The London School of Economics and Political Science

Essays on International Finance and Monetary Economics

Laura Castillo Martinez

A thesis submitted to the Department of Economics of the London School of Economics for the degree of Doctor of Philosophy, London, July 2019.

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). The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. I warrant that this authorization does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of approximately 34,042 words, excluding graphs and tables.

Statement of inclusion of previous work

I confirm that Chapter 3 is a revised version of the paper I submitted at the end of my MRes in 2014.

Statement of conjoint work

Chapter 2 is joint work with Ricardo Reis and I contributed 50% of this work.

Abstract

This thesis examines three dimensions of monetary policy: its implementation through exchange rates, its role in price level determination and its interaction with fiscal policy.

Following a sudden stop, real exchange rates can adjust through a nominal exchange rate depreciation, lower domestic prices, or a combination of both. Chapter 1 makes four contributions to understand how the type of adjustment shapes the response of macroeconomic variables, in particular productivity, to such an episode. First, it documents that aggregate TFP systematically collapses after a sudden stop under a flexible exchange rate arrangement while it moderately improves within a currency union. Second, using firm-level data for two sudden stops in Spain, it shows that the difference in the productivity response is largely driven by firm entry and exit dynamics. Third, it proposes a small open economy DSGE framework with firm selection into production and endogenous markups to explain the empirical findings. Fourth, it uses a quantitative version of the model to revisit the relative performance of exchange rate policies after a sudden stop.

Different theories of the price level are too often presented in terms of opposing camps and conflicting policymakers. As a result, most economists resort to undergraduate monetarist insights to explain inflation, even though for the most part they do not apply to modern advanced economies. Chapter 2 investigates how central banks control inflation in terms of one unified theory that allows for different policy choices by the central bank.

Chapter 3 examines the incentives that shape monetary policy in the context of dual mandates. I present a simple model of optimal monetary policy based on a multitask principal-agent problem framework with two agents. I show there is risk of excessive hawkishness if the compensation scheme is not appropriately designed.

Acknowledgements

I owe a debt of gratitude to my advisors at the LSE for patiently guiding me through the process of becoming a researcher. Silvana Tenreyro has been a role model at both personal and professional levels. Her focus on the big picture, her ability to communicate assertively and her natural optimism, in addition to a solid knowledge of the field, are attributes I hope to take with me. Ricardo Reis' move to London was a real game changer for my PhD experience. He has been incredibly generous in time and advice. Being around him has not only taught me loads about economic theory but also about the profession, the art of teaching and life in general.

I am grateful to many other faculty members with whom I've had the opportunity to interact repeatedly. Francesco Caselli, Ethan Ilzetzki and Per Krusell have often challenged my understanding of economics and encouraged me to go the extra mile. Gianluca Benigno and Gianmarco Ottaviano have given detailed feedback on Chapter 1 and contributed hugely to a successful job market. I also thank Pol Antrás, Charlie Bean, Andrew Bernard, Wouter Den Haan, Keyu Jin, Ben Moll, Catherine Thomas, John Van Reenen and Shengxing Zhang as well as many participants at seminars and conferences for useful questions and comments.

I have been very lucky to share this PhD journey with an inspiring group of young economists. Tiloka and Shan have been referees, editors and enthusiastic supporters in addition to very good friends during these years. I will miss them greatly! My days at the LSE would not have been the same without the Total Football Productivity team, including Andrea, Daniel, Adrien, Charles, Thomas, Andy, Martina, Friedrich, Chao and Sevim. A special thanks goes to Miguel for many many coffees and laps around Lincoln Inn Fields.

Happiness has been a key input in this process. Somewhere in between academia, friendship and family, Miguel Almunia has been able to jump from conversations about economics to music, politics and running. Having his experience some years ahead has been extremely comforting. Friends from outside the PhD have often helped me breath whenever I needed a break from research. In particular, Tere and Moi have always been there despite the distance, while Sonia and Miles have made Uxbridge Road feel even more home.

My family has always supported firmly in all my projects, even when I first chose to study Economics. My parents are an example of resoluteness, courage and integrity and a source of confidence, strength and warmth. My role as Victor's older sister has forced me to be half as good as he thought I was.

Finally, I am fully indebted to Alvaro. While he probably knows more about sudden stops, monetary policy and the academic job market than he would have ever wished, I could not have completed this journey without his unconditional love. This thesis is dedicated to him.

Contents

List of Tables

List of Figures

Chapter 1

Sudden Stops, Productivity and the Exchange Rate

1.1 Introduction

The procyclicality of productivity is a well-established empirical fact in macroeconomics; productivity rises in booms and falls in recessions (Basu and Fernald (2001)).[1] This is not only an essential feature of closed-economy business cycle models, but a crucial ingredient in balance of payment crises. The disruption of economic activity that accompanies unexpected reversals in capital flows, also known as sudden stops, is associated with declines in aggregate total factor productivity (TFP).[2] The recent European sovereign debt crisis challenges this interpretation: in this instance, productivity improved moderately as the external adjustment unfolded. The fact that it took place within a currency union, another singularity of this episode, emerges as a potential explanatory factor and raises the following question: what is the relationship between sudden stops, productivity and the prevailing exchange rate regime? How does accounting for this complement the inherited wisdom on fixed versus floating regimes?

This paper studies how the type of real exchange rate realignment shapes the response of macroeconomic variables, in particular productivity, to a sudden stop. It shows that it shows that the difference in the productivity response is largely driven by firm entry and exit dynamics. Internal devaluations, as opposed to nominal depreciations, lead to greater exit of unproductive firms, contributing to positive aggregate TFP growth through a so-called cleansing effect. This implies that incorporating firm dynamics and allowing for demand effects in an otherwise standard open economy framework is key to capturing the observed cyclicalities under different exchange rate regimes. The result is that the type of external adjustment,

[1] Several authors have argued that the cyclicality of total factor productivity and labor productivity relative to inputs has reduced or even reversed since the mid-1980s in the US; see Galí and Van Rens (2014) and Fernald and Wang (2016).

[2] This paper uses the terms productivity and TFP interchangeably.

whether nominal or real, affects macroeconomic performance even in the absence of nominal rigidities.

Section 1.3 provides systematic evidence on the behavior of macroeconomic variables during a sudden stop for both developed and developing economies during the 1990-2015 period. I use a new criterion to identify sudden stops that captures both the episodes discussed previously in the literature as well as the recent Southern-European cases. I classify sudden stops according to the prevalent exchange rate regime and evaluate the response of macroeconomic variables in each regime using an event study approach. Two regularities stand out: first, aggregate TFP systematically collapses under flexible exchange rate arrangements, while it improves, albeit moderately, within currency unions. Second, in a currency union, there is a larger decline in employment and a greater contraction in consumption as a share of GDP following a sudden stop. In order words, the increase in aggregate TFP comes at the expense of a greater domestic contraction.

