sector (%)	0.110
	(0.088)
N	731

All variables are demeaned and normalized to have unit variance. Independent variables refer to the year 2000. FIRE refers to Finance, Insurance, and Real Estate. Standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.4: Summary statistics on pre-treated municipalities by skill concentration level

	Treated Low Skill Concentration	Treated High Skill Concentration
N	92	570
Population (th.)	72.1 (77.2)	124.8 (469.0)
Share Earning Min Wage (%)	3.8 (4.7)	4.1 (8.4)
Real Wage	1,469.2 (492.4)	1,583.1 (575.7)
Unemployment Rate (%)	12.4 (4.6)	13.3 (5.0)
Share Earning $<$ 0.25 x Min Wage (%)	23.5 (16.0)	26.0 (17.6)
Share in Agriculture (%)	2.0 (4.1)	1.8 (4.0)
Share in Mining (%)	0.5 (3.5)	0.6 (3.6)
Share in Manufacturing (%)	10.6 (10.5)	12.1 (13.7)
Share in Utilities/Transportation (%)	3.5 (4.0)	3.3 (4.5)
Share in Construction (%)	2.3 (5.1)	1.5 (3.9)
Share in Retail/Wholesale (%)	13.9 (7.7)	11.3 (8.3)
Share in FIRE (%)	14.4 (11.3)	12.4 (9.3)
Share in Public Sector (%)	23.0 (17.8)	33.6 (20.7)
Share in Other Services (%)	29.8 (14.7)	23.3 (14.5)
Illiterate Share (%)	2.6 (2.6)	2.6 (3.0)
$< 5^{\text{th}}$ Grade Share (%)	10.9 (8.1)	11.3 (8.4)
$= 5^{\text{th}}$ Grade Share (%)	13.3 (6.9)	14.5 (7.5)
6^{th} to $< 9^{\text{th}}$ Grade Share (%)	16.9 (6.7)	16.4 (6.2)
$= 9^{\text{th}}$ Grade Share (%)	16.7 (5.6)	16.6 (6.4)
Incomplete High-School Share (%)	10.0 (3.4)	9.0 (3.7)
High-School Share (%)	22.0 (9.5)	22.0 (9.1)
Incomplete College Share (%)	2.4 (1.4)	2.2 (1.7)
College+ Share (%)	5.3 (3.4)	5.5 (3.8)

Table reports sample means and standard errors, the latter in parenthesis. Statistics are for the year 2000. High-skill concentration is the local share of high-skilled workers at large firms over total local supply. Low (high) concentration municipalities are defined as those below (above) the 14^{th} percentile of high-skill concentration averaged within the three-year period pre-treatment, for those treated, or below (above) the same concentration level averaged in the initial two periods for the non-treated so as to get similar threshold levels in both cases. FIRE refers to Finance, Insurance, and Real Estate.

Table A.5: Effect of large and small firm SSIV on high-skill and non-high-skill hiring, and energy consumption)

	# High-Skill		# CT Workers		# Non-High-Skill	Energy Consumption
	(1)	(2)	(3)	(4)	(5)	(6)
SSIV - Large Firms	5.078***	0.542***	3.033***	0.594***	-2.208***	0.0653
	(0.263)	(0.068)	(0.146)	(0.047)	(0.316)	(0.146)
SSIV - Small Firms	9.311***	7.413***	4.953***	4.296***	5.416*	-0.559
	(0.845)	(0.469)	(0.400)	(0.219)	(2.214)	(0.531)
SSIV - Total						-1.576***
						(0.236)
N	74,090	74,090	73,684	73,684	74,090	74,090
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
$\mathbb{1}_{\{HS\ Conc_{i,t-1}>p\}}$	0	1	0	1		

High-skill concentration is the local share of high-skilled workers at large firms over total local supply. CT (critical-thinking) workers are those with at least some college education who are also employed in occupations at the top skill quartile for one of the following: Math, Science, Critical Thinking, Active Learning, and Complex Problem Solving. # refers to workers per capita. SSIV - Total refers to the SSIV constructed using loans to both small and large firms. Energy Consumption refers to the capital proxy variable calculated using the change in local electricity consumption and is winsorized at the 1^{st} and 99^{th} percentiles. All specifications control for the lagged local sum of shares $s_{i,t-2}$. Local-level controls (all lagged to be contemporaneous with the SSIV): log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. $\mathbb{1}\{HSConc_{i,t-1} > p\}$ indicates whether the dependent variable and the SSIVs are interacted with a dummy for being below (0) or above (1) the high-skill concentration threshold which is set at the 16^{th} percentile. Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.6: Model estimation and moment fit (without the non-innovative, outside firm)

Parameter	Value	Parameter	Value
γ	1.05	κ	61,252
b	0.62	h_l	1.61
ρ	2,069	h_c	0.32
A_l	5.97		
A_λ	62.4		

Moments	Data	Model
Growth Rate (%)	1.31	1.01
Skill Premium, Large Firms	2.76	2.77
Labor Market Tightness	0.48	0.00
High-Skill Wage, Non-Large Firms	0.58	0.33
High-Skill Concentration	0.59	0.64
Firm Profitability	0.20	0.25
R&D Investing-to-Sales Ratio (%)	0.19	0.12
Cost-per-Hire	0.12	0.00
High-Skill Unemployment	0.19	1.00
Share of High-Skill Concentration > 80%	0.38	0.32

Table A.7: Model estimation and moment fit (1999-2004 period)

Parameter	Value	Parameter	Value
γ	1.05	κ	1.34
b	0.64	h_l	1.00
ρ	5,656	h_c	0.32
A_l	2.21	ν	0.20
A_λ	38.6		

Moments	Data	Model
Growth Rate (%)	1.30	1.29
Skill Premium, Large Firms	2.94	2.95
Labor Market Tightness	0.48	0.48
High-Skill Wage, Non-Large Firms	0.62	0.59
High-Skill Concentration	0.45	0.52
Firm Profitability	0.20	0.20
R&D Investing-to-Sales Ratio (%)	0.19	0.19
Cost-per-Hire	0.12	0.11
High-Skill Unemployment	0.19	0.22
Share of High-Skill Concentration $\leq$ 50%	0.47	0.51

Table A.8: Summary statistics for treated, last-treated, early-treated, and late-treated municipalities

	Last-Treated	Treated	Early Treated	Late Treated
N	32	133	340	325
Population (th.)	490.7 (1,980.4)	171.5 (414.8)	103.9 (156.2)	130.7 (602.4)
Share Earning Min Wage (%)	7.1 (7.2)	5.8 (6.9)	3.0 (5.8)	5.3 (9.9)
Real Wage	1,778.0 (530.9)	1,755.3 (468.2)	1,636.4 (587.3)	1,491.9 (532.8)
Unemployment Rate (%)	7.4 (2.1)	7.2 (2.8)	13.3 (5.1)	13.2 (4.9)
Share Earning $<$ 0.25 x Min Wage (%)	18.9 (14.6)	18.5 (14.7)	22.2 (15.1)	29.5 (18.9)
Share in Agriculture (%)	1.1 (1.4)	1.2 (2.8)	1.9 (3.9)	1.7 (4.1)
Share in Mining (%)	0.1 (0.3)	1.0 (4.3)	0.6 (3.1)	0.6 (4.0)
Share in Manufacturing (%)	9.5 (10.8)	8.9 (10.7)	12.4 (12.6)	11.3 (13.9)
Share in Utilities/Transportation (%)	2.8 (2.5)	2.1 (2.1)	3.3 (4.6)	3.3 (4.3)
Share in Construction (%)	2.2 (7.7)	1.4 (2.2)	1.7 (3.5)	1.5 (4.6)
Share in Retail/Wholesale (%)	9.8 (5.7)	10.8 (5.7)	11.5 (6.9)	11.7 (9.5)
Share in FIRE (%)	10.3 (6.6)	9.8 (6.3)	13.0 (8.8)	12.6 (10.9)
Share in Public Sector (%)	48.9 (23.6)	48.7 (22.1)	30.7 (18.7)	33.6 (22.8)
Share in Other Services (%)	15.4 (10.5)	16.2 (10.4)	24.8 (13.3)	23.5 (15.9)
Illiterate Share (%)	0.6 (0.5)	0.8 (0.8)	2.4 (3.2)	2.7 (2.7)
$<$ 5$^{\text{th}}$ Grade Share (%)	4.6 (3.3)	5.7 (4.8)	10.4 (7.4)	12.1 (9.2)
= 5$^{\text{th}}$ Grade Share (%)	5.9 (6.1)	5.5 (3.0)	14.0 (7.1)	14.8 (7.6)
6$^{\text{th}}$ to $<$ 9$^{\text{th}}$ Grade Share (%)	8.4 (3.5)	10.0 (5.1)	16.9 (5.4)	16.1 (6.9)
= 9$^{\text{th}}$ Grade Share (%)	13.2 (4.6)	13.9 (5.8)	17.5 (6.3)	15.6 (6.1)
Incomplete High-School Share (%)	8.9 (3.3)	9.1 (4.6)	9.7 (3.5)	8.6 (3.7)
High-School Share (%)	42.7 (8.8)	42.1 (10.4)	21.2 (8.3)	22.8 (9.9)
Incomplete College Share (%)	3.0 (1.3)	2.8 (1.4)	2.3 (1.5)	2.1 (1.9)
College+ Share (%)	12.8 (7.0)	10.1 (4.4)	5.6 (3.4)	5.2 (4.1)

Table reports sample means and standard errors, the latter in parenthesis. Statistics are for the year 2010 for Last Treated and Treated and for the year 2000 for Early Treated and Late Treated. Last-treated cohort receives treatment in 2019. Last Treated and Treated groups consist of municipalities that have not yet been treated in 2010. Early Treated places consist of those that will be treated by 2005. Late Treated places consist of those that will be treated after 2005. For the early- and late-treated cases, municipalities have not been treated yet.

Table A.9: Effect of high-skill concentration in large firms on local skill premium in places with high and low concentration

	Skill Premium						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}$=0 × $HS\ Conc._{t-1}$	9.286**	9.669**	8.294**	9.630**	9.641**	9.692**	9.107**
	(2.964)	(2.989)	(2.936)	(3.174)	(2.980)	(3.024)	(3.107)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}$=1 × $HS\ Conc._{t-1}$	-1.280**	-1.358**	-1.237**	-1.349*	-1.381**	-1.341**	-1.260*
	(0.471)	(0.478)	(0.477)	(0.533)	(0.481)	(0.482)	(0.534)
N	68,582	68,582	68,582	68,582	68,582	68,582	68,582
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes
Informality			Yes				
Non-High-Skill				Yes			Yes
Capital Proxy					Yes		Yes
Employment HHI						Yes	
LIML							
Joint F-statistic	42.9	42.1	37.1	14.9	17.6	31.9	6.7
J-test, p-value	0.51	0.96	0.08	.	0.55	0.60	.
OP F-statistic, $\mathbb{1}_{\{HS\ Conc._{t-1}>p\}} = 0$	32.9	32.8	32.3	35.0	31.7	37.2	36.2
OP Critical Value, $\mathbb{1}_{\{HS\ Conc._{t-1}>p\}} = 0$	17.5	17.5	17.2	23.1	16.4	23.1	23.1
OP F-statistic, $\mathbb{1}_{\{HS\ Conc._{t-1}>p\}} = 1$	72.7	71.4	63.9	85.2	65.6	103.8	79.4
OP Critical Value, $\mathbb{1}_{\{HS\ Conc._{t-1}>p\}} = 1$	10.6	10.6	9.2	23.1	8.5	23.1	23.1

High-skill concentration (*HS Conc*) is the local share of high-skilled workers at large firms over total local supply. Threshold p is set at the 10^{th} percentile. Skill premium is defined as the ratio between the average high-skill wage over the average non-high-skill wage within a municipality and it is winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(4) use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as IVs, the former interacted with $\mathbb{1}\{HSConc_{i,t-1} > p\}$ and the latter with $\mathbb{1}\{\{HSConc_{i,t-1} > p\} = 0\}$. Columns (5) and (6) add a SSIV calculated using both small and large shocks together, and using total employment shares. All specifications control for the lagged local sum of shares $s_{i,t-2}$. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. Non-High-Skill refers to total non-high-skill hiring. Capital Proxy refers to the proxy variable calculated using the change in local electricity consumption. Employment HHI refers to the HHI measure of concentration calculated for total employment. J-test refers to the overidentification test in Hansen (1982). OP F-statistic and Critical Value refer, respectively, to the Olea-Pflueger effective F-statistic and the critical value for a 5% significance level and a 10% "worst-case" bias. Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.10: Summary statistics of shocks and shares

	Mean	St. Dev.	IQR	Max	Obs
Shock - Large Firms	0.908	2.216	1.102	7.794	594
Shock - Large Firms - Residual	0.000	1.117	0.530	8.214	594
Share - Large Firms	0.001	0.007	0.000	0.051	594
Shock - Small Firms	0.263	0.770	0.804	2.887	622
Shock - Small Firms - Residual	0.000	0.660	0.615	3.122	622
Share - Small Firms	0.001	0.002	0.001	0.011	622
Effective Sample Size - Large	.	.	.	.	26
Effective Sample Size - Small	.	.	.	.	190
Number of Sectors	.	.	.	.	60

Shocks consist of the yearly change at the national level of BNDES loans by sector and firm size. Shares consist of the local-level lagged high-skill employment shares by sector and firm size. Shock statistics are weighted by the shares. Residual statistics refer to shocks residualized on year fixed-effects. The effective sample size is measured as the inverse renormalized Herfindahl index of the shares.

Table A.11: Shock balance tests and pre-trend tests

	GDP Growth		Skill Premium		Log(Wage)		Log(Population)		% High-Skill		% Min. Wage		Net Hiring	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Shock - Large Firms	0.0942		0.134		0.171		0.0145		-0.00764		-0.0382*		0.0198	
	(0.061)		(0.072)		(0.092)		(0.025)		(0.024)		(0.019)		(0.026)	
Shock - Small Firms		-0.0560		0.0130		-0.0189		-0.0317		0.0379		-0.00920		0.000947
		(0.047)		(0.077)		(0.095)		(0.055)		(0.019)		(0.016)		(0.026)
N	593	622	593	622	593	622	593	622	593	622	593	622	593	622
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Dependent variable is defined as $\overline{q}_{n,t-3} = \frac{\sum_i s_{int} q_{i,t-3}}{\sum_i s_{int}}$ where s_{int} are the exposure shares and $q_{i,t-3}$ is one of the controls used in Equation 1.3. Regressions are weighted by sector high-skill employment shares. In Columns (1)-(4), $q_{i,t-3}$ is replaced by $y_{i,t-3}$ where $y_{i,t-3}$ is GDP per-capita growth in Columns (1)-(2) and skill premium in Columns (3)-(4), both winsorized at the 1^{st} and 99^{th} percentiles. Local-level controls (all lagged to be contemporaneous with shocks): log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. In each specification the variable used as the dependent variable is excluded from the list of controls. Sector-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.12: Sector-level survey data balance check

Balance Variable	Coef	SE	Obs.
Revenue Growth	-0.161	(0.134)	495
Value Added Growth	-0.026	(0.028)	480
Wages-to-Value Added Ratio	0.041	(0.046)	495
Intermediate Inputs-to-Value Added Ratio	-0.014	(0.033)	495
Fuel and Electricity-to-Value Added Ratio	0.039	(0.089)	473
Production Workers' Share of Employment (on 12/31)	0.020	(0.050)	359
Production Workers' Share of Employment (yearly avg.)	0.054	(0.055)	359

Table reports the regression coefficients of each sector-level variable on shocks $g_{n,t}$ weighted by high-skill shares and controlling for year fixed-effects. Variables are set to the shocks' initial period ($t-1$). Standard errors are sector-clustered. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.13: Effect of high-skill concentration in large firms on local GDP growth (polynomial regressors and instruments)

	GDP Growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$HS\ Conc._{t-1}^{3}$	-0.463***	-0.513***	-0.503***	-0.524**	-0.455***	-0.455***	-0.410**	-0.400*
	(0.109)	(0.116)	(0.134)	(0.159)	(0.113)	(0.137)	(0.152)	(0.166)
$HS\ Conc._{t-1}$	0.523***	0.553***	0.546***	0.554***	0.539***	0.537***	0.531***	0.531***
	(0.098)	(0.097)	(0.095)	(0.098)	(0.100)	(0.114)	(0.103)	(0.105)
Vertex	0.61	0.60	0.60	0.59	0.63	0.63	0.66	0.67
Std. Error	0.06	0.06	0.06	0.08	0.06	0.06	0.10	0.11
N	74,090	74,090	74,090	74,090	74,090	74,090	74,090	74,090
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Informality			Yes					
Non-High-Skill				Yes			Yes	Yes
Capital Proxy					Yes		Yes	Yes
Employment HHI						Yes		
LIML								Yes
Joint F-statistic	10.3	8.9	7.4	11.8	11.8	8.1	7.3	7.3
J-test, p-value	0.73	0.47	0.84	0.22	0.37	0.36	0.21	0.21

High-skill concentration (*HS Conc*) is the local share of high-skilled workers at large firms over total local supply. GDP Growth is real per-capita local GDP growth winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(4) use $B_{i,t-2,large}$, $B^3_{i,t-2,large}$, $B_{i,t-2,small}$, and $B^3_{i,t-2,small}$ as instruments (shocks winsorized at the 5^{th} and 95^{th} percentiles). Columns (5)-(7) add a SSIV calculated using both small and large shocks together, and using total employment shares. Vertex refers to the point in the domain where the derivative with respect to the regressor of interest is zero (i.e. the point where the slope changes sign). Std. Error refers to the standard error of the vertex point-estimate. All specifications control for the lagged local sum of shares $s_{i,t-2}$. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. Non-High-Skill refers to total non-high-skill hiring. Capital Proxy refers to the proxy variable calculated using the change in local electricity consumption. Employment HHI refers to the HHI measure of concentration calculated for total employment. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.14: Effect of high-skill concentration in large firms on local GDP growth (polynomial instruments)

	GDP Growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}$=0 × $HS\ Conc._{t-1}$	0.779***	0.803***	0.836***	0.836***	0.782***	0.758***	0.841***	0.861***
	(0.181)	(0.184)	(0.186)	(0.192)	(0.184)	(0.188)	(0.190)	(0.198)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}$=1 × $HS\ Conc._{t-1}$	-0.541***	-0.611***	-0.645***	-0.540*	-0.547**	-0.511**	-0.401	-0.375
	(0.162)	(0.172)	(0.193)	(0.242)	(0.166)	(0.181)	(0.232)	(0.275)
N	74,090	74,090	74,090	74,090	74,090	74,090	74,090	74,090
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Informality			Yes					
Non-High-Skill				Yes			Yes	Yes
Capital Proxy					Yes		Yes	Yes
Employment HHI						Yes		
LIML								Yes
Joint F-statistic	21.0	19.3	15.1	11.7	15.7	14.5	7.6	7.6
J-test, p-value	0.36	0.14	0.49	0.05	0.14	0.13	0.10	0.10

High-skill concentration (*HS Conc*) is the local share of high-skilled workers at large firms over total local supply. Threshold p is set at the 25^{th} percentile. GDP Growth is real per-capita local GDP growth winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(4) use $B_{i,t-2,large}$, $B^3_{i,t-2,large}$, $B_{i,t-2,small}$, and $B^3_{i,t-2,small}$ as instruments (shocks winsorized at the 5^{th} and 95^{th} percentiles). Columns (5)-(7) add a SSIV calculated using both small and large shocks together, and using total employment shares. Vertex refers to the point in the domain where the derivative with respect to the regressor of interest is zero (i.e. the point where the slope changes sign). Std. Error refers to the standard error of the vertex point-estimate. All specifications control for the lagged local sum of shares $s_{i,t-2}$. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. Non-High-Skill refers to total non-high-skill hiring. Capital Proxy refers to the proxy variable calculated using the change in local electricity consumption. Employment HHI refers to the HHI measure of concentration calculated for total employment. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.15: Effect of concentration of critical thinking workers in large firms on local GDP growth in places with high and low concentration

	GDP Growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{1}_{\{CT\ Conc._{t-1}>p\}}$=0 × $CT\ Conc._{t-1}$	1.683***	1.580***	1.709***	2.030***	1.577***	1.589***	2.464**	2.486**
	(0.475)	(0.475)	(0.503)	(0.612)	(0.476)	(0.474)	(0.768)	(0.779)
$\mathbb{1}_{\{CT\ Conc._{t-1}>p\}}$=1 × $CT\ Conc._{t-1}$	-0.354**	-0.339**	-0.356**	-0.395**	-0.340**	-0.343**	-0.485**	-0.489**
	(0.125)	(0.125)	(0.132)	(0.141)	(0.126)	(0.126)	(0.176)	(0.178)
N	73,684	73,684	73,684	73,684	73,684	73,684	73,684	73,684
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Informality			Yes					
Non-High Critical-Thinking				Yes			Yes	Yes
Capital Proxy					Yes		Yes	Yes
Employment HHI						Yes		
LIML								Yes
Joint F-statistic	24.1	22.0	25.8	13.6	15.1	17.3	4.1	4.1
J-test, p-value	0.20	0.08	0.21	0.08	0.08	0.09	0.46	0.46

High critical-thinking concentration ($CT\ Conc$) is the local share of high critical-thinking people working at large firms over total local supply. High critical-thinking workers are those with at least some college education who are employed in occupations at the top skill quartile for one of the following: Math, Science, Critical Thinking, Active Learning, and Complex Problem Solving. High critical-thinking concentration threshold p is set at the 16^{th} percentile. GDP Growth is real per-capita local GDP growth winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(4) use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as IVs (shocks winsorized at the 5^{th} and 95^{th} percentiles), each interacted with $\mathbb{1}\{CT\ Conc_{i,t-1} > p\}$. Columns (5)-(7) add a SSIV calculated using both small and large shocks together, and using total employment shares. All specifications control for the lagged local sum of shares $s_{i,t-2}$. Lagged local-level controls: log of population, log of average real wage, the percentage of high critical-thinking workers in the population, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. Non-High Critical-Thinking refers to the total non-high critical-thinking hiring instrumented with the SSIVs. Capital Proxy refers to the proxy variable calculated using the change in local electricity consumption. Employment HHI refers to the HHI measure of concentration calculated for total employment. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.16: Effect of concentration of critical thinking workers in large firms on local skill premium in places with high and low concentration

	Skill Premium						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\mathbb{1}_{\{CT\ Conc._{t-1}>p\}}$=0 × $CT\ Conc._{t-1}$	44.42***	44.66***	38.03***	42.48***	44.06***	44.95***	35.18**
	(11.481)	(11.585)	(10.517)	(12.251)	(11.533)	(12.014)	(11.908)
$\mathbb{1}_{\{CT\ Conc._{t-1}>p\}}$=1 × $CT\ Conc._{t-1}$	-2.037*	-2.089*	-1.818*	-1.902	-2.121*	-2.108*	-1.332
	(0.961)	(0.980)	(0.913)	(1.003)	(0.970)	(0.990)	(1.031)
N	68,122	68,122	68,122	68,122	68,122	68,122	68,122
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes
Informality			Yes				
Non-High Critical-Thinking				Yes			Yes
Capital Proxy					Yes		Yes
Employment HHI						Yes	
LIML							
Joint F-statistic	13.2	13.1	11.2	15.1	15.6	7.9	1.9
J-test, p-value	0.66	0.50	0.04	.	0.15	0.22	.

High critical-thinking concentration ($CT\ Conc$) is the local share of high critical-thinking people working at large firms over total local supply. High critical-thinking workers are those with at least some college education who are employed in occupations at the top skill quartile for one of the following: Math, Science, Critical Thinking, Active Learning, and Complex Problem Solving. High critical-thinking concentration threshold p is set at the 10^{th} percentile. Skill premium is defined as the ratio between the average high critical-thinking wage over the average wage of non-high critical-thinking workers and is winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(4) use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as IVs (shocks winsorized at the 5^{th} and 95^{th} percentiles), the former interacted with $\mathbb{1}\{CT\ Conc_{i,t-1} > p\}$ and the latter with $\mathbb{1}\{\{CT\ Conc_{i,t-1} > p\} = 0\}$. Columns (5) and (6) add a SSIV calculated using both small and large shocks together, and using total employment shares. All specifications control for the lagged local sum of shares $s_{i,t-2}$ and are weighted by the twice lagged log of local population. Lagged local-level controls: log of population, log of average real wage, the percentage of high critical-thinking workers in the population, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. Non-High Critical-Thinking refers to the total non-high critical-thinking hiring instrumented with the SSIVs. Capital Proxy refers to the proxy variable calculated using the change in local electricity consumption. Employment HHI refers to the HHI measure of concentration calculated for total employment. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.17: Effect of high-skill concentration in large firms on local skill premium in places with low and mid-level concentration

	Skill Premium				
	(1)	(2)	(3)	(4)	(5)
$\mathbb{1}_{\{HS\ Conc._{t-1}<p_1\}}$=1 × $HS\ Conc._{t-1}$	9.878***	10.47***	10.75***	11.01***	9.412**
	(2.877)	(2.915)	(2.912)	(2.942)	(2.926)
$\mathbb{1}_{\{p_1<HS\ Conc._{t-1}<p_2\}}$=1 × $HS\ Conc._{t-1}$	-4.960**	-3.774***	-2.868***	-2.493***	-2.762***
	(1.733)	(1.098)	(0.795)	(0.702)	(0.681)
Top Threshold p_2	30%	40%	50%	60%	60%
N	68,582	68,582	68,582	68,582	68,582
Time FE	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
# IVs	4	4	4	4	3
Joint F-statistic	19.2	24.7	36.1	38.3	55.1
J-test, p-value	0.10	0.08	0.04	0.01	0.21

High-skill concentration (*HS Conc*) is the local share of high-skilled workers at large firms over total local supply. High-skill concentration threshold p_1 is set at the 10^{th} percentile. Skill premium is defined as the ratio between the average high-skill wage over the average non-high-skill wage within a municipality and is winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(4) use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as IVs, each interacted with $(\mathbb{1}\{HSConc_{i,t-1} < p_1\} = 1)$ and $(\mathbb{1}\{p_1 < HSConc_{i,t-1} < p_2\} = 1)$. Column (5) removes $B_{i,t-2,small}(\mathbb{1}\{p_1 < HSConc_{i,t-1} < p_2\} = 1)$ from the set of instruments. All specifications control for the lagged local sum of shares $s_{i,t-2}$. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.18: Effect of high-skill concentration in large firms on local skill premium in places with high and low concentration (non-interacted $B_{i,t-2,small}$)

	Skill Premium						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=0 \times HS\ Conc._{t-1}$	13.73***	13.85***	13.86***	14.00***	13.83***	13.82***	13.95***
	(2.532)	(2.503)	(2.641)	(3.033)	(2.494)	(2.500)	(3.390)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=1 \times HS\ Conc._{t-1}$	-2.347**	-2.422**	-2.434**	-2.456*	-2.418**	-2.412**	-2.446*
	(0.903)	(0.904)	(0.931)	(0.997)	(0.903)	(0.908)	(1.064)
N	68,582	68,582	68,582	68,582	68,582	68,582	68,582
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes
Informality			Yes				
Non-High-Skill				Yes			Yes
Capital Proxy					Yes		Yes
Employment HHI						Yes	
LIML							
Joint F-statistic	16.5	17.1	15.2	11.3	17.5	15.0	4.7
J-test, p-value	0.66	0.92	0.18	.	0.95	0.95	.