To disentangle the drivers of aggregate TFP performance, Section 1.4 presents micro evidence in the form of firm-level data from the manufacturing sector. More specifically, I exploit survey data during two sudden stop episodes that occurred in Spain: the 1992-93 Exchange Rate Mechanism crisis and the 2009-13 European sovereign debt crisis. During the former, the national currency, the *peseta*, depreciated on multiple occasions and TFP fell by over 10 percent. During the second episode, Spain was a member of a currency union and could only regain competitiveness by lowering wages. In this case, TFP increased by 10 percent.

The firm-level analysis of these two episodes uncovers the following empirical findings. First, changes in productivity are concentrated in the lower tail of the firm productivity distribution in both episodes. Second, while productivity was declining at the firm-level during both crises, unproductive exiting firms contributed substantially more to TFP growth in the 2009-2013 sudden stop. Third, the 2009-13 sudden stop had a cleansing effect on productivity while the 1992-93 sudden stop did not. Finally, there is suggestive evidence that links heterogeneity in price markups with changes in allocative efficiency throughout the full sample.

Based on the previous evidence, Section 1.5 develops a dynamic stochastic general equilibrium (DSGE) model with a micro-structure borrowed from Melitz and Ottaviano (2008) to study the behavior of productivity, output and employment during a sudden stop. The model features a small open economy with quasi-linear quadratic preferences and firm heterogeneity in productivity. This gives rise to firm selection into production and endogenous variable markups. I extend this model by including leisure in the consumer's utility function, thereby deriving the labor supply decision. This provides a new channel through which the wage level and individual firm profits interact.

To allow a role for policy, I introduce nominal rigidities in the wage setting process. The central bank chooses the nominal exchange rate as its main policy tool and two extreme regimes are discussed: a currency union, characterized by a credible commitment to keep the nominal exchange rate constant; and a strict wage inflation targeting regime, where the flexible wage equilibrium is always implemented. A sudden stop is an exogenous shock to the risk premium component of the interest rate that domestic consumers pay for international borrowing. By increasing the return on foreign denominated bonds, the domestic economy is forced to save internationally and increase net exports through a real exchange rate depreciation.

Section 1.6 uses a partial equilibrium version of the model to build intuition on how it is able to generate the observed TFP patterns following a sudden stop. The key insight is that aggregate productivity is proportional to the domestic productivity threshold. The threshold represents the minimum productivity level at which a firm can generate positive profits and, thus, select into the domestic market. It therefore suffices to understand how the threshold moves after a sudden stop to learn about its effect on aggregate productivity.

In equilibrium, the domestic threshold is entirely determined by the number of active firms in the market and the wage level. Therefore, there are three mechanisms through which a shock can affect productivity. First, the threshold increases with the number of active firms, as greater competition lowers profit margins for all firms and then requires a higher level of productivity to remain profitable. This is the pro-competitive channel. Second, higher wages increase the costs of production for all firms, lowering again their profit margin and calling for a higher productivity level. This is the cost channel. Third, higher wages also increase the demand for overall consumption by increasing households' labor income. This, instead, increases the firm profit margin and relaxes the productivity requirement. This is the demand channel.

The effect of a sudden stop on the domestic productivity threshold will hinge on the relative strength of these conflicting forces. This, in turn, depends on how the real exchange rate adjustment is conducted. More precisely, whether it takes place through the depreciation of nominal exchange rates or a lower wage level. Consider the two polar cases, in the first case the nominal exchange rate bears the full brunt of the adjustment: only the pro-competitive channel works, fewer firms import, and productivity falls unambiguously. In contrast, when the nominal exchange rate is fixed, the wage adjusts completely and all three channels operate, resulting in a quantitatively ambiguous overall effect. The model delivers conditions under which the demand channel dominates, allowing a sudden stop to generate a productivity improvement in a currency union.

Section 1.7 extends the analysis to general equilibrium by calibrating the model using Spanish macroeconomic data as well as the firm-level evidence presented in Section 1.4. I simulate the response of the economy to an unexpected exogenous increase in the country risk premium component of the interest rate under the two alternative exchange rate policy regimes. The model predictions mimic a sudden stop episode: the economy runs a current account surplus and the real exchange rate depreciates. The responses of macroeconomic variables match the empirics along many dimensions: GDP and consumption fall, with a more pronounced drop of the latter under the currency union regime. Moreover, relative employment dynamics are correctly captured by the generated impulse response functions with the currency union experiencing greater volatility, although the model does not generate an absolute decline in employment under a floating regime.

The baseline general equilibrium framework fully captures the empirically documented TFP fact: productivity falls when the real exchange rate depreciation translates one-to-one into a nominal depreciation, while it increases as the devaluation takes place through wages instead. This result is not only robust to alternative parameterizations of the model, but also to a range of extensions to the baseline set-up presented in Section 1.8.

Finally, to evaluate the overall performance of exchange rate regimes following a sudden stop, Section 1.9 studies how TFP improvements translate into welfare gains by computing the cumulative output loss under a currency union and a floating arrangement. These findings show that while higher nominal rigidities are more harmful in a currency union, at low levels of wage stickiness, the floating arrangement performs worse. The latter effect is driven by the opposing effect of lower wages and appreciated exchange rates on the cost of entering the market, which can partly compensate for the drop in domestic demand in the currency union.

The results contrast with the standard view that both exchange rate regimes would lead to similar economic outcomes in the event of an external adjustment in a perfectly flexible world. To better understand this finding, I explore further how the relative performance of regimes depends also on the labor income share, the degree of firm heterogeneity and the complementarity of foreign and domestic labor inputs. I also show that results are robust to evaluating performance either by the cumulative consumption loss or by a utility-based welfare measure.

1.2 Relation to the literature

This paper combines several strands of the literature at the intersection of international finance, trade theory and firm dynamics.

First, it focuses on sudden stops, as defined by Calvo (1998), abrupt and unexpected reversals in foreign capital inflows. It follows the empirical research that documents regularities among historical sudden stop episodes including Calvo et al. (2004), Guidotti et al. (2004), Calvo and Talvi (2005) and Kehoe and Ruhl (2009) and contributes to their previous analysis in three ways: by modifying the Calvo sudden stop identification methodology to account for gradualism, by expanding the time frame and the set of economies traditionally considered, and by classifying episodes according to the flexibility of the nominal exchange rate.[3] The results show that previous findings - current account adjustment, depreciation of the real exchange rate and fall in output and TFP - apply to economies with flexible exchange rates but not fully to currency unions.

Related, several articles propose amendments to the standard open economy neoclassical model with flexible exchange rates in order to reconcile theoretical predictions with the observed behavior of macroeconomic variables, especially TFP, during a sudden stop. For example Meza and Quintin (2007) allow for endogenous factor utilization, Neumeyer and Perri (2005), Christiano et al. (2004) and Mendoza (2006) introduce advanced payments of inputs and Mendoza (2010) directly assumes that exogenous productivity shocks trigger the collateral constraints that drive sudden stops. I bring an alternative explanation to the table: selection into production.

Second, while the goal of the paper is not the normative analysis *per se*, the implications of the model link to the floating versus fixed exchange rate debate initiated by Friedman (1953). Cúrdia (2007), Braggion et al. (2009) and Fornaro (2015), among others, argue that policy rules that target inflation are superior in terms of welfare to those that prioritize exchange rate stability. More recently, Farhi et al. (2013) Schmitt-Grohé and Uribe (2016) and Galí and Monacelli (2016) have emphasized the interaction of wage flexibility and the exchange rate regime. Accounting for firm dynamics preserves the importance of nominal rigidities in evaluating the relative performance of policy rules, while it allows for different outcomes across regimes even in their absence.

The third strand of the literature to which this paper closely relates is trade models of heterogeneous firms à la Melitz (2003), which emphasize firm selection into domestic and international markets by featuring fixed production and exporting costs.[4] Even though the main focus of these papers is on the welfare effects of trade liberalizations, the real exchange rate depreciation that results from a balance of payment crisis resembles an asymmetric trade liberalization.[5] This makes the New New Trade Theory framework a suitable starting point for this analysis.

Unlike the canonical Melitz (2003) model, I do not restrict attention to competition in the labor market and incorporate pro-competitive effects of trade by departing from constant elasticity of substitution (CES) preferences. In doing so, I follow Bernard et al. (2003), Feenstra (2003), Behrens and Murata (2006) and, more specifically, Melitz and Ottaviano (2008), to feature endogenous variable markups. This allows me to defuse, at least in the short-run, the negative dependence between domestic and exporting cutoffs that drives the baseline result in the Melitz

[3]This is not, however, the first paper to classify the massive reversals that Southern-European countries experienced between late 2009 and 2011 as sudden stop episodes (see Merler and Pisani-Ferry (2012) and Gros and Alcidi (2015) for an earlier discussion).

[4]For a review of the literature, refer to Melitz and Redding (2014).

[5]More precisely, a real exchange rate depreciation can be modelled as a simultaneous increase in import tariffs and export subsidies.

(2003) model. I preserve, however, the role that wages play in such model within the Melitz and Ottaviano (2008) framework. My approach is to modify the quasi-linear-quadratic preferences proposed by Ottaviano et al. (2002) and assume that leisure is the homogeneous good. The result is that the demand for differentiated varieties is no longer independent of labor income and a new demand effect of wages emerges.

While most of the standard trade models are static, this paper is closer to the subset within the trade literature interested in firm dynamics and business cycles.[6] The most notable reference is Melitz and Ghironi (2005), which embeds the steady-state version of the Melitz (2003) model into a two country DSGE setting. To gain tractability, however, Melitz and Ghironi (2005) assume that all firms that enter the market generate positive profits and, thus, firm exit is exclusively determined by exogenous death shocks.[7] In addition, because their focus is on the Balassa-Samuelson effect, the main driver of the business cycle is an aggregate productivity shock. My paper, instead, incorporates selection into production and studies the effects of exogenous current account shocks on productivity.[8] Moreover, it is, to the best of my knowledge, the first study to incorporate nominal rigidities and, thus, to be able to discuss the role for monetary policy in a similar setting.[9]

Finally, this paper is connected to the literature that studies the contribution of reallocation to TFP growth. Two theoretical arguments have been put forward to date. On the one hand, Hsieh and Klenow (2009) show that increases in allocative efficiency that involve closing gaps in the return of inputs increase aggregate productivity. On the other hand, Caballero et al. (1994)'s interpretation of Schumpeter et al. (1939)'s creative destruction emphasizes the role of reallocation among new and incumbent firms as an important factor of growth. The paper presents a model based on the second current, emphasizing the cleansing effect of internal devaluations, but discusses both conjectures in the empirical analysis.

The pre-crisis slowdown of productivity in Southern Europe has prompted an increasingly popular narrative that links declining TFP and enhanced misallocation with capital inflows. Papers on this topic are often grouped according to the margin of misallocation suggested: Benigno and Fornaro (2014) considers a model of misallocation between a tradable and a non-tradable sector; Dias et al. (2016), García-Santana et al. (2016), Reis (2013) and Gopinath et al. (2017) show, and formalize in the latter two cases, that resources were also misallocated within sectors. My work contributes to these hypotheses in two dimensions: from the empirical side, I show

[6]Alternatively, models are commonly assumed to be dynamic but with a stationary equilibrium featuring constant aggregate variables.

[7]As opposed to a framework in which the presence of fixed production costs drives firms generating negative profit out of the market.

[8]Bilbiie et al. (2012) and Ottaviano (2012), among others, already consider the effect of endogenous entry and/or new product variety in a business cycle model. They do so, however, in a closed economy setting.

[9]Bilbiie et al. (2008) and Bilbiie et al. (2014) introduce price adjustment costs in a closed-economy DSGE model with endogenous entry and product variety to study optimal monetary policy.

that the negative relationship between capital flows and TFP growth in Southern Europe is symmetric, that is, productivity improved as the crisis hit and foreign capital retrenched. From the theoretical side, I provide an alternative framework that reconciles both sets of papers by endogenizing firms' decision to export and, thus, the size of the exporting sector. In addition, I abstract from the theoretical emphasis on capital that characterizes these papers and consider an alternative dimension of misallocation: variable markups.

1.3 Aggregate productivity during a sudden stop

How do unexpected capital flow reversals affect macroeconomic performance? To answer this question I proceed in three steps: (i) establishing a criterion to identify sudden stops; (ii) classifying episodes by exchange rate regime; and (iii) characterizing the behavior of several macroeconomic variables using a standard event study approach.