High-skill concentration (*HS Conc*) is the local share of high-skilled workers at large firms over total local supply. Threshold p is set at the 15^{th} percentile. Skill premium is defined as the ratio between the average high-skill wage over the average non-high-skill wage within a municipality and it is winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(4) use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as IVs, the former interacted with $\mathbb{1}\{HSConc_{i,t-1} > p\}$ and the latter with $\mathbb{1}\{\{HSConc_{i,t-1} > p\} = 0\}$. Columns (5) and (6) add a SSIV calculated using both small and large shocks together, and using total employment shares. All specifications control for the lagged local sum of shares $s_{i,t-2}$. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. Non-High-Skill refers to total non-high-skill hiring. Capital Proxy refers to the proxy variable calculated using the change in local electricity consumption. Employment HHI refers to the HHI measure of concentration calculated for total employment. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.19: Effect of high-skill concentration in large firms on local GDP growth in places with high and low concentration (non-tradables only)

	GDP Growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=0 \times HS\ Conc._{t-1}$	1.270***	1.207***	1.259***	1.244***	1.209***	1.213***	1.290***	1.385***
	(0.276)	(0.274)	(0.280)	(0.287)	(0.271)	(0.271)	(0.316)	(0.382)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=1 \times HS\ Conc._{t-1}$	-0.576**	-0.561**	-0.583**	-0.571**	-0.559**	-0.565**	-0.578**	-0.615**
	(0.187)	(0.184)	(0.184)	(0.186)	(0.182)	(0.180)	(0.188)	(0.205)
N	73,879	73,879	73,879	73,879	73,879	73,879	73,879	73,879
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Informality			Yes					
Non-High-Skill				Yes			Yes	Yes
Capital Proxy					Yes		Yes	Yes
Employment HHI						Yes		
LIML								Yes
Joint F-statistic	16.3	16.0	16.5	11.7	11.7	14.4	4.8	4.8
J-test, p-value	0.23	0.15	0.20	0.06	0.15	0.16	0.07	0.08

High-skill concentration ($HS\ Conc$) is the local share of high-skilled workers at large firms over total local supply. Threshold p is set at the 16^{th} percentile. GDP Growth is real per-capita local non-tradables GDP growth winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(4) use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as IVs, each interacted with $\mathbb{1}\{HSConc_{i,t-1} > p\}$. Columns (5)-(7) add a SSIV calculated using both small and large shocks together, and using total employment shares. All specifications control for the lagged local sum of shares $s_{i,t-2}$. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. Non-High-Skill refers to total non-high-skill hiring. Capital Proxy refers to the proxy variable calculated using the change in local electricity consumption. Employment HHI refers to the HHI measure of concentration calculated for total employment. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.20: Effect of high-skill concentration in large firms on local skill premium in places with high and low concentration (non-tradables only)

	Skill Premium						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=0 \times HS\ Conc._{t-1}$	11.51***	11.57***	10.77**	11.60***	10.60**	11.69***	10.37**
	(3.376)	(3.387)	(3.273)	(3.402)	(3.253)	(3.503)	(3.279)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=1 \times HS\ Conc._{t-1}$	-2.366*	-2.422*	-2.322*	-2.447*	-2.381*	-2.129	-2.238
	(1.151)	(1.153)	(1.098)	(1.211)	(1.092)	(1.213)	(1.148)
N	68,466	68,466	68,466	68,466	68,466	68,466	68,466
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes
Informality			Yes				
Non-High-Skill				Yes			Yes
Capital Proxy					Yes		Yes
Employment HHI						Yes	
LIML							
Joint F-statistic	16.7	16.1	15.2	14.2	13.8	13.6	7.4
J-test, p-value	0.67	0.92	0.17	.	0.63	0.17	.

High-skill concentration (*HS Conc*) is the local share of high-skilled workers at large firms over total local supply. Threshold p is set at the 13^{th} percentile. Skill premium is defined as the ratio between the average high-skill wage over the average non-high-skill wage within a municipality and it is winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(4) use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as IVs, the former interacted with $\mathbb{1}\{HSConc_{i,t-1} > p\}$ and the latter with $\mathbb{1}\{\{HSConc_{i,t-1} > p\} = 0\}$. Columns (5) and (6) add a SSIV calculated using both small and large shocks together, and using total employment shares. All specifications control for the lagged local sum of shares $s_{i,t-2}$ and are weighted by the twice lagged log of local population. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. Non-High-Skill refers to total non-high-skill hiring. Capital Proxy refers to the proxy variable calculated using the change in local electricity consumption. Employment HHI refers to the HHI measure of concentration calculated for total employment. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.21: Effect of high-skill concentration in large firms on local GDP growth in places with high and low concentration (weighted by log of population)

	GDP Growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=0 \times HS\ Conc._{t-1}$	0.813**	0.705**	0.855**	1.044**	0.688*	0.732*	1.199**	1.196**
	(0.259)	(0.261)	(0.284)	(0.330)	(0.268)	(0.290)	(0.401)	(0.400)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=1 \times HS\ Conc._{t-1}$	-0.559**	-0.565**	-0.600***	-0.681***	-0.554**	-0.564**	-0.716**	-0.715**
	(0.180)	(0.177)	(0.182)	(0.201)	(0.177)	(0.182)	(0.219)	(0.218)
N	74,090	74,090	74,090	74,090	74,090	74,090	74,090	74,090
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Informality			Yes					
Non-High-Skill				Yes			Yes	Yes
Capital Proxy					Yes		Yes	Yes
Employment HHI						Yes		
LIML								Yes
Joint F-statistic	19.0	17.9	17.7	12.9	13.4	11.5	4.6	4.6
J-test, p-value	0.12	0.12	0.21	0.72	0.11	0.10	0.87	0.87

High-skill concentration ($HS\ Conc$) is the local share of high-skilled workers at large firms over total local supply. Threshold p is set at the 24^{th} percentile. GDP Growth is real per-capita local GDP growth winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(4) use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as IVs, each interacted with $\mathbb{1}\{HSConc_{i,t-1} > p\}$. Columns (5)-(7) add a SSIV calculated using both small and large shocks together, and using total employment shares. All specifications control for the lagged local sum of shares $s_{i,t-2}$ and are weighted by the twice lagged log of local population. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. Non-High-Skill refers to total non-high-skill hiring. Capital Proxy refers to the proxy variable calculated using the change in local electricity consumption. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.22: Effect of high-skill concentration in large firms on local skill premium in places with high and low concentration (weighted by log of population)

	Skill Premium						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=0 \times HS\ Conc._{t-1}$	9.767**	10.08**	8.867**	10.27**	10.11**	10.17**	9.611**
	(3.123)	(3.142)	(3.097)	(3.382)	(3.142)	(3.198)	(3.272)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=1 \times HS\ Conc._{t-1}$	-1.432**	-1.502**	-1.395**	-1.546**	-1.536**	-1.488**	-1.415**
	(0.477)	(0.484)	(0.486)	(0.555)	(0.488)	(0.490)	(0.549)
N	68,582	68,582	68,582	68,582	68,582	68,582	68,582
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes
Informality			Yes				
Non-High-Skill				Yes			Yes
Capital Proxy					Yes		Yes
Employment HHI						Yes	
LIML							
Joint F-statistic	35.9	35.5	30.9	11.1	18.9	26.4	5.0
J-test, p-value	0.75	0.82	0.09	.	0.57	0.70	.

High-skill concentration ($HS\ Conc$) is the local share of high-skilled workers at large firms over total local supply. Threshold p is set at the 10^{th} percentile. Skill premium is defined as the ratio between the average high-skill wage over the average non-high-skill wage within a municipality and it is winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(4) use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as IVs, the former interacted with $\mathbb{1}\{HSConc_{i,t-1} > p\}$ and the latter with $\mathbb{1}\{\{HSConc_{i,t-1} > p\} = 0\}$. Columns (5) and (6) add a SSIV calculated using both small and large shocks together, and using total employment shares. All specifications control for the lagged local sum of shares $s_{i,t-2}$ and are weighted by the twice lagged log of local population. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. Non-High-Skill refers to total non-high-skill hiring. Capital Proxy refers to the proxy variable calculated using the change in local electricity consumption. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.23: Effect of high-skill concentration in large firms on local GDP growth in places with high and low concentration (twice lagged shares)

	GDP Growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=0 \times HS\ Conc._{t-1}$	1.204**	1.105**	1.218**	1.565***	1.055**	1.056**	1.601**	1.631**
	(0.375)	(0.371)	(0.394)	(0.459)	(0.363)	(0.366)	(0.499)	(0.512)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=1 \times HS\ Conc._{t-1}$	-0.462**	-0.448**	-0.477**	-0.586**	-0.429**	-0.421**	-0.600**	-0.611**
	(0.157)	(0.154)	(0.163)	(0.181)	(0.154)	(0.156)	(0.192)	(0.197)
N	73,864	73,864	73,864	73,864	73,864	73,864	73,864	73,864
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Informality			Yes					
Non-High-Skill				Yes			Yes	Yes
Capital Proxy					Yes		Yes	Yes
Employment HHI						Yes		
LIML								Yes
Joint F-statistic	14.2	14.3	12.7	12.7	13.9	13.5	5.3	5.3
J-test, p-value	0.10	0.09	0.17	0.52	0.12	0.12	0.54	0.54

High-skill concentration (*HS Conc*) is the local share of high-skilled workers at large firms over total local supply. Threshold p is set at the 16^{th} percentile. GDP Growth is real per-capita local GDP growth winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(4) use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as IVs, each interacted with $\mathbb{1}\{HSConc_{i,t-1} > p\}$. Columns (5)-(7) add a SSIV calculated using both small and large shocks together, and using total employment shares. All specifications control for the lagged local sum of shares $s_{i,t-2}$. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. Non-High-Skill refers to total non-high-skill hiring. Capital Proxy refers to the capital proxy variable calculated using the change in local electricity consumption for the service sector instrumented with the SSIVs. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.24: Effect of high-skill concentration in large firms on local skill premium in places with high and low concentration (twice lagged shares)

	Skill Premium						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\mathbb{1}_{\{HS\ Conc._{t-1}<p_1\}=1} \times HS\ Conc._{t-1}$	8.088*	8.660**	6.955*	11.49**	8.582**	7.126*	12.09**
	(3.307)	(3.357)	(3.311)	(4.037)	(3.319)	(3.128)	(4.606)
$\mathbb{1}_{\{p_1<HS\ Conc._{t-1}<p_2\}=1} \times HS\ Conc._{t-1}$	-1.429*	-1.498*	-1.247	-2.054**	-1.492*	-1.136	-2.172**
	(0.669)	(0.682)	(0.673)	(0.780)	(0.681)	(0.612)	(0.779)
N	68,363	68,363	68,363	68,363	68,363	68,363	68,363
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes
Informality			Yes				
Non-High-Skill				Yes			Yes
Capital Proxy					Yes		Yes
Employment HHI						Yes	
LIML							
Joint F-statistic	29.1	28.9	25.6	6.1	16.7	22.3	2.6
J-test, p-value	0.10	0.20	0.19	0.82	0.19	0.00	.

High-skill concentration (*HS Conc*) is the local share of high-skilled workers at large firms over total local supply. High-skill concentration thresholds p_1 and p_2 are set at the 10^{th} and 75^{th} percentiles, respectively. Skill premium is defined as the ratio between the average high-skill wage over the average non-high-skill wage within a municipality and it is winsorized at the 1^{st} and 99^{th} percentiles. Columns (1)-(3) use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as instruments, both interacted with $(\mathbb{1}\{HSConc_{i,t-1} < p_1\} = 1)$ and the large one also with $(\mathbb{1}\{p_1 < HSConc_{i,t-1} < p_2\} = 1)$. Column (4) adds to the IV set a SSIV calculated using both small and large shocks together, and using total high-skill employment shares as exposure shares. Columns (5) and (6) add a SSIV calculated using both small and large shocks together, and using total employment shares. All specifications control for the lagged local sum of shares $s_{i,t-2}$. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. Non-High-Skill refers to total non-high-skill hiring. Capital Proxy refers to the capital proxy variable calculated using the change in local electricity consumption for the service sector instrumented with the SSIVs. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.25: Effect of high-skill concentration in large firms on local growth in places with high and low concentration for different thresholds p

	GDP Growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=0 \times HS\ Conc._{t-1}$	1.091***	0.739**	0.723**	0.630***	0.505**	0.366**	0.300*	0.374*
	(0.328)	(0.259)	(0.231)	(0.187)	(0.160)	(0.130)	(0.148)	(0.156)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=1 \times HS\ Conc._{t-1}$	-0.488***	-0.480***	-0.546**	-0.368	-0.198	-0.104	-0.180	-0.231
	(0.110)	(0.136)	(0.185)	(0.228)	(0.291)	(0.383)	(0.572)	(0.773)
Threshold p	15%	20%	25%	30%	35%	40%	45%	50%
N	74,090	74,090	74,090	74,090	74,090	74,090	74,090	74,090
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Joint F-statistic	30.1	25.0	16.2	13.2	11.4	8.3	4.9	3.7
J-test, p-value	0.05	0.25	0.32	0.39	0.44	0.56	0.37	0.54

High-skill concentration (*HS Conc*) is the local share of high-skilled workers at large firms over total local supply. GDP Growth is real per-capita local GDP growth winsorized at the 1^{st} and 99^{th} percentiles. All columns use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as instruments, each interacted with $\mathbb{1}\{HSConc_{i,t-1} > p\}$. All specifications control for the lagged local sum of shares $s_{i,t-2}$. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

Table A.26: Effect of high-skill concentration in large firms on low-skill wages and labor supply

	Non-High-Skill Wage		Log(# Non-High-Skill)	
	(1)	(2)	(3)	(4)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=0 \times HS\ Conc._{t-1}$	16.54	9.225	3.791*	5.202***
	(677.848)	(690.900)	(1.528)	(1.560)
$\mathbb{1}_{\{HS\ Conc._{t-1}>p\}}=1 \times HS\ Conc._{t-1}$	-8.525	21.18	-0.201	-0.344
	(121.003)	(126.195)	(0.269)	(0.291)
N	68,607	68,607	68,607	68,607
Time FE	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Informality		Yes		Yes
Joint F-statistic	40.3	35.6	40.3	35.6
J-test, p-value	0.70	0.09	0.93	0.13

High-skill concentration (*HS Conc*) is the local share of high-skilled workers at large firms over total local supply. Threshold p is set at the 10^{th} percentile. All columns use $B_{i,t-2,large}$ and $B_{i,t-2,small}$ as instruments, the former interacted with $\mathbb{1}\{HSConc_{i,t-1} > p\}$ and the latter with $\mathbb{1}\{\{HSConc_{i,t-1} > p\} = 0\}$. Lagged local-level controls: log of population, log of average real wage, the population share of high-skilled workers, the population share of workers receiving minimum wage or less, and the ratio of net hiring over population. Informality refers to the 2000 ratio of informal workers over total employment interacted with year fixed-effects. J-test refers to the overidentification test in Hansen (1982). Municipality-clustered standard errors are in parentheses. *, **, *** denote significance at the 5%, 1%, and 0.1% levels, respectively.

References

Acemoglu, D., & Akcigit, U. (2012). Intellectual property rights policy, competition and innovation. *Journal of the European Economic Association*, *10*(1), 1–42.