1.3.1 Data and methodology

Following Cavallo and Frankel (2008), I define a sudden stop as an episode in which there is a substantial decline in the capital account surplus together with a reduction in the current account deficit and a simultaneous recession. I develop an algorithm that classifies as a sudden stop a period that contains at least one year during which (i) the financial account surplus has fallen at least one standard deviation below its rolling average; (ii) there is a simultaneous decline in the current account deficit (or an equivalent decline in foreign reserves); and (iii) GDP per capita contracts.[10] The start and end of each episode is marked by the first and last year within the period in which the financial account surplus is half a standard deviation below the rolling average.[11]

The two latter requirements ensure that the capital flow reversals captured by the algorithm strictly qualify as sudden stops. First, by requiring that the financing disruption is accompanied by an appropriate macroeconomic adjustment. Second, by ruling out booming episodes that display similar characteristics, for example a positive trade shock.

Annual data on the current and capital accounts for all available countries comes from the IMF's International Financial Statistics Database (IFS) for the period 1990-2015 and complemented with data on GDP per capita growth from the World Bank's World Development Indicators Database.[12]

[10] This contrasts with Cavallo and Frankel (2008), who consider a reduction in the financial account surplus that is two standard deviations above the mean standard deviation for the corresponding decade.

[11] Refer to Appendix 1.11.1 for further details.

[12] I do not consider countries which are small, both in terms of population (below one million inhabitants) and in terms of GDP (below one billion USD). The final sample covers 119 countries.

The total number of episodes is 78, representing 5.2% of total available country/year observations in the sample.[13] The criterion successfully captures all traditional sudden stop episodes previously discussed by the literature - mostly occurring around the 1994/5 Tequila crisis, the 1997 Asian Financial Crisis, the 1998 Russian default - as well as the most recent balance of payment crisis in the peripheral economies of the European Union.[14]

I build on Ilzetzki et al. (2017) updated *de facto* coding system as opposed to relying on declared exchange rate regime reported to the IMF in order to classify episodes. In my baseline results, I consider as prevalent the exchange rate regime that is in place during the last year of the sudden stop. There are four different cases: a currency union, a hard peg, a soft peg and a floating arrangement.[15] Out of the 78 episodes identified, 11 occur within a currency union (8 in the Euro Area and 3 in the West African Economic and Monetary Union), 14 in a hard peg system, 26 in a soft peg regime and 25 in a floating arrangement.

To characterize the behavior of the macro-economy as a sudden stop unfolds I use data on GDP, final private consumption, employment, total factor productivity, current account deficit and real exchange rate. All variables are compiled from the World Development Indicators except for Total Factor Productivity, TFP, that is collected from the Conference Board's Total Economy Database and the current account deficit from the IMF's World Economic Outlook Database.

Figures 1.1 and 1.2 show the mean and median path of each of these aggregate variables during the episodes conditional on their exchange rate classification together with standard error bands. In order to capture the buildup and end phase of each episode, the plot depicts six-year windows that begin two years before the start of each reversal and marks the start and the average duration of a sudden stop with vertical lines. As is standard in much of the literature, I focus on the cyclical component of most of the variables by looking at its percentage deviation from an extrapolated pre-crisis linear trend.[16]

Figure 1.1: Sudden stops in a floating arrangement

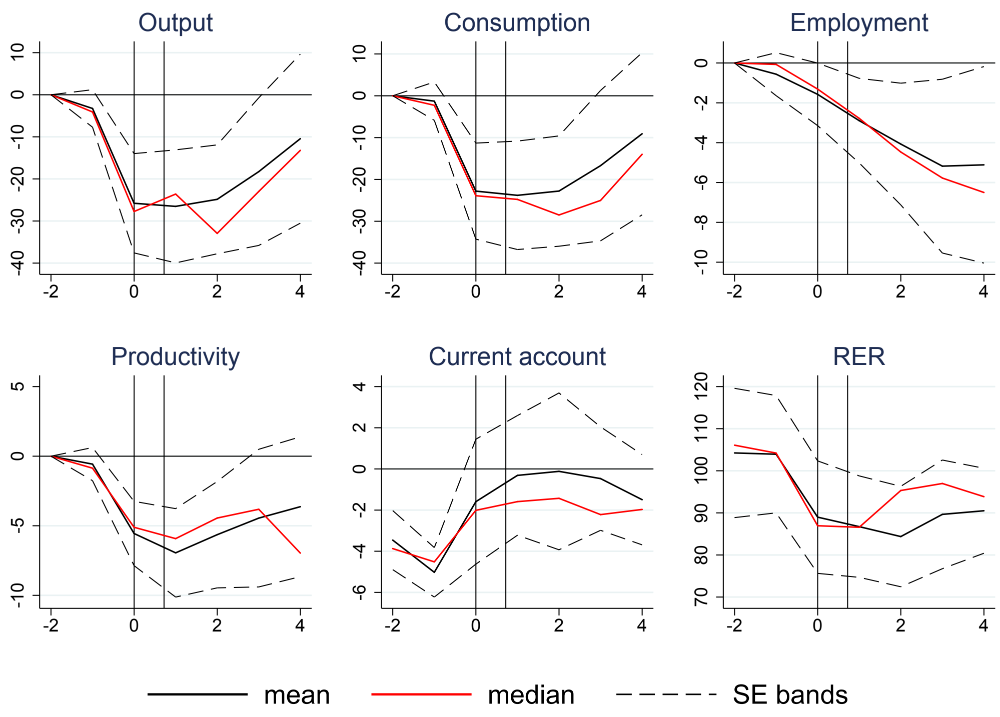

Notes: This figure plots the response of macroeconomic variables to a sudden stop under a floating arrangement. The black and red solid lines depict the mean and median path of the corresponding variables while the black dashed lines represent standard error bands. The two vertical lines show the start and end of an average episode. Output, consumption, employment and productivity are expressed in terms of percentage deviations from an extrapolated linear trend calculated from periods $t-5$ to $t-2$. Current account is expressed as a share of GDP and the real exchange rate (RER), calculated as an index, is expressed in levels.

Source: IFS, WDI, Total Economy Database and own calculations.

1.3.2 Results

Figure 1.1 illustrates how domestic variables respond to an unexpected reversal of capital flows when the exchange rate is allowed to adjust freely. First, a sudden

[13] The total number of episodes is 105. However, I drop one-year long episodes that start in 2009 as these are explained by the global trade collapse rather than by a country-specific reversal of flows. The full list of episodes per country, plus exchange rate classification, can be found in Appendix 1.11.1.