Adão, R., Kolesár, M., & Morales, E. (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, *134*(4), 1949-2010.

Aghion, P., Harris, C., Howitt, P., & Vickers, J. (2001). Competition, imitation and growth with step-by-step innovation. *The Review of Economic Studies*, *68*(3), 467-492.

Aghion, P., Harris, C., & Vickers, J. (1997). Competition and growth with step-by-step innovation: An example. *European Economic Review*, *41*(3-5), 771–782.

Azar, J., Marinescu, I., & Steinbaum, M. (2022). Labor market concentration. *Journal of Human Resources*, *57*(S), S167-S199.

Betarelli Junior, A. A., Faria, W. R., Gonçalves Montenegro, R. L., Bahia, D. S., & Gonçalves, E. (2020). Research and development, productive structure and economic effects: Assessing the role of public financing in brazil. *Economic Modelling*, *90*, 235-253.

Borusyak, K., Hull, P., & Jaravel, X. (2021). Quasi-experimental shift-share research designs. *The Review of Economic Studies*, *89*(1), 181-213.

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200-230.

Cattaneo, M. D., Crump, R. K., Farrell, M. H., & Feng, Y. (2024). On binscatter. *American Economic Review*, *114*(5), 1488–1514.

Chikis, C. A., Goldberg, J., & López-Salido, D. (2023). A comment on: "low interest rates, market power, and productivity growth". *Econometrica*, *91*(6), 2457-2461.

Damodaran, A. (2023). Equity risk premiums (erp): Determinants, estimation and implications - the 2023 edition. *NYU Stern School of Business*.

de Chaisemartin, C., & D'Haultfœuille, X. (2024). Difference-in-differences estimators of intertemporal treatment effects. *The Review of Economics and Statistics*, 1-45.

Hall, B., & Lerner, J. (2010). The financing of r&d and innovation. In *Handbook of the economics of innovation* (Vol. 1, p. 609-639). Elsevier.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, *50*(4), 1029–1054.

Iacus, S. M., King, G., & Porro, G. (2012). Causal inference without balance checking: Coarsened exact matching. *Political Analysis*, *20*(1), 1–24.

Jaeger, D. A., Ruist, J., & Stuhler, J. (2018). *Shift-share instruments and the impact of immigration* (Working Paper No. 24285). National Bureau of Economic Research.

Jones, C. I. (1995). R & d-based models of economic growth. *Journal of Political Economy*, *103*(4), 759–784.

Liu, E., Mian, A., & Sufi, A. (2022). Low interest rates, market power, and productivity growth. *Econometrica*, *90*(1), 193-221.

Nimier-David, E. (2023). Local human capital and firm creation: Evidence from the massification of higher education in france. *Working Paper*.

Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, *37*(2), 187-204.

Pessoa, A. (2005, 01). "ideas" driven growth: the oecd evidence. *Portuguese Economic Journal*, *4*(1), 46-67.

Pires, M., Marques, P. R., & Bergamin, J. (2023). *A tributação da renda corporativa no brasil: estimativas da carga tributária efetiva a partir das demonstrações de resultado no período 2012-2022* (Research Working Paper). FGV.

Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, *98*(5), S71–S102.

Roth, J., Sant'Anna, P. H., Bilinski, A., & Poe, J. (2023). What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *Journal of Econometrics*, *235*(2), 2218-2244.

Schubert, G., Stansbury, A., & Taska, B. (2024). Employer concentration and outside options [Working Paper]. *Available at SSRN 3599454*.

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, *225*(2), 175-199.

UNDP, IPEA, & FJP. (2024). *Atlas do desenvolvimento humano no brasil.* http://www.atlasbrasil.org.br/. (Accessed June 12, 2024.)

Veloso, F., de Holanda Barbosa Filho, F., & Peruchetti, P. (2022). *Impactos da educação no mercado de trabalho* (Tech. Rep.). IBRE - FGV.

Appendix B

Appendix to Chapter 2: Heterogeneous Local Fiscal Multipliers: New Shift-Share Evidence From The UK

B.1 Figures and Tables

Figure B1: *First-Stage (left) Binned Scatter Plot and Reduced Form (right) Binned Plot*

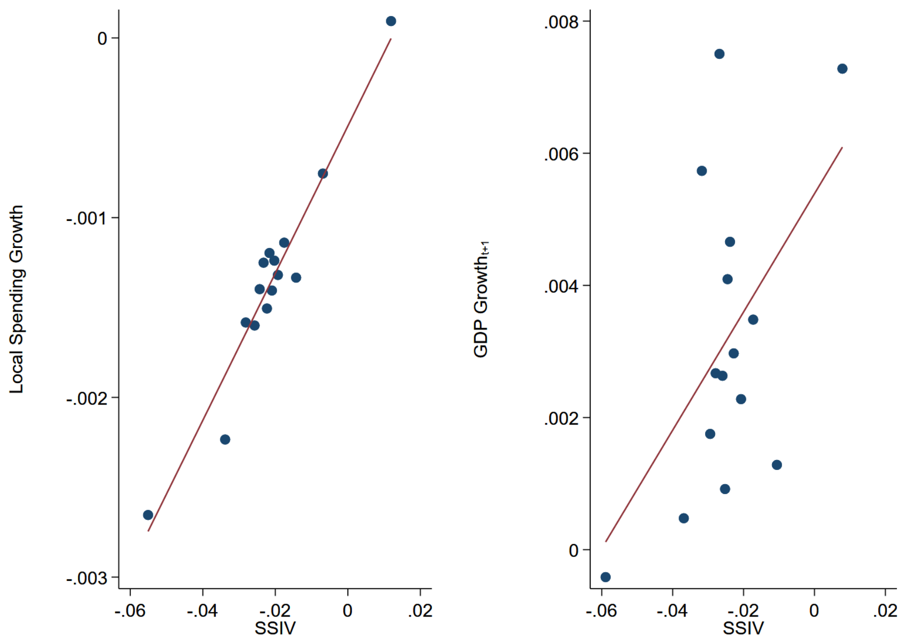

Notes: Variables correspond to the residual of regressing each one on location and year fixed-effects. SSIV was multiplied by a factor of 1000 for clarity.

Figure B2: *Rotemberg Weights: Heterogeneity of $\hat{\beta}_t$.*

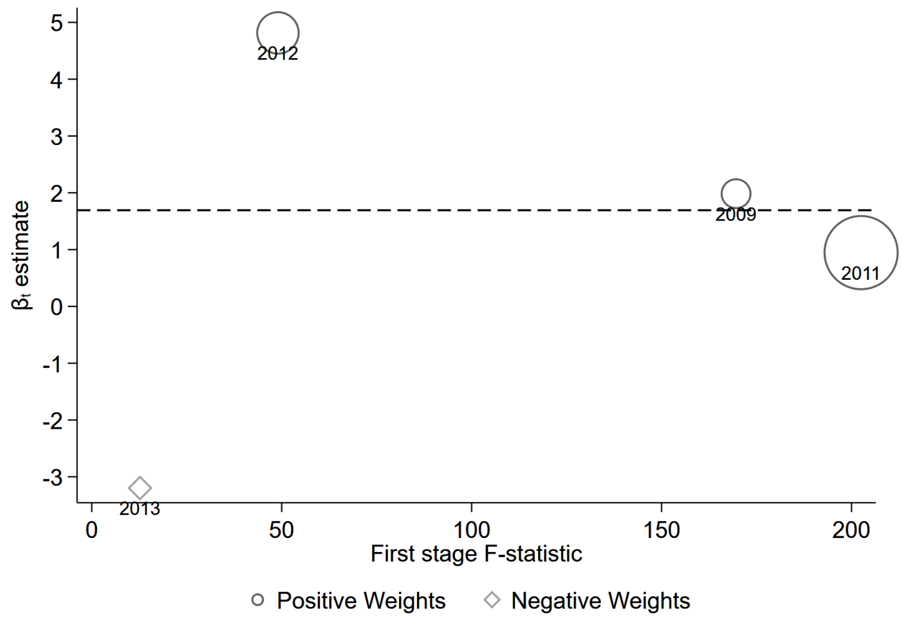

Just-identified coefficient estimates $\hat{\beta}_t$ and first-stage F-statistics are calculated using the specification of column (4) in Table 2.1. The horizontal dashed line indicates the benchmark estimate in Table 2.1 calculated using the SSIV. The size of the points is scaled by the magnitude of the respective Rotemberg weight. The figure excludes instruments with first-stage F-statistic below 10.

Figure B3: *Local Spending Breakdown*

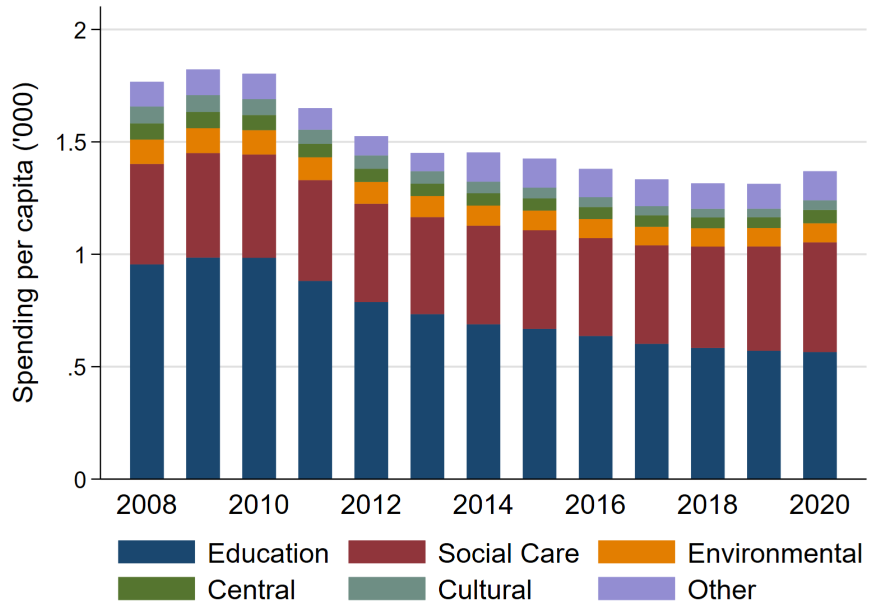

Figure B4: *Average Expected Local Fiscal Multiplier*

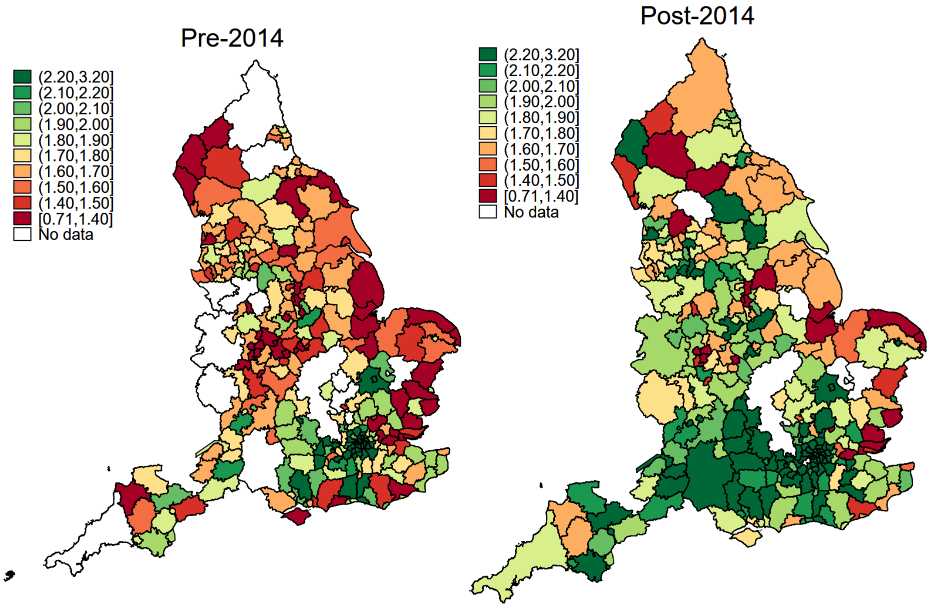

Note: The fiscal multiplier is the average local fiscal multiplier estimated in Table 2.3 using the local-level values for inactivity per capita, share of those in inactivity who want to work, and per-capita low-skilled labor.

Figure B5: *Optimal Spending and Optimal to Actual Spending Ratio (London)*

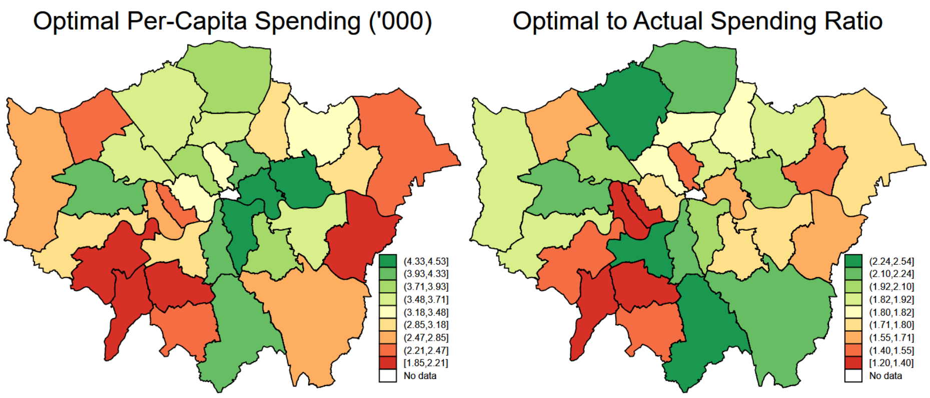

Note: Optimal Per-Capita Spending refers to the planner's solution in the model for optimal local spending. Optimal to Actual Spending Ratio is the ratio between the model-estimated optimal spending and actual spending by LADs. Data and model-generated estimates are for 2010 and in pounds.

Figure B6: *Optimal to Actual Spending Ratio and Grant Allocation*

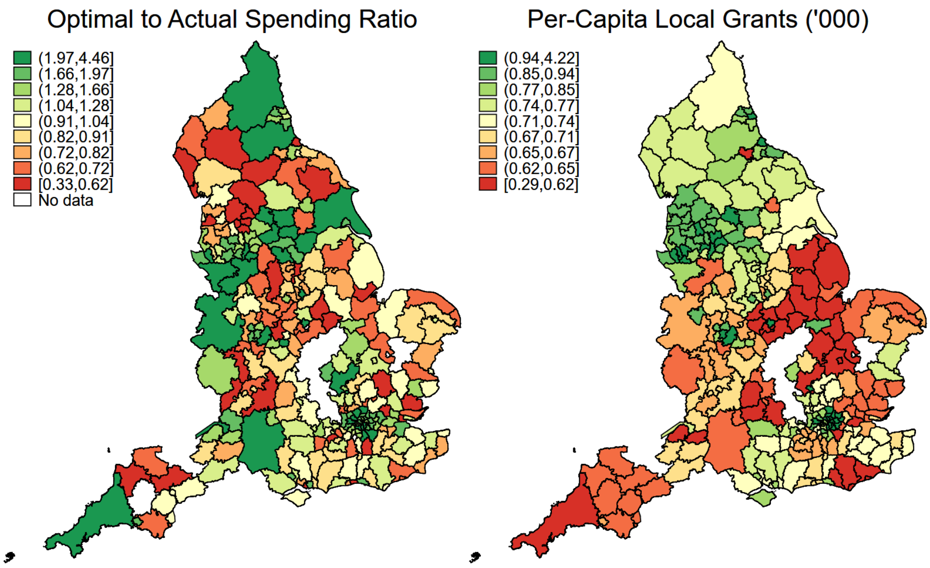

Note: Optimal to Actual Spending Ratio is the ratio between the model-estimated optimal spending and actual spending by LADs. Data and model-generated estimates are for 2010 and in pounds.

Table B.1: Summary Statistics

	Mean	St. Dev.	Obs	Min	Max
GDP Per Capita Growth	0.031	0.181	74,209	-0.840	12.7
Skill Premium	2.114	0.601	68,691	0.293	11.7
Skill Premium - CT Workers	2.086	0.682	68,171	0.355	11.8
High-Skill Concentration	0.624	0.266	74,209	0.000	1.0
CT Worker Concentration	0.608	0.287	73,733	0.000	1.0
High-Skill Workers (th.)	2.153	31.338	74,209	0.010	2319.5
CT Workers (th.)	1.188	15.337	74,209	0.000	1076.3
Non-High-Skill Workers (th.)	9.453	86.051	74,209	0.001	5768.9
Electricity Consumption Growth	-0.013	0.243	74,209	-0.942	7.4
Real Wages (th.)	1.543	0.478	74,209	0.232	9.4
Population (mm.)	0.036	0.209	74,209	0.001	12.0
Total Workers (th.)	11.299	114.655	74,209	0.003	8042.9
Net New Workers Per Capita	0.002	0.017	74,209	-0.867	1.1
Minimum Wage Population Share	0.010	0.013	74,209	0.000	0.6
High-Skill Population Share	0.024	0.024	74,209	0.000	1.8
CT Workers Population Share	0.014	0.015	74,209	0.000	1.6
Informality Share (2000)	0.512	0.165	74,209	0.073	1.0

Table B.2: Local Spending Fiscal Multiplier Estimates (Different Specifications)

	GDP_{t+1}									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Multiplier	1.738**	1.648**	1.635	1.856*	2.247**	1.769	1.789**	1.714**	3.017***	2.878***
	(0.805)	(0.795)	(0.992)	(1.036)	(0.908)	(1.183)	(0.790)	(0.790)	(0.912)	(0.892)
N	3,235	3,235	2,938	2,938	3,235	3,235	3,235	3,235	3,235	3,235
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Council FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes		Yes		Yes		Yes		Yes
Outside AEF Grants	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Robust F-statistic	101.8	110.6	102.1	123.2	94.8	63.2	101.0	109.8	75.5	75.4

Notes: Specifications: (1)-(2): weighted regression; (3)-(4): two-year lagged shares; (5)-(6): first-period fixed shares; (7)-(8): dependent variable is the one-year ahead yearly change in real local GDP per-capita over current real GDP per-capita; (9)-(10): SSIV using shares by spending category. Main regressor corresponds to growth in real local authority total service expenditure per-capita. GDP_{t+1} corresponds to growth one period ahead. Local-level controls (first-period interacted with year fixed effects in columns (5) and (6), one-year lagged for the rest): share of NVQ3 awards (equivalent to a high-school diploma), unemployment rate, median age, child poverty rate, average full-time wage, per-capita number of people receiving central government benefits (for Disability Living Allowance, Incapacity Benefits, Housing Benefit, Universal Credit, Personal Independence Payment, and Child Benefit) where each claimant counts by the number of individual benefits they receive, dummies for the political party controlling the local council, LAD reserves per-capita, per-capita amount of funds from non-domestic rates, and average council tax. Outside AEF Grants is the per-capita amount of all grants outside the AEF which are not included in the SSIV. Standard errors are clustered by counties for non-metropolitan districts and by individual LAD for the rest. *, **, *** denote significance at 10%, 5%, and 1% levels, respectively.

Table B.3: Local Spending Fiscal Multiplier Estimates (TTWA level)

	GDP_{t+1}			
	(1)	(2)	(3)	(4)
Multiplier	1.925*	1.513	1.852*	1.583
	(0.996)	(0.971)	(1.025)	(1.023)
N	1,317	1,317	1,317	1,317
Time FE	Yes	Yes	Yes	Yes
Council FE	Yes	Yes	Yes	Yes
Controls		Yes		Yes
Outside AEF Share			Yes	Yes
Robust F-statistic	29.2	33.4	27.0	30.9

Notes: Main regressor corresponds to growth in real local authority total service expenditure per-capita. GDP_t corresponds to local GDP per-capita growth and GDP_{t+1} to growth one period ahead. Local-level controls (one-year lagged): share of NVQ3 awards (equivalent to a high-school diploma), unemployment rate, median age, child poverty rate, average full-time wage, per-capita number of people receiving central government benefits (for Disability Living Allowance, Incapacity Benefits, Housing Benefit, Universal Credit, Personal Independence Payment, and Child Benefit) where each claimant counts by the number of individual benefits they receive, LAD reserves per-capita, per-capita amount of funds from non-domestic rates, and average council tax. Outside AEF Grants is the per-capita amount of all grants outside the AEF which are not included in the SSIV. Standard errors are clustered by counties for non-metropolitan districts and by individual LAD for the rest. *, **, *** denote significance at 10%, 5%, and 1% levels, respectively.

Table B.4: Local Total Spending Fiscal Multiplier Estimates

	GDP_{t+1}							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Total Spending	1.390**	1.320**	1.427**	1.346**				
	(0.549)	(0.557)	(0.562)	(0.566)				
Total Spending$_{2y}$					1.015***	1.062***	1.066***	1.089***
					(0.337)	(0.341)	(0.338)	(0.343)
N	3,235	3,235	3,235	3,235	2,938	2,938	2,938	2,938
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Council FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes		Yes		Yes		Yes
Outside AEF Grants			Yes	Yes			Yes	Yes
Robust F-statistic	81.2	79.6	83.5	82.7	96.6	97.5	99.1	104.3

Notes: Main regressors correspond to growth in real local authority total expenditure (services and capital combined) per-capita. Subscript $2y$ indicates that the change is over two years. GDP_{t+1} corresponds to local GDP per-capita growth one period ahead. Local-level controls (one-year lagged in columns (1)-(4) and two-year lagged in columns (5)-(8)): share of NVQ3 awards (equivalent to a high-school diploma), unemployment rate, median age, child poverty rate, average full-time wage, per-capita number of people receiving central government benefits (for Disability Living Allowance, Incapacity Benefits, Housing Benefit, Universal Credit, Personal Independence Payment, and Child Benefit) where each claimant counts by the number of individual benefits they receive, dummies for the political party controlling the local council, LAD reserves per-capita, per-capita amount of funds from non-domestic rates, and average council tax. Outside AEF Grants is the per-capita amount of all grants outside the AEF which are not included in the SSIV. Standard errors are clustered by counties for non-metropolitan districts and by individual LAD for the rest. *, **, *** denote significance at 10%, 5%, and 1% levels, respectively.

Table B.5: Shift-Share Balance Test

	DSG Share	GDP Growth	Within AEF Share	GDP Growth
	(1)	(2)	(3)	(4)
High School level	0.003	-0.076*	0.018	-0.063*
	(0.015)	(0.039)	(0.013)	(0.037)
Unemployment Rate	0.037***	-0.040	0.032***	-0.038
	(0.012)	(0.024)	(0.010)	(0.023)
Age (median)	-0.115	-0.281*	-0.075	-0.358**
	(0.113)	(0.165)	(0.101)	(0.162)
Child Poverty Rate	-0.117**	0.003	-0.123**	-0.010
	(0.058)	(0.055)	(0.050)	(0.050)
Independent	0.239	-0.015	0.357*	-0.052
	(0.186)	(0.121)	(0.198)	(0.107)
Liberal Democrat	-0.076	-0.202*	-0.128*	-0.232**
	(0.077)	(0.104)	(0.069)	(0.095)
Labour	0.090	-0.149*	-0.005	-0.124*
	(0.059)	(0.088)	(0.059)	(0.073)
No Control	-0.030	-0.074	-0.027	-0.073
	(0.031)	(0.059)	(0.035)	(0.049)
Transfers (People)	-0.006	0.207***	0.055	0.241***
	(0.063)	(0.058)	(0.042)	(0.066)
Wage	-0.121***	0.061	-0.112***	0.038
	(0.043)	(0.106)	(0.037)	(0.097)
Reserves	0.059*	-0.047	0.050	-0.038
	(0.031)	(0.030)	(0.031)	(0.028)
Non-Domestic Rates	0.037	0.050	0.070	0.037
	(0.050)	(0.057)	(0.047)	(0.052)
Council Tax	-0.450	0.859*	-0.228	0.693*
	(0.275)	(0.453)	(0.249)	(0.416)
N	2,949	2,949	3,235	3,235
R-squared	0.937	0.431	0.943	0.420
Time FE	Yes	Yes	Yes	Yes
Council FE	Yes	Yes	Yes	Yes

DSG Share is the ratio between the Dedicated Schools Grant and local authority spending. Within AEF Share is the ratio between the sum of all grants inside the AEF and local authority spending. Regressors are one-year lagged in columns (2) and (4). Standard errors are clustered by LAD. *, **, *** denote significance at 10%, 5%, and 1% levels, respectively.

Table B.6: Local Spending Fiscal Multiplier Estimates (Overidentified Shares-Only Specification)

	GDP_{t+1}			
	(1)	(2)	(3)	(4)
Multiplier	0.906	0.495	0.907	0.519
	(0.658)	(0.633)	(0.661)	(0.639)
N	3,235	3,235	3,235	3,235
Time FE	Yes	Yes	Yes	Yes
Council FE	Yes	Yes	Yes	Yes
Controls		Yes		Yes
Outside AEF Grants			Yes	Yes
Robust F-test	21.8	21.9	22.1	22.9
J-test, p-val	0.03	0.02	0.03	0.02

Notes: Estimates are calculated using the Limited Information Maximum Likelihood (LIML) estimator. Main regressor corresponds to real local authority total service expenditure per-capita growth. GDP_{t+1} corresponds to growth one period ahead. Local-level controls (one-year lagged): share of NVQ3 awards (equivalent to a high-school diploma), unemployment rate, median age, child poverty rate, average full-time wage, per-capita number of people receiving central government benefits (for Disability Living Allowance, Incapacity Benefits, Housing Benefit, Universal Credit, Personal Independence Payment, and Child Benefit) where each claimant counts by the number of individual benefits they receive, dummies for the political party controlling the local council, LAD reserves per-capita, per-capita amount of funds from non-domestic rates, and average council tax. Outside AEF Grants is the per-capita amount of all grants outside the AEF which are not included in the SSIV. Standard errors are clustered by counties for non-metropolitan districts and by individual LAD for the rest. *, **, *** denote significance at 10%, 5%, and 1% levels, respectively.

Table B.7: Summary of Rotemberg Weights

Panel A: Negative and positive weights

	Sum	Mean	Share
Negative	-0.040	-0.020	0.037
Positive	1.040	0.116	0.963

Panel B: Correlations

	$\hat{\alpha}_t$	g_t	$\hat{\beta}_t$	$\hat{F}_t$	Var(s_{t-1})
$\hat{\alpha}_t$	1				
g_t	-0.832	1			
$\hat{\beta}_t$	0.120	-0.010	1		
$\hat{F}_t$	0.816	-0.417	0.163	1	
Var(s_{t-1})	0.561	-0.282	0.027	0.628	1

Panel C: Top 4 Rotemberg weight years

	$\hat{\alpha}_t$	Share of LADs with negative grant grw.	g_t	$\hat{\beta}_t$	95 % CI
2011	0.670	0.969	-3.506	0.947	(-1.40,3.10)
2012	0.213	0.896	-2.498	4.814	(1.00,10.80)
2009	0.102	0.021	0.625	1.984	(0.10,4.20)
2013	-0.030	0.587	-0.716	-3.197	(-14.00,5.30)

Panel D: Estimates of $\hat{\beta}_t$ for positive and negative weights

	$\hat{\alpha}$-weighted sum	Share of overall $\hat{\beta}_t$	Mean
Negative	-0.088	-0.052	7.203
Positive	1.780	1.052	-9.511

Note: This table reports the summary statistics about the Rotemberg weights using the specification of column (8) in Table 2.1. Panel A reports the sum, mean, and share of weights for both positive and negative weights. Panel B reports the correlations between the weights ($\hat{\alpha}_t$), the national changes to grants (g_t), the just-identified coefficients estimates calculated for each year ($\hat{\beta}_t$), the first-stage F-statistic of the year shares ($\hat{F}_t$), and the variation in the year shares across LADs (Var(s_{t-1})). Panel C reports the top four years according to the Rotemberg weights (annual shock g_t was multiplied by a factor of 1000 for clarity). The 95% confidence interval is the weak instrument robust confidence interval using the method from Chernozhukov & Hansen (2008) over a range of -20 to 20. Panel D reports how the values of $\hat{\beta}_t$ vary with the positive and negative Rotemberg weights.