[14] The methodology does not account for changes in TARGET2 balances in the Eurozone and, thus, prevents me from measuring private capital flows accurately. However, this is not problematic for my purposes as the algorithm already identifies the GIIPS episodes.

[15] In terms of the Ilzetzki et al. (2017) fine classification, I deviate as follows: (1) I manually divide code 1 into currency union and no separate legal tender, (2) I group codes 2 to 4 under the hard peg category, (3) I group codes 5 to 11 under the soft peg category, (4) I group codes 12 to 14 under the floating arrangement and (5) I rename group 15 as 5 *i.e.* other categories.

[16] The current account deficit, expressed as a share of GDP, and the real exchange rate index, with base t-2, are the exception.

stop is associated with a contraction in output and consumption, with most of the decline occurring on impact or shortly after. There is also a smooth decline in employment levels, measured as the number of employed workers, and a significant collapse in total factor productivity. The last two graphs capture the response of the external sector: capital outflows coincide with a depreciation of the real exchange rate, represented by a decline in the index plotted in Figure 1.1. The current account deficit is reduced sharply, almost reaching trade balance as soon as one year after the start of the episode. Finally, the average duration is slightly less than two years.

The results for a currency union are summarized by Figure 1.2. The response of all variables but TFP is similar, in qualitative terms, to that depicted in the flexible exchange rate case. The unexpected reversal of flows is associated with a decline in output, consumption and employment. There is a gradual reduction in the current account deficit that yet persists four years after the onset of the crisis. In line with this result, the real depreciation is more gentle than in the previous case and the episodes last longer, on average, two and a half years.

The most notable difference across the plots is the behavior of TFP: whereas productivity clearly falls in the first case, in line with the findings of the literature, it remains unchanged or, if anything, improves slightly within currency unions.[17] A closer look into individual episodes shows that sudden stops in currency unions are preceded by periods of worsening TFP performance which, at least, slow down as the capital flows reversal materializes. In contrast, periods of capital inflows in free-floating regimes are characterized by increasing TFP records that completely flip as soon as the sudden stop hits the economy.

There are additional, although arguably minor, differences in responses across regimes that are worth highlighting. Although a quantitative comparison is beyond the scope of this exercise, the decline in employment is more pronounced in Figure 1.2. This holds in both absolute and relative to GDP terms. Moreover, controlling for the size of output contraction, the fall in private consumption is larger in the currency union.

1.3.3 Robustness

I conduct a battery of robustness checks to evaluate the consistency of these findings.[18] Regarding the exchange rate classification, I consider alternative *de facto* coding systems, such as Shambaugh (2004) and Klein and Shambaugh (2008), that allow for regime changes in higher frequency. I also redo the analysis taking as given the exchange rate regime prevalent at the start of the sudden stop. This is motivated by the fact that, although in most cases countries abandoned pre-existing pegs

[17] Given the reduced sample size in Figure 1.2, standard error bands are admittedly large to be able to conclude that TFP increases significantly.

[18] Results available upon request.

Figure 1.2: Sudden stops in a currency union

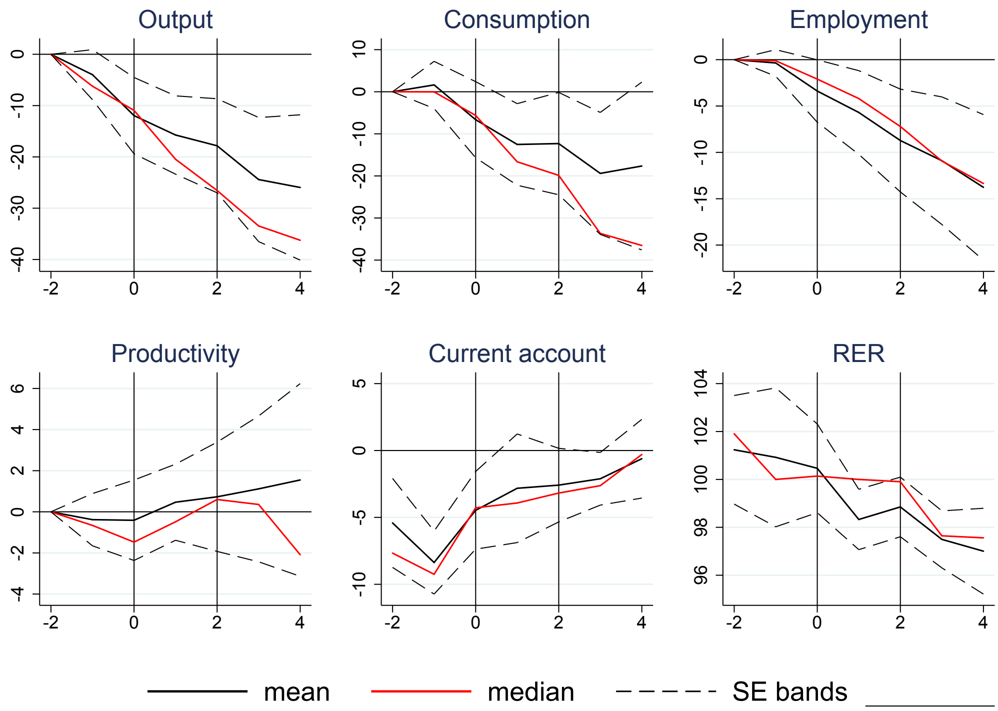

Notes: This figure plots the response of macroeconomic variables to a sudden stop under a currency union. The black and red solid lines depict the mean and median path of the corresponding variables while the black dashed lines represent standard error bands. The two vertical lines show the start and end of an average episode. Output, consumption, employment and productivity are expressed in terms of percentage deviations from an extrapolated linear trend calculated from periods $t-5$ to $t-2$. Current account is expressed as a share of GDP and the real exchange rate (RER), calculated as an index, is expressed in levels.

Source: IFS, WDI, Total Economy Database and own calculations.

because of a sudden stop, there are cases in which failed currency pegs led to capital outflows. Moreover, I remove episodes in which the exchange rate regime changed more than once as this is exclusively due to missing data. None of the alternative methods change the main conclusions discussed above.

Regarding the event study, I explore different ways of detrending the data including a one-sided HP filter and alternative pre-crisis sample lengths. I also measure labor input as total hours worked instead of employment. To control for changes in the composition of the sample, I redo the analysis including only episodes for which all six years of data are available. Finally, I control for the degree of economic development and show that results are not driven by advanced versus developing structural differences. Results hold for all of these specifications too.