Table B.8: Heterogeneity in the Effect of Local Fiscal Spending by Spending Category

	GDP_{t+1}		
	(1)	(2)	(3)
Education Spending Share	-0.763*		
	(0.414)		
Social Care Spending Share		1.116*	
		(0.597)	
Planning Spending Share			3.344*
			(1.850)
N	3,235	3,235	3,235
Time FE	Yes	Yes	Yes
Council FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Outside AEF Grants	Yes	Yes	Yes
Spending Category	Yes	Yes	Yes
Robust F-test	13.1	14.4	9.7

Notes: Main regressors correspond to change in the spending-specific share of total local spending in services. GDP_t corresponds to local GDP per-capita growth and GDP_{t+1} to growth one period ahead. Local-level controls (one-year lagged): share of NVQ3 awards (equivalent to a high-school diploma), unemployment rate, median age, child poverty rate, average full-time wage, per-capita number of people receiving central government benefits (for Disability Living Allowance, Incapacity Benefits, Housing Benefit, Universal Credit, Personal Independence Payment, and Child Benefit) where each claimant counts by the number of individual benefits they receive, dummies for the political party controlling the local council, LAD reserves per-capita, per-capita amount of funds from non-domestic rates, and average council tax. Outside AEF Grants is the per-capita amount of all grants outside the AEF which are not included in the SSIV. Spending Category controls for the one-year lagged share of spending in transportation, education, social-care, housing, cultural, planning, central, and environmental. Standard errors are clustered by counties for non-metropolitan districts and by individual LAD for the rest. *, **, *** denote significance at 10%, 5%, and 1% levels, respectively.

References

Chernozhukov, V., & Hansen, C. (2008). Instrumental variable quantile regression: A robust inference approach. *Journal of Econometrics*, *142*(1), 379-398.

Appendix C

Appendix to Chapter 3: The Fiscal Multiplier of Education Expenditures

C.1 Main Legislative Changes to the Pell Grant Program

First, in 1978, the Middle Income Student Assistance Act (MISAA) expanded student eligibility by limiting the rate at which parental discretionary income was assessed under the EFC formula. This act was repealed two years later, in 1980. In 1990, the Omnibus Budget Reconciliation Act eliminated student aid eligibility at high default schools. In 1992, the Higher Education Act was reauthorized and changed the definition of an independent student. In 1994, the Violent Crime Control and Law Enforcement Act eliminated Pell grants for prisoners. In 2007 Congress passed the College Cost Reduction and Access Act (CCRAA), which supplemented the grant funding and changed Pell eligibility by increasing the amount and types of income excluded from the EFC formula. A renewed set of legislative measures paired with the countercyclicality of the enrollment effect caused a significant increase in Pell grant disbursements. These legislative measures include: the Higher Education Opportunity Act (HEOA) of 2008, which authorized year-round Pell grants and limited eligibility to 18 full-time semesters or the equivalent; the American Recovery and Reinvestment Act (ARRA) of 2009, which provided additional funding to the Pell Grant Program (ARRA raised the maximum Pell grant by more than $400); the Health Care and Education Reconciliation Act of 2010, which increased the maximum Pell grant by over $600 and expanded eligibility by increasing the income threshold (from $20,000 to $30,000) for an automatic EFC of zero. Pell grant disbursements started to decline in 2011, once the economy gained momentum and undergraduate enrollment returned to pre-crisis levels.[148] Congress eliminated the year-round Pell grant eligibility established in 2008, when it provided supplemental funding to the program

[148] During economic recovery, fewer individuals qualify to receive Pell grants. Enrollment is counter-cyclical as people opt for employment instead of education.

and lowered the income threshold for an automatic EFC of zero to $23,000. In 2012, the Consolidated Appropriations Act provided additional funding to the Pell Grant Program and reduced Pell lifetime eligibility to 12 semesters.

C.2 Additional Tables

Table C.1: Effect of Pell Grants on Local Income Per Capita for 1-Year Horizon

	Full Sample				Post 1999						Full Sample
	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS	(9) 2SLS	(10) 2SLS	(11) OLS
Panel A: Income Growth											
Multiplier	2.232*	2.221*	2.371*	2.358*	2.930*	3.011**	2.874*	2.775*	2.858*	2.836*	-0.785
	(1.233)	(1.254)	(1.234)	(1.254)	(1.499)	(1.524)	(1.528)	(1.519)	(1.539)	(1.505)	(0.728)
Panel B: Employment Growth											
Multiplier	0.888	0.984	1.022	1.107	1.996*	2.151**	2.135**	1.793	1.939*	1.809	-0.992
	(0.977)	(0.956)	(0.973)	(0.957)	(1.108)	(1.065)	(1.077)	(1.118)	(1.073)	(1.099)	(0.617)
Observations	8,793	8,793	8,793	8,793	4,781	4,781	4,781	4,781	4,781	4,781	8,793
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test	96.4	94.1	96.4	94.2	65.3	64.2	63.6	64.2	63.3	68.2	-
Joint F-test	-	54.0	-	53.9	-	36.4	36.3	-	36.0	35.3	-

Notes: SSIV strategy for the Pell grants regressor uses the one-year lagged population share of recipients. SSIV strategy for appropriations uses the one-year lagged appropriation share of income. Controls are one-year lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test is the robust F-statistic of the first-stage regression of Pell grants. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.2: Effect of Pell Grants on Local Income Per Capita Weighting by Two-Year Lagged MSA Population Logarithm

	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
Panel A: Income Growth											
Multiplier	2.898**	2.816*	3.056**	2.940**	3.723**	3.708**	3.175*	3.085*	3.080*	3.219**	-1.662*
	(1.434)	(1.471)	(1.443)	(1.479)	(1.690)	(1.719)	(1.667)	(1.637)	(1.662)	(1.627)	(0.923)
Panel B: Employment Growth											
Multiplier	1.711	1.895*	1.842*	1.978*	3.191**	3.302**	2.935**	2.727**	2.850**	2.749**	-1.577**
	(1.091)	(1.103)	(1.095)	(1.115)	(1.292)	(1.304)	(1.256)	(1.233)	(1.239)	(1.231)	(0.761)
Observations	8,436	8,436	8,436	8,436	4,447	4,447	4,447	4,447	4,447	4,447	8,436
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test	103.4	95.8	103.4	95.9	74.8	71.3	70.8	74.0	70.7	77.3	-
Joint F-test	-	70.1	-	70.4	-	38.7	38.4	-	38.3	45.1	-

Notes: SSIV strategy for the Pell grants regressor uses the two-year lagged population share of recipients. SSIV strategy for appropriations uses the two-year lagged appropriation share of income. Controls are one-year lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test is the robust F-statistic of the first-stage regression of Pell grants. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.3: Balance Test for the Specification with Fiscal Transfers

	Full Sample				Post 1999			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Shares	Income Growth	Empl. Growth	Pass	Shares	Income Growth	Empl. Growth	Pass
Spending Growth	0.006	-0.013	-0.038**	✓	-0.005	-0.022	-0.048**	✓
	(0.014)	(0.023)	(0.017)		(0.008)	(0.033)	(0.019)	
Log(Tuition)	-0.006	-0.068	-0.019	✓	-0.022	-0.237***	-0.118*	✓
	(0.048)	(0.051)	(0.050)		(0.049)	(0.063)	(0.064)	
D.Log(Students)	0.017	0.001	-0.007	✓	0.021*	0.001	-0.009	✓
	(0.010)	(0.009)	(0.009)		(0.011)	(0.011)	(0.007)	
For Profit	0.052	-0.002	0.026	✓	0.006	0.027	0.092*	✓
	(0.035)	(0.040)	(0.028)		(0.028)	(0.057)	(0.047)	
Share Black	0.157	-0.026	0.246	✓	0.593	0.046	0.427	✓
	(0.233)	(0.159)	(0.158)		(0.388)	(0.524)	(0.360)	
Share Hisp.	-0.172	0.558***	0.780***	✓	-0.287	0.353	0.794***	✓
	(0.142)	(0.144)	(0.117)		(0.223)	(0.363)	(0.278)	
Share Bach.	-0.115	0.074	0.271***	✓	0.026	0.654***	0.771***	✓
	(0.133)	(0.102)	(0.081)		(0.134)	(0.166)	(0.147)	
Risk Score					-0.041	-0.022	-0.090*	✓
					(0.029)	(0.042)	(0.048)	
Age					0.002	-0.016	0.001	✓
					(0.037)	(0.041)	(0.038)	
Debt to Income					-0.052	0.238	-0.082*	✓
					(0.032)	(0.191)	(0.044)	
Card Util.					0.026*	-0.037	-0.014	✓
					(0.016)	(0.030)	(0.024)	
Mort. Delinq.					0.013	-0.009	-0.022	✓
					(0.015)	(0.019)	(0.019)	
Observations	8,436	8,436	8,436		4,447	4,447	4,447	
R-square	0.7	0.3	0.5		0.8	0.4	0.6	
Time FE	Yes	Yes	Yes		Yes	Yes	Yes	
MSA FE	Yes	Yes	Yes		Yes	Yes	Yes	

Notes: Independent variables are twice lagged in columns (2), (3), (6), and (7) except for spending growth. Spending Growth refers to the change in the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.4: Balance Test for Appropriations

	Full Sample				Post 1999			
	(1) Shares	(2) Income Growth	(3) Empl. Growth	(4) Pass	(5) Shares	(6) Income Growth	(7) Empl. Growth	(8) Pass
Log(Tuition)	-0.015	-0.069	-0.021	✓	-0.037	-0.239***	-0.123*	✓
	(0.028)	(0.051)	(0.050)		(0.034)	(0.063)	(0.064)	
D.Log(Students)	-0.000	0.001	-0.006	✓	-0.003	0.001	-0.009	✓
	(0.003)	(0.009)	(0.009)		(0.006)	(0.011)	(0.007)	
For Profit	0.026	-0.002	0.026	✓	0.030**	0.029	0.097**	×
	(0.018)	(0.040)	(0.028)		(0.014)	(0.057)	(0.048)	
Share Black	0.183	-0.027	0.244	✓	-0.147	0.044	0.423	✓
	(0.125)	(0.159)	(0.158)		(0.156)	(0.523)	(0.359)	
Share Hisp.	0.347***	0.557***	0.777***	×	0.470***	0.349	0.787***	×
	(0.076)	(0.143)	(0.117)		(0.148)	(0.362)	(0.279)	
Share Bach.	0.271***	0.075	0.274***	×	0.342***	0.662***	0.788***	×
	(0.067)	(0.101)	(0.081)		(0.082)	(0.164)	(0.146)	
Risk Score					0.007	-0.024	-0.093*	✓
					(0.017)	(0.043)	(0.048)	
Age					0.015	-0.014	0.005	✓
					(0.020)	(0.041)	(0.038)	
Debt to Income					0.005	0.237	-0.084*	✓
					(0.017)	(0.191)	(0.045)	
Card Util.					-0.014	-0.038	-0.015	✓
					(0.010)	(0.030)	(0.024)	
Mort. Delinq.					0.003	-0.009	-0.022	✓
					(0.005)	(0.019)	(0.019)	
Observations	8,436	8,436	8,436		4,447	4,447	4,447	
R-square	1.0	0.3	0.5		1.0	0.4	0.6	
Time FE	Yes	Yes	Yes		Yes	Yes	Yes	
MSA FE	Yes	Yes	Yes		Yes	Yes	Yes	

Notes: Independent variables are twice lagged in columns (2), (3), (6), and (7). MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.5: Balance Test for Fiscal Transfers

	Post 1999				Full Cases Only			
	(1) Shares	(2) Income Growth	(3) Empl. Growth	(4) Pass	(5) Shares	(6) Income Growth	(7) Empl. Growth	(8) Pass
Log(Tuition)	-0.019	-0.239***	-0.123*	✓	-0.003	-0.121	-0.066	✓
	(0.035)	(0.063)	(0.064)		(0.041)	(0.125)	(0.109)	
D.Log(Students)	0.000	0.001	-0.009	✓	0.009	-0.009	-0.022	✓
	(0.006)	(0.011)	(0.007)		(0.013)	(0.028)	(0.014)	
For Profit	0.038**	0.029	0.097**	×	-0.000	0.030	0.119*	✓
	(0.015)	(0.057)	(0.048)		(0.020)	(0.079)	(0.067)	
Share Black	-0.020	0.044	0.423	✓	-0.209	0.971	1.050	✓
	(0.182)	(0.523)	(0.359)		(0.256)	(1.384)	(1.004)	
Share Hisp.	0.541***	0.349	0.787***	×	0.739***	0.245	0.834	✓
	(0.166)	(0.362)	(0.279)		(0.259)	(0.714)	(0.643)	
Share Bach.	0.407***	0.662***	0.788***	×	0.349***	0.578	0.698**	×
	(0.083)	(0.164)	(0.146)		(0.112)	(0.389)	(0.281)	
Risk Score	0.013	-0.024	-0.093*	✓	0.014	-0.115**	-0.107*	✓
	(0.019)	(0.043)	(0.048)		(0.022)	(0.053)	(0.061)	
Age	0.041*	-0.014	0.005	✓	0.054*	-0.040	0.050	✓
	(0.022)	(0.041)	(0.038)		(0.028)	(0.051)	(0.046)	
Debt to Income	0.017	0.237	-0.084*	✓	0.041**	0.527	-0.143*	×
	(0.016)	(0.191)	(0.045)		(0.018)	(0.338)	(0.083)	
Card Util.	-0.012	-0.038	-0.015	✓	-0.009	-0.067	0.019	✓
	(0.011)	(0.030)	(0.024)		(0.009)	(0.041)	(0.028)	
Mort. Delinq.	0.019***	-0.009	-0.022	✓	0.022***	0.024	-0.000	✓
	(0.006)	(0.019)	(0.019)		(0.007)	(0.028)	(0.027)	
Observations	4,447	4,447	4,447		2,523	2,523	2,523	
R-square	1.0	0.4	0.6		1.0	0.5	0.7	
Time FE	Yes	Yes	Yes		Yes	Yes	Yes	
MSA FE	Yes	Yes	Yes		Yes	Yes	Yes	

Notes: Independent variables are twice lagged in columns (2), (3), (6), and (7). Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.6: Effect of Pell Grants on Local Income Per Capita and Employment (Three-Times Lagged Shares)

	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
Panel A: Income Growth											
Multiplier	3.411**	3.406**	3.407**	3.338**	4.215**	4.218**	3.230**	3.269**	3.249**	3.470**	-1.732*
	(1.452)	(1.506)	(1.436)	(1.494)	(1.674)	(1.714)	(1.601)	(1.610)	(1.642)	(1.592)	(0.938)
Panel B: Employment Growth											
Multiplier	1.830	2.164*	1.871	2.133*	3.259**	3.297**	2.463**	2.389*	2.452**	2.347*	-1.766**
	(1.157)	(1.148)	(1.144)	(1.150)	(1.332)	(1.354)	(1.250)	(1.221)	(1.230)	(1.240)	(0.727)
Observations	8,062	8,062	8,062	8,062	4,118	4,118	4,118	4,118	4,118	4,118	8,062
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes	Yes	Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test	96.4	89.3	96.8	89.9	69.6	65.8	66.4	71.0	66.8	73.0	-
Joint F-test	-	59.0	-	58.6	-	35.8	35.7	-	35.9	62.0	-

Notes: SSIV strategy for the Pell grants regressor uses the thrice-lagged share of recipients in MSA population (see eq. 3.3). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test is the robust F-statistic of the first-stage regression of Pell grants. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.7: Effect of Pell Grants on Local Income Per Capita Using Different Estimators (Additional Specifications)

Post 1999	Income Growth			Employment Growth		
	(1)	(2)	(3)	(4)	(5)	(6)
2SLS (SSIV)	3.640**	3.077*	3.125**	3.120**	2.722**	2.657**
	(1.691)	(1.641)	(1.638)	(1.312)	(1.268)	(1.257)
2SLS	4.418**	3.811**	3.896**	2.707**	2.465**	2.223*
	(1.726)	(1.699)	(1.664)	(1.236)	(1.220)	(1.184)
LIML	4.540**	3.892**	4.022**	2.941**	2.631**	2.393*
	(1.777)	(1.744)	(1.714)	(1.315)	(1.286)	(1.238)
HFUL	4.882***	4.146***	4.328***	3.089***	2.758***	2.469**
	(1.467)	(1.480)	(1.483)	(1.007)	(0.994)	(0.970)
Observations	4,447	4,447	4,447	4,447	4,447	4,447
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes
Approp.						
MSA Controls		Yes	Yes		Yes	Yes
Financial Controls			Yes			Yes
Fiscal Transfers			Yes			Yes

Notes: 2SLS uses each yearly share as a separate IV. LIML uses the limited information maximum likelihood estimation with the same set of instruments. Finally, HFUL uses the estimator from Hausman et al. (2012) also with the same set of instruments. Controls are contemporaneous to the respective timing of shares. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.8: Pell Grants and Student Loans (Additional Specifications)

Post 1999 Sample	Student Loan Growth			Income Growth			Employment Growth		
	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS	(9) 2SLS
Multiplier	2.019***	2.006***	1.995***						
	(0.450)	(0.447)	(0.448)						
Δ (Pell Grants + Loans)				1.206**	0.992*	1.044**	1.033***	0.876**	0.888**
				(0.531)	(0.533)	(0.530)	(0.392)	(0.389)	(0.388)
Observations	4,447	4,447	4,447	4,447	4,447	4,447	4,447	4,447	4,447
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.									
MSA Controls		Yes	Yes		Yes	Yes		Yes	Yes
Financial Controls		Yes	Yes		Yes	Yes		Yes	Yes
Fiscal Transfers			Yes			Yes			Yes
Δ Pell Grants F-test	78.0	77.2	80.9	80.3	80.1	79.5	80.3	80.1	79.5
Joint F-test	-	-	45.7		-	36.3		-	36.3

Notes: SSIV strategy for the Pell grants regressor and the sum of Pell grants and loan disbursements uses the twice-lagged share of recipients in MSA population (see eq. 3.3). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test is the robust F-statistic of the first-stage regression of Pell grants. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.9: Effect of Pell Grants on Education Expenditures by For-Profit and Non-Profit Colleges

Education Exp. Growth	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
Non-Profit Δ Pell Grants	0.466	0.729	0.473	0.733	0.608	0.824	0.845	0.648	0.862	0.728	0.855***
	(0.589)	(0.570)	(0.588)	(0.568)	(0.665)	(0.574)	(0.576)	(0.662)	(0.570)	(0.622)	(0.226)
For-Profit Δ Pell Grants	1.500***	1.469***	1.490***	1.460***	1.457***	1.436***	1.432***	1.434***	1.417***	1.442***	1.397***
	(0.160)	(0.158)	(0.161)	(0.159)	(0.187)	(0.180)	(0.180)	(0.189)	(0.182)	(0.193)	(0.190)
Difference	-1.034*	-0.740	-1.017*	-0.727	-0.849	-0.613	-0.588	-0.786	-0.556	-0.714	-0.542*
Std. Error	(0.530)	(0.529)	(0.530)	(0.527)	(0.576)	(0.510)	(0.510)	(0.575)	(0.507)	(0.546)	(0.285)
Observations	8,432	8,432	8,432	8,432	4,443	4,443	4,443	4,443	4,443	4,443	8,432
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, NP	107.3	97.1	107.1	97.1	85.0	79.9	80.2	85.5	80.3	92.3	-
Δ Pell Grants F-test, FP	28.8	29.8	29.0	29.8	33.4	33.3	33.6	33.0	33.0	35.0	-
Joint F-test	53.7	39.9	53.7	40.3	42.4	20.6	20.6	42.7	20.5	27.3	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at for- and non-profit institutions (see eq. 3.8). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, NP is the robust F-statistic of the first-stage regression of Pell grants at non-profit colleges. Δ Pell Grants F-test, FP is the robust F-statistic of the first-stage regression of Pell grants at for-profit colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.10: Effect of Pell Grants on Non-Education Expenditures by For-Profit and Non-Profit Colleges

Non-Education Exp. Growth	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
Non-Profit Δ Pell Grants	0.296*	0.396**	0.304*	0.402**	0.507**	0.395*	0.382*	0.510***	0.399*	0.458**	0.121
	(0.160)	(0.195)	(0.161)	(0.195)	(0.197)	(0.228)	(0.228)	(0.195)	(0.223)	(0.197)	(0.108)
For-Profit Δ Pell Grants	0.007	-0.005	0.008	-0.003	-0.006	0.005	-0.002	-0.033	-0.025	-0.038	-0.019
	(0.086)	(0.085)	(0.086)	(0.085)	(0.086)	(0.088)	(0.088)	(0.083)	(0.083)	(0.085)	(0.051)
Difference	0.289*	0.401**	0.296*	0.405**	0.513***	0.390*	0.383*	0.543***	0.424*	0.496***	0.140
Std. Error	(0.166)	(0.203)	(0.167)	(0.204)	(0.189)	(0.222)	(0.222)	(0.190)	(0.216)	(0.190)	(0.125)
Observations	8,432	8,432	8,432	8,432	4,443	4,443	4,443	4,443	4,443	4,443	8,432
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, NP	107.3	97.1	107.1	97.1	85.0	79.9	80.2	85.5	80.3	92.3	-
Δ Pell Grants F-test, FP	28.8	29.8	29.0	29.8	33.4	33.3	33.6	33.0	33.0	35.0	-
Joint F-test	53.7	39.9	53.7	40.3	42.4	20.6	20.6	42.7	20.5	27.3	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at for- and non-profit institutions (see eq. 3.8). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, NP is the robust F-statistic of the first-stage regression of Pell grants at non-profit colleges. Δ Pell Grants F-test, FP is the robust F-statistic of the first-stage regression of Pell grants at for-profit colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.11: Effect of Pell Grants on Research Expenditures by For-Profit and Non-Profit Colleges

Research Exp. Growth	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
Non-Profit Δ Pell Grants	0.167	0.231*	0.173	0.236*	0.340**	0.257*	0.249	0.338**	0.255*	0.300**	0.047
	(0.115)	(0.125)	(0.117)	(0.126)	(0.155)	(0.151)	(0.154)	(0.157)	(0.152)	(0.149)	(0.062)
For-Profit Δ Pell Grants	0.016	0.008	0.017	0.010	0.004	0.012	0.006	-0.006	0.000	-0.010	-0.036
	(0.071)	(0.070)	(0.071)	(0.070)	(0.071)	(0.071)	(0.070)	(0.071)	(0.069)	(0.071)	(0.029)
Difference	0.151	0.222	0.156	0.226	0.336**	0.244	0.243	0.344**	0.255	0.310**	0.083
Std. Error	(0.138)	(0.143)	(0.139)	(0.144)	(0.161)	(0.156)	(0.158)	(0.164)	(0.158)	(0.156)	(0.075)
Observations	8,432	8,432	8,432	8,432	4,443	4,443	4,443	4,443	4,443	4,443	8,432
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, NP	107.3	97.1	107.1	97.1	85.0	79.9	80.2	85.5	80.3	92.3	-
Δ Pell Grants F-test, FP	28.8	29.8	29.0	29.8	33.4	33.3	33.6	33.0	33.0	35.0	-
Joint F-test	53.7	39.9	53.7	40.3	42.4	20.6	20.6	42.7	20.5	27.3	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at for- and non-profit institutions (see eq. 3.8). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, NP is the robust F-statistic of the first-stage regression of Pell grants at non-profit colleges. Δ Pell Grants F-test, FP is the robust F-statistic of the first-stage regression of Pell grants at for-profit colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.12: Effect of Pell Grants on Instruction Expenditures by For-Profit and Non-Profit Colleges

Instruction Exp. Growth	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
Non-Profit Δ Pell Grants	0.392	0.635	0.397	0.639	0.480	0.680	0.699	0.511	0.710	0.585	0.764***
	(0.538)	(0.520)	(0.537)	(0.518)	(0.611)	(0.526)	(0.528)	(0.608)	(0.522)	(0.571)	(0.217)
For-Profit Δ Pell Grants	0.493***	0.464***	0.485***	0.457***	0.452**	0.433***	0.430***	0.433**	0.418***	0.440***	0.477***
	(0.188)	(0.167)	(0.186)	(0.167)	(0.177)	(0.161)	(0.162)	(0.173)	(0.157)	(0.162)	(0.097)
Difference	-0.102	0.171	-0.088	0.182	0.027	0.247	0.269	0.077	0.292	0.144	0.287
Std. Error	(0.497)	(0.490)	(0.496)	(0.488)	(0.531)	(0.466)	(0.467)	(0.529)	(0.462)	(0.497)	(0.230)
Observations	8,432	8,432	8,432	8,432	4,443	4,443	4,443	4,443	4,443	4,443	8,432
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, NP	107.3	97.1	107.1	97.1	85.0	79.9	80.2	85.5	80.3	92.3	-
Δ Pell Grants F-test, FP	28.8	29.8	29.0	29.8	33.4	33.3	33.6	33.0	33.0	35.0	-
Joint F-test	53.7	39.9	53.7	40.3	42.4	20.6	20.6	42.7	20.5	27.3	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at for- and non-profit institutions (see eq. 3.8). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, NP is the robust F-statistic of the first-stage regression of Pell grants at non-profit colleges. Δ Pell Grants F-test, FP is the robust F-statistic of the first-stage regression of Pell grants at for-profit colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.13: Effect of Pell Grants on Public Service Expenditures by For-Profit and Non-Profit Colleges

Public Service Exp. Growth	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
Non-Profit Δ Pell Grants	0.130	0.165	0.131	0.166	0.167	0.138	0.133	0.172	0.144	0.158	0.074
	(0.131)	(0.157)	(0.131)	(0.157)	(0.131)	(0.158)	(0.158)	(0.130)	(0.153)	(0.134)	(0.084)
For-Profit Δ Pell Grants	-0.009	-0.013	-0.009	-0.013	-0.010	-0.008	-0.008	-0.027	-0.025	-0.028	0.017
	(0.031)	(0.030)	(0.031)	(0.030)	(0.033)	(0.035)	(0.036)	(0.031)	(0.031)	(0.033)	(0.035)
Difference	0.139	0.178	0.140	0.179	0.177	0.146	0.141	0.199*	0.169	0.186	0.057
Std. Error	(0.120)	(0.151)	(0.120)	(0.150)	(0.114)	(0.146)	(0.146)	(0.118)	(0.138)	(0.118)	(0.092)
Observations	8,432	8,432	8,432	8,432	4,443	4,443	4,443	4,443	4,443	4,443	8,432
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, NP	107.3	97.1	107.1	97.1	85.0	79.9	80.2	85.5	80.3	92.3	-
Δ Pell Grants F-test, FP	28.8	29.8	29.0	29.8	33.4	33.3	33.6	33.0	33.0	35.0	-
Joint F-test	53.7	39.9	53.7	40.3	42.4	20.6	20.6	42.7	20.5	27.3	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at for- and non-profit institutions (see eq. 3.8). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, NP is the robust F-statistic of the first-stage regression of Pell grants at non-profit colleges. Δ Pell Grants F-test, FP is the robust F-statistic of the first-stage regression of Pell grants at for-profit colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.14: Effect of Pell Grants on Student Service Expenditures by For-Profit and Non-Profit Colleges

Student Serv. Exp. Grw.	Full Sample				Post 1999						Full Sample
	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS	(9) 2SLS	(10) 2SLS	(11) OLS
Non-Profit Δ Pell Grants	0.075	0.094	0.075	0.095	0.128*	0.143**	0.146**	0.137*	0.152**	0.143**	0.091***
	(0.067)	(0.064)	(0.067)	(0.064)	(0.071)	(0.065)	(0.066)	(0.072)	(0.067)	(0.069)	(0.033)
For-Profit Δ Pell Grants	1.007***	1.005***	1.005***	1.003***	1.004***	1.003***	1.002***	1.001***	0.999***	1.001***	0.920***
	(0.250)	(0.251)	(0.250)	(0.251)	(0.252)	(0.252)	(0.250)	(0.251)	(0.252)	(0.252)	(0.258)
Difference	-0.932***	-0.911***	-0.930***	-0.908***	-0.876***	-0.860***	-0.856***	-0.864***	-0.847***	-0.858***	-0.829***
Std. Error	(0.248)	(0.248)	(0.248)	(0.248)	(0.247)	(0.247)	(0.244)	(0.246)	(0.245)	(0.246)	(0.256)
Observations	8,432	8,432	8,432	8,432	4,443	4,443	4,443	4,443	4,443	4,443	8,432
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, NP	107.3	97.1	107.1	97.1	85.0	79.9	80.2	85.5	80.3	92.3	-
Δ Pell Grants F-test, FP	28.8	29.8	29.0	29.8	33.4	33.3	33.6	33.0	33.0	35.0	-
Joint F-test	53.7	39.9	53.7	40.3	42.4	20.6	20.6	42.7	20.5	27.3	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at for- and non-profit institutions (see eq. 3.8). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, NP is the robust F-statistic of the first-stage regression of Pell grants at non-profit colleges. Δ Pell Grants F-test, FP is the robust F-statistic of the first-stage regression of Pell grants at for-profit colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.15: Effect of College Spending on Local Income Per Capita