1.4 Firm-level productivity during a sudden stop

What lies behind the observed difference in TFP performance across exchange rate regimes? This section resorts to micro-evidence to document the role of selection into production and firm-level TFP in explaining aggregate productivity patterns.

1.4.1 Spain: a tale of two sudden stops

Given limited availability in firm-level data, I use a case study approach to study firm dynamics during a sudden stop. In order to control for country fixed effects, it is preferable to compare episodes occurred within the same economy. Based on the results in Section 1.3, Spain emerges as a natural candidate. It has experienced two sudden stops in its recent economic history under the two exchange rate regimes of interest: the first coincides with the 1992-93 Exchange Rate Mechanism (ERM) crisis and the second with the 2009-2013 European sovereign debt crisis.[19] Moreover, the Spanish economy has been previously analyzed within the misallocation debate (see Gopinath et al. (2017)) and, thus, makes for an interesting benchmark for comparison.

There are clear parallels between the two episodes regarding the onset. Both were preceded by periods of increasing capital inflows, declining international competitiveness and widening current account deficits. Capital inflows were abruptly reverted following a confidence shock affecting the European integration project: the negative outcome of the Danish referendum on the Maastricht Treaty in the first case, and the Greek announcement of substantial upward revisions in the government budget deficit more recently. The flight of international investment led to an urgent correction of misaligned real exchange rates and a boost in exports in order to close the trade gap.

The response of exchange rate policy to these events, however, diverted significantly. While the *peseta* was devalued in three occasions during the 1992-93 crisis, Spain already shared a common currency with its largest trading partners by 2009 and underwent a process of internal devaluation.[20] Consistent with my previous results, TFP fell following the nominal depreciation in 1992, while it increased during the 2009-13 period. I take these episodes as representative of sudden stops under floating arrangements and currency unions, respectively, and use firm-level data to explore what is driving the observed aggregate TFP pattern.[21]

[19]There are two other countries which have experienced sudden stops under different exchange rate regimes: Finland and Italy. The episodes also correspond to the ERM and the European sovereign debt crisis. The focus on Spain is partly driven by data availability.

[20]In 1992, the *peseta* was first devalued by 5% on September 17^{th}, known as Black Wednesday, when the pound and the lira abandoned the ERM altogether. A further 6% was devalued on November 23^{rd}, with a third devaluation taking place in May 1993.

[21]It can be argued that Spain does not strictly classify as a floating exchange rate regime in 1992-93 as it remains a member of the Exchange Rate Mechanism, a multilateral party grid of exchange rates established in 1979. However, the repeated realignments of its central rate against

1.4.2 Data

I use firm-level data from the Survey on Business Strategies (Encuesta sobre Estrategias Empresariales, ESEE, in Spanish) managed by the SEPI Foundation, a public entity linked to the Spanish Ministry of Finance and Public Administrations. The ESEE surveys all manufacturing firms operating in Spain with more than 200 workers and a sample of firms between 10 and 200 workers, providing a rich panel dataset with over 1,800 firms for the period 1990-2014.[22] It covers around 35 percent of value added in Spanish manufacturing and provides information on each firm's balance sheet together with its profit and loss statement.

The main advantage of ESEE, especially over the ORBIS dataset compiled by Bureau van Dijk Electronic Publishing (BvD), is that it closely captures the extensive margin of production.[23] This is particularly true for the exit of firms as the dataset clearly differentiates between firms that decide not to collaborate in a given year, firms that exit the market and firms that are affected by a split-up, a merger or an acquisition process. In addition, firms that resume production or collaboration with the survey are re-included in the sample and properly recorded. As for entry, new firms are incorporated every year in order to minimize the deterioration of the initial sample. These include all entrants with more than 200 workers and a random selection representing 5% of those with 10 to 200 workers.

There are other advantages of the ESEE dataset that are also worth highlighting. It is the only dataset with reliable financial information going back as early as the beginning of the 1990s, allowing me to study the 1992-93 episode. It also provides firm-level records of the value of exports which is most often subject to stringent confidentiality rules in Spain.[24]

Finally, I drop the entire firm record, instead of the corresponding firm-year observation, when conducting standard consistency checks on the data. The reason is that I want to prevent firms disappearing (and maybe then reappearing) in the

the *deutsche mark* implied that the overall devaluation of its currency was even larger than that of the floating currencies such as the pound. In order words, despite the formal membership of the ERM, the exchange rate effectively behaved as flexible.

[22]Large firms will be overrepresented in my sample. Given that firms that enter are typically small while those that exit range from small to medium-sized, this could potentially weaken the role of extensive margin in my analysis. Keeping this caveat in mind, my findings should be interpreted as a lower bound.

[23]The other existing firm-level dataset, as used in García-Santana et al. (2016), is the Central Balance Sheet Data (Central de Balances Integrada, CBI, in Spanish) owned by the Bank of Spain and only accessible to in-house economists. This alternative dataset, however, is built using the same source of data that constitutes the Spanish input for ORBIS, annual financial statements that firms are obliged to submit to the Commercial Registry, and, thus, is subjected to the same criticism. Please check Almunia et al. (2018b) for more details.

[24]To the best of my knowledge, the only available dataset is the foreign transactions registry collected by the Bank of Spain containing transaction-level data which can be aggregated to the firm-level using the firm's fiscal identifier as done in Almunia et al. (2018a). Given the administrative nature of the dataset, however, only large operations are recorded. Moreover, the minimum reporting threshold changed from 12,500 to 50,000 euros in 2008, hindering the possibility of correctly measuring the extensive margin of exports.

sample strictly due to the cleaning procedure. This is important to correctly capture entry and exit to the market. The efforts devoted to ensure consistency and accuracy during the ESEE data collection process minimize the loss of observations resulting from this requirement. I only leave out firms that report zero or negative values of value added or capital stock.

1.4.3 Estimating firm-level TFP

I measure real output as nominal value added divided by an output price deflator. Obtaining an appropriate industry-specific output price deflator series has proved to be challenging for two reasons. First, the data needs to go back in time at least until 1990, while Eurostat series, the standard source, only start around 2000. Instead, I use the producer price index provided by the Spanish National Statistics Institute (NSI). Second, the ESEE provides its own industry classification based on the sum of the three-digit NACE Rev.2 codes to 20 manufacturing industries. Given that the mapping is not strictly one-to-one, deriving corresponding industry-specific deflators requires implementing a weighting strategy.[25] My approach is to use sector contribution to total manufacturing value added in 2018, also provided by the NSI, as the relevant weight.[26]

I follow the literature in using the wage bill, deflated by the above price series, instead of employment to measure the labor input, in order to control for heterogeneity in labor quality across firms. To measure capital stock I use two different variables given existing data restrictions: for the 1990-1999 period I use total real net capital stock whereas for the 2000-2014 period I use the book value of fixed assets deflated by the price of investment goods from the Spanish National Statistics Institute.[27,28]

The standard practice is to estimate industry output elasticities for capital and labor by regressing value added on input choices and to compute firm-level productivity as the Solow residual.[29] When performing the first step, two potential problems emerge. First, productivity is unobservable and strongly correlated with input choices. A simple OLS regression will therefore deliver biased estimates of the desired elasticities because of simultaneity. Second, there is a selection bias due to

[25] For example, manufacturing industry with ESEE code 7 (paper) corresponds to NACE Rev.2 codes 171 and 172.

[26] The NSI provides weightings for the 2010-2018 period only. I use 2018 figures, as opposed to taking an average or an alternative year, because 2018 is the only year for which there are no missing values.

[27] Total real net capital stock is defined as the value of the stock of total net capital at 1990 constant prices which I simply convert into base year (2015) prices.

[28] I conduct several robustness exercises in order to check whether the change in the capital stock measure has an impact on the results. First, for the years for which the two series overlap, 1993-1999, I estimate that the correlation coefficient at the firm-level is 0.9. Second, for the 1993-1999 period, I estimate the production function using the two series separately and then compare resulting coefficients - for 18 out of 20 industries the differences are of magnitude ± 0.5 on average. Finally, I redo the analysis splitting the sample before and after 1999 such that the two series do not interact in any way during the production function estimation stage.

[29] See Appendix 1.11.2 for a more detailed review of production function estimation techniques.

the fact that firm survival is related to the unobserved productivity level: firms that remain in the sample tend to be the most productive ones.

To overcome the former issue, I follow the proxy variable approach (see Olley and Pakes (1996) and Levinsohn and Petrin (2003)) among the possibilities offered by the literature.[30] Intuitively, this method substitutes unobserved productivity by a proxy variable in the original regression; where a proxy variable is an observable input or choice variable for which the mapping with respect to productivity is assumed to be invertible. Coefficients of the inputs that do not enter this mapping, mainly labor, can be non-parametrically estimated using OLS in a first stage. The remaining coefficients, capital, are estimated next by exploiting the zero correlation assumption between the unexpected component of productivity and the input choice using GMM. I use materials deflated by the output price deflator as the proxy variable. To account for labor dynamics, however, I implement the refinement introduced by Ackerberg et al. (2015) that consists of identifying all coefficients in the second stage by using conditional (as opposed to unconditional) moments.[31]

To control for attrition, I include an intermediate stage in which the probability of survival is estimated by fitting a probit model on materials, labor and capital in the spirit of Olley and Pakes (1996). This probability is then included as a regressor in the final stage.

1.4.4 Analysis

Aggregate TFP, defined as the labor-weighted average of firm-level TFP, decreased by 10.87% during the 1992-1993 episode while increased by 10.02% in the 2009-2013 period.[32] While consistent with the results of the event study, the granularity of the data allows for a more detailed investigation regarding the drivers of productivity.

The lower tail I first document changes in the distribution of firm-level productivity before and after each of the crises. Figure 1.9 plots a kernel probability distribution estimate of log TFP before and after a sudden stop for both the 1992-1993 and the 2009-2013 episodes. A number of patterns stand out. First, there is ample heterogeneity in TFP levels among firms in any given year as already highlighted

[30] The other alternatives are fixed effects, instrumental variables, first order conditions and dynamic panel approach.

[31] In addition to accounting for labor dynamics, Ackerberg et al. (2015) improves on the Wooldridge (2009)'s extension of the Levinsohn and Petrin (2003) approach by allowing for unobserved serially correlated shocks to wages. Their framework also overcomes Gandhi et al. (2016)'s concern regarding the non-identification result of the proxy variable approach by assuming a Leontief production function in materials. As a robustness check, nevertheless, I show that these two alternative methodologies generate firm-level TFP series which are highly correlated with my baseline TFP.

[32] I consider labor, as opposed to value added, weights when aggregating TFP for two reasons. On the one hand, I will be presenting a theoretical model with labor as the only factor of production where labor shares are the appropriate weight. On the other, large firms in terms of employment are overstated in my sample, as explained above, and, thus, labor weights are consistent with the interpretation of my results as a lower bound.

by the literature. Second, the shape of the distribution is reasonably similar and remains unchanged throughout both crisis periods. Third, changes in TFP are not explained by major shifts in the distribution. A visual inspection suggests that the lower tail concentrates most, if not all, of the action: it lengthens as TFP decreases in the former crisis while shortens considerably as TFP increases in the latter case.

Figure 1.3 summarizes graphically the predominant role of the lower tail by presenting the mean change in log TFP per percentile of the distribution. On average, the change in productivity is close to zero during both episodes across the entire distribution, with the notable exception of the 1% percentile where TFP decreases by 70% during 1991-1993 while increases by 73% during 2009-2013.[33] Although the standard errors are admittedly large for the 1% percentiles in both cases, the difference relative to other percentiles is large enough to remain relevant.

Estimated moments of the distribution confirm the above hypothesis with higher-order moments experiencing the largest swings.[34] During the 1992-93 crisis firms display lower productivity on average and the dispersion of log TFP increases. The increase in dispersion, however, is asymmetric. The distribution of unproductive firms expands while that of productive changes little with the coefficient of skewness declining from -0.40 to -1.24. Moreover, increasing kurtosis, 7.04 versus 10.42, is associated with fatter tails as the probability mass moves away from the shoulders of the distribution. Although the behavior of TFP exactly reverses during the 2009-2013 crisis - productivity increases while dispersion drops - it is still the tails, and especially, the lower tail, that changes the most. In this case, skewness increases from -2.37 to -0.89 while kurtosis shrinks from 27.92 to 7.13.

Decomposing TFP growth While the above findings already support a narrative of shifting productivity cutoffs, there is yet room for skepticism. It is often the case that firms at the lower end of the productivity scale are small in size and, thus, have negligible effects on the aggregate. A more formal test of growth patterns would therefore consider weighted measures. Moreover, it should aim to disentangle the role of incumbent, entering and exiting firms in shaping TFP changes.