Income Growth	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
Δ Expenditure	0.473	0.695	0.301	0.531	0.904	0.940	0.619	0.658	0.686	0.686	0.422*
	(0.646)	(0.721)	(0.621)	(0.688)	(0.931)	(0.932)	(0.885)	(0.892)	(0.894)	(0.911)	(0.217)
Observations	8,436	8,436	8,436	8,436	4,447	4,447	4,447	4,447	4,447	4,447	8,436
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes		Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test	10.7	8.2	11.2	8.5	6.1	7.6	9.1	7.0	9.1	10.3	-
Joint F-test	10.7	4.3	-	4.5	6.1	3.2	3.7	-	3.8	3.8	-

Notes: SSIV strategy for the college spending regressor uses the twice-lagged share of students in MSA population as shares and yearly changes to total national spending as shocks (see eq. 3.3). SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Robust F-test statistic is for the college spending SSIV. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test is the robust F-statistic of the first-stage regression of Pell grants. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.16: Effect of Pell Grants on Education Expenditures by Four-Year and Two-Year Colleges

Education Exp. Growth	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
4-year Δ Pell Grants	0.850	1.061**	0.854	1.062**	0.945*	1.104**	1.113**	0.960*	1.118**	1.026**	0.914***
	(0.528)	(0.491)	(0.526)	(0.489)	(0.552)	(0.464)	(0.465)	(0.550)	(0.462)	(0.515)	(0.195)
2-year Δ Pell Grants	0.447	0.329	0.444	0.328	0.423	0.334	0.353	0.458	0.364	0.322	1.075**
	(0.301)	(0.281)	(0.301)	(0.281)	(0.333)	(0.306)	(0.307)	(0.334)	(0.307)	(0.331)	(0.420)
Difference	0.403	0.731	0.410	0.734	0.522	0.770*	0.760*	0.502	0.754*	0.703	-0.161
Std. Error	(0.519)	(0.514)	(0.518)	(0.512)	(0.500)	(0.435)	(0.437)	(0.501)	(0.437)	(0.479)	(0.453)
Observations	8,436	8,436	8,436	8,436	4,447	4,447	4,447	4,447	4,447	4,447	8,436
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes	Yes	Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, 4-Year	133.8	119.0	133.9	119.3	95.3	86.0	84.7	93.8	84.5	95.0	-
Δ Pell Grants F-test, 2-Year	65.2	65.0	65.2	65.2	54.0	54.1	54.8	54.9	55.2	57.2	-
Joint F-test	32.5	43.1	32.6	43.4	29.5	26.0	25.6	29.9	25.6	20.1	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at four-year and two-year institutions. We estimate the four-year share based on the MSA four-year penetration and the MSA-level number of recipients where missing. SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, 4-Year is the robust F-statistic of the first-stage regression of Pell grants at 4-year colleges. Δ Pell Grants F-test, 2-Year is the robust F-statistic of the first-stage regression of Pell grants at 2-year colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.17: Effect of Pell Grants on Non-Education Expenditures by Four-Year and Two-Year Colleges

Non-Education Exp. Growth	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
4-year Δ Pell Grants	0.257*	0.331*	0.267*	0.340*	0.418**	0.342	0.329	0.406**	0.330	0.368*	0.239*
	(0.145)	(0.177)	(0.147)	(0.178)	(0.200)	(0.208)	(0.206)	(0.199)	(0.207)	(0.201)	(0.123)
2-year Δ Pell Grants	0.086	0.045	0.079	0.039	-0.054	-0.012	-0.027	-0.059	-0.014	0.020	-0.223*
	(0.195)	(0.201)	(0.195)	(0.201)	(0.203)	(0.195)	(0.197)	(0.208)	(0.199)	(0.214)	(0.128)
Difference	0.171	0.286	0.187	0.301	0.472**	0.353	0.356	0.465**	0.344	0.347	0.461**
Std. Error	(0.191)	(0.220)	(0.190)	(0.219)	(0.212)	(0.242)	(0.243)	(0.213)	(0.244)	(0.236)	(0.181)
Observations	8,436	8,436	8,436	8,436	4,447	4,447	4,447	4,447	4,447	4,447	8,436
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes	Yes	Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, 4-Year	133.8	119.0	133.9	119.3	95.3	86.0	84.7	93.8	84.5	95.0	-
Δ Pell Grants F-test, 2-Year	65.2	65.0	65.2	65.2	54.0	54.1	54.8	54.9	55.2	57.2	-
Joint F-test	32.5	43.1	32.6	43.4	29.5	26.0	25.6	29.9	25.6	20.1	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at four-year and two-year institutions. We estimate the four-year share based on the MSA four-year penetration and the MSA-level number of recipients where missing. SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, 4-Year is the robust F-statistic of the first-stage regression of Pell grants at 4-year colleges. Δ Pell Grants F-test, 2-Year is the robust F-statistic of the first-stage regression of Pell grants at 2-year colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.18: Effect of Pell Grants on Research Expenditures by Four-Year and Two-Year Colleges

Research Exp. Growth	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
4-year Δ Pell Grants	0.173*	0.222*	0.181*	0.229*	0.299**	0.244*	0.235*	0.291**	0.236*	0.264*	0.143*
	(0.104)	(0.117)	(0.106)	(0.119)	(0.145)	(0.135)	(0.136)	(0.144)	(0.135)	(0.140)	(0.079)
2-year Δ Pell Grants	0.003	-0.025	-0.001	-0.028	-0.096	-0.065	-0.074	-0.103	-0.070	-0.046	-0.206***
	(0.169)	(0.175)	(0.168)	(0.175)	(0.183)	(0.169)	(0.174)	(0.188)	(0.174)	(0.184)	(0.074)
Difference	0.170	0.246	0.182	0.256	0.395**	0.308*	0.309*	0.394**	0.305*	0.310*	0.349***
Std. Error	(0.159)	(0.168)	(0.159)	(0.168)	(0.168)	(0.166)	(0.167)	(0.170)	(0.168)	(0.170)	(0.125)
Observations	8,436	8,436	8,436	8,436	4,447	4,447	4,447	4,447	4,447	4,447	8,436
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes	Yes	Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, 4-Year	133.8	119.0	133.9	119.3	95.3	86.0	84.7	93.8	84.5	95.0	-
Δ Pell Grants F-test, 2-Year	65.2	65.0	65.2	65.2	54.0	54.1	54.8	54.9	55.2	57.2	-
Joint F-test	32.5	43.1	32.6	43.4	29.5	26.0	25.6	29.9	25.6	20.1	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at four-year and two-year institutions. We estimate the four-year share based on the MSA four-year penetration and the MSA-level number of recipients where missing. SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, 4-Year is the robust F-statistic of the first-stage regression of Pell grants at 4-year colleges. Δ Pell Grants F-test, 2-Year is the robust F-statistic of the first-stage regression of Pell grants at 2-year colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.19: Effect of Pell Grants on Instruction Expenditures by Four-Year and Two-Year Colleges

Instruction Exp. Growth	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
4-year Δ Pell Grants	0.461	0.647	0.465	0.648	0.521	0.660	0.667	0.530	0.668	0.588	0.622^{***}
	(0.492)	(0.484)	(0.491)	(0.484)	(0.526)	(0.470)	(0.473)	(0.526)	(0.471)	(0.500)	(0.155)
2-year Δ Pell Grants	0.270	0.166	0.268	0.166	0.158	0.081	0.096	0.185	0.103	0.065	0.931^{**}
	(0.252)	(0.243)	(0.253)	(0.244)	(0.285)	(0.277)	(0.279)	(0.288)	(0.279)	(0.289)	(0.416)
Difference	0.191	0.481	0.197	0.482	0.363	0.579	0.571	0.346	0.565	0.523	-0.309
Std. Error	(0.495)	(0.500)	(0.494)	(0.498)	(0.472)	(0.425)	(0.426)	(0.474)	(0.426)	(0.456)	(0.431)
Observations	8,436	8,436	8,436	8,436	4,447	4,447	4,447	4,447	4,447	4,447	8,436
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes	Yes	Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, 4-Year	133.8	119.0	133.9	119.3	95.3	86.0	84.7	93.8	84.5	95.0	-
Δ Pell Grants F-test, 2-Year	65.2	65.0	65.2	65.2	54.0	54.1	54.8	54.9	55.2	57.2	-
Joint F-test	32.5	43.1	32.6	43.4	29.5	26.0	25.6	29.9	25.6	20.1	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at four-year and two-year institutions. We estimate the four-year share based on the MSA four-year penetration and the MSA-level number of recipients where missing. SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, 4-Year is the robust F-statistic of the first-stage regression of Pell grants at 4-year colleges. Δ Pell Grants F-test, 2-Year is the robust F-statistic of the first-stage regression of Pell grants at 2-year colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.20: Effect of Pell Grants on Public Service Expenditures by Four-Year and Two-Year Colleges

Public Service Exp. Growth	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
4-year Δ Pell Grants	0.084	0.109	0.086	0.111	0.118	0.098	0.094	0.115	0.094	0.104	0.096
	(0.107)	(0.126)	(0.107)	(0.125)	(0.113)	(0.130)	(0.129)	(0.113)	(0.129)	(0.117)	(0.079)
2-year Δ Pell Grants	0.084	0.070	0.080	0.066	0.042	0.053	0.046	0.044	0.056	0.067	-0.017
	(0.123)	(0.118)	(0.123)	(0.117)	(0.120)	(0.123)	(0.123)	(0.122)	(0.126)	(0.132)	(0.103)
Difference	0.000	0.040	0.006	0.045	0.077	0.045	0.047	0.071	0.038	0.037	0.113
Std. Error	(0.149)	(0.162)	(0.149)	(0.159)	(0.148)	(0.171)	(0.171)	(0.150)	(0.173)	(0.168)	(0.114)
Observations	8,436	8,436	8,436	8,436	4,447	4,447	4,447	4,447	4,447	4,447	8,436
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes	Yes	Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, 4-Year	133.8	119.0	133.9	119.3	95.3	86.0	84.7	93.8	84.5	95.0	-
Δ Pell Grants F-test, 2-Year	65.2	65.0	65.2	65.2	54.0	54.1	54.8	54.9	55.2	57.2	-
Joint F-test	32.5	43.1	32.6	43.4	29.5	26.0	25.6	29.9	25.6	20.1	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at four-year and two-year institutions. We estimate the four-year share based on the MSA four-year penetration and the MSA-level number of recipients where missing. SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, 4-Year is the robust F-statistic of the first-stage regression of Pell grants at 4-year colleges. Δ Pell Grants F-test, 2-Year is the robust F-statistic of the first-stage regression of Pell grants at 2-year colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table C.21: Effect of Pell Grants on Student Service Expenditures by Four-Year and Two-Year Colleges

Student Services Exp. Growth	Full Sample				Post 1999						Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	OLS
4-year Δ Pell Grants	0.389*	0.414*	0.389*	0.414*	0.424*	0.444*	0.447*	0.430*	0.450*	0.438*	0.291*
	(0.231)	(0.229)	(0.230)	(0.229)	(0.235)	(0.234)	(0.233)	(0.232)	(0.232)	(0.233)	(0.164)
2-year Δ Pell Grants	0.177	0.163	0.176	0.162	0.265	0.253	0.257	0.274	0.262	0.257	0.144***
	(0.135)	(0.128)	(0.134)	(0.128)	(0.170)	(0.162)	(0.161)	(0.169)	(0.161)	(0.164)	(0.035)
Difference	0.211	0.251*	0.213	0.252*	0.159	0.192	0.190	0.156	0.188	0.181	0.148
Std. Error	(0.139)	(0.141)	(0.139)	(0.141)	(0.120)	(0.122)	(0.121)	(0.118)	(0.120)	(0.120)	(0.148)
Observations	8,436	8,436	8,436	8,436	4,447	4,447	4,447	4,447	4,447	4,447	8,436
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Approp.		Yes	Yes	Yes		Yes	Yes		Yes		Yes
MSA Controls			Yes	Yes			Yes	Yes	Yes	Yes	Yes
Financial Controls								Yes	Yes	Yes	
Fiscal Transfers										Yes	
Δ Pell Grants F-test, 4-Year	133.8	119.0	133.9	119.3	95.3	86.0	84.7	93.8	84.5	95.0	-
Δ Pell Grants F-test, 2-Year	65.2	65.0	65.2	65.2	54.0	54.1	54.8	54.9	55.2	57.2	-
Joint F-test	32.5	43.1	32.6	43.4	29.5	26.0	25.6	29.9	25.6	20.1	-

Notes: SSIV strategy for the Pell grants regressor uses the twice-lagged share of recipients in MSA population at four-year and two-year institutions. We estimate the four-year share based on the MSA four-year penetration and the MSA-level number of recipients where missing. SSIV strategy for appropriations uses the twice-lagged appropriation share of income. Controls are twice-lagged. MSA controls: change in undergraduate students (log) in the last 2 years, average tuition fee (log), for-profit penetration, percentage of population black, percentage Hispanic, percentage with at least a bachelor's degree. Data on financial controls is from Federal Reserve Bank of New York/Equifax Consumer Credit Panel and is available from 1999 to 2015. It includes median Equifax Risk Score, age, debt-to-income ratio, credit card utilization, and 30-day mortgage delinquency rate. Fiscal Transfers refers to the total amount of fiscal transfers due to state appropriations, SNAP, UI, and HUD programs. We instrument the fiscal transfers variable with an SSIV analogous to the appropriations SSIV. Δ Pell Grants F-test, 4-Year is the robust F-statistic of the first-stage regression of Pell grants at 4-year colleges. Δ Pell Grants F-test, 2-Year is the robust F-statistic of the first-stage regression of Pell grants at 2-year colleges. Joint F-test is the robust F-statistic of the joint IV set. MSA-clustered standard errors are in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

References

Hausman, J. A., Newey, W. K., Woutersen, T., Chao, J. C., & Swanson, N. R. (2012). Instrumental variable estimation with heteroskedasticity and many instruments. *Quantitative Economics*, *3*(2), 211-255.