Define aggregate productivity, Z_t, as a weighted average of firm-level TFP. Given that the focus is on firm dynamics, I express overall aggregate productivity as the weighted sum of the aggregate productivities of incumbents, Z_t^C, entrants, Z_t^N, and exiters, Z_t^X,

$$Z_t \equiv \sum_{i \in N_t} s_{i,t} Z_{i,t} = s_t^C Z_t^C + s_t^N Z_t^N + s_t^E Z_t^E ,$$

where $s_{i,t}$ is the labor share of firm i and N_t the total number of firms in the economy, both at time t. In addition, s_t^j is the total labor share and $Z_t^j \equiv \sum_{i \in j} s_{i,t}^j Z_{i,t}^j$ is the aggregate productivity of firms pertaining to group j, where $j = \{C, N, E\}$.

[33]In the former case, the 5% percentile also shrinks although by a smaller magnitude, 36%.

[34]Refer to Table 1.5 for further details.

Figure 1.3: Average change in log (TFP) by percentile

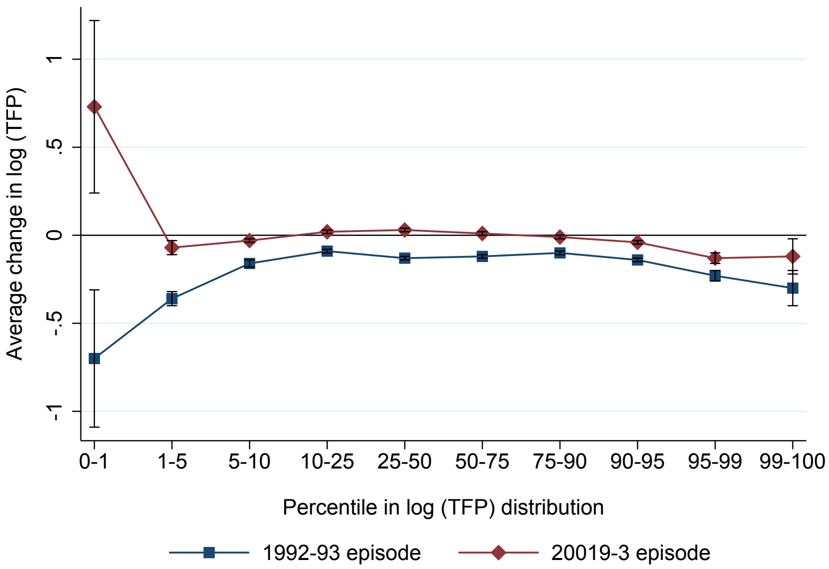

Notes: This graph plots the average change in log (TFP) by percentile of the distribution. It compares the average TFP of firms in a given quantile before and after each of the two sudden stops. As this is an unbalanced panel, firms are allowed to change quantiles and even exit the sample during the transition. The corresponding base and end years are 1991 and 1993 for the first episode; 2009 and 2013 for the second episode. To account for variability, the vertical lines represent error bands.

Source: ESEE data and own calculations.

The variable of interest is the change in aggregate productivity from period $t-1$ to period t, ΔZ_t. It follows that the relevant groups for the analysis are: incumbents in both periods, firms exiting at period $t-1$ and firms entering in period t. This implies that $s^E_{t-1} = s^X_t = 0$. By exploiting the fact that $s^C_{t-1} + s^X_{t-1} = 1$ and $s^C_t + s^N_t = 1$ and using the expression above, I can rewrite the change in aggregate productivity as

$$\Delta Z_t = Z^C_t - Z^C_{t-1} + s^N_t \left(Z^N_t - Z^C_t\right) + s^X_{t-1} \left(Z^X_{t-1} - Z^C_{t-1}\right) .$$

The interpretation of the above decomposition partly coincides with that of Melitz and Polanec (2015): entrants (exiters) contribute positively to TFP growth when their average productivity is higher (lower) than the incumbents' counterpart. These contributions are weighted by the labor share of entrants, s^N_t, and exiters, s^X_{t-1}, respectively.[35] I abstract, however, from decomposing the contribution of incumbents

[35] This version differs from the widely used Foster et al. (2001) decomposition in allowing for differences in the reference productivity for entrants, exiters and incumbents. Intuitively, the contribution of entrants (exiters) is now equal to the change in productivity one would observe if entry (exit) was elided. Moreover, it has a direct mapping into a theoretical model of firm productivity heterogeneity, circumventing the recent criticism to accounting exercises measuring reallocation posed by Hsieh and Klenow (2017). Even so, results are robust to considering Foster et al. (2001)

further using Olley and Pakes (1996)'s approach.[36] Instead, I follow Dias and Marques (2018) in tracking individual incumbent firms over time so that I can distinguish between the contributions of firm-level productivity growth and labor share reallocation among them.

Given the definition of Z_t^C, the change in aggregate productivity can be further decomposed as:

$$\Delta Z_t = \sum_{i \in C} s_{i,t-1} \Delta Z_{i,t} + \sum_{i \in C} Z_{i,t-1} \Delta s_{i,t} + \sum_{i \in C} \Delta s_{i,t} \Delta Z_{i,t} \\ + s_t^N \left(Z_t^N - Z_t^C \right) + s_{t-1}^X \left(Z_{t-1}^X - Z_{t-1}^C \right) .$$

The contribution by incumbents maps exactly into that in Foster et al. (2016). The first term measures the contribution of within-firm productivity changes of incumbents weighted by their initial share. The second term captures the contribution of market share reallocation. The third term is known as the cross-effect, it is the covariance of market share and productivity changes for the individual firm.

Table 1.1: Decomposition of TFP growth

	Episode	
	1992-1993	2009-2013
Productivity growth (%)	-10.87	10.02
Shares of productivity growth		
Incumbent firms share	**-11.20**	**3.05**
Within firm share	-9.69	-2.41
Between firm share	0.47	3.75
Cross-term share	-1.98	1.71
Net entry share	**0.33**	**6.96**
Entrants' share	-0.77	-0.72
Exiters' share	1.10	7.68

Notes: Productivity growth refers to accumulated growth for the considered period. Base and final years are 1991 and 1993 for the first episode; 2009 and 2013 for the second episode. Incumbent and net entry shares add up to productivity growth. Within firm, between firm and cross-term shares add up to incumbent shares. Entrants' and exiters' shares add up to net entry share.

Source: ESEE data and own calculations.

The results of the TFP growth decomposition for the two sudden stops are summarized in Table 1.1. The decline in TFP in the 1992-1993 crisis is entirely driven

and Griliches and Regev (1995) alternative decompositions.

[36] Olley and Pakes (1996) would simply set:

$$Z_t^C - Z_{t-1}^C = \Delta \bar{Z}_t^C + \Delta Cov \left(s_{i,t}^C, Z_{i,t}^C \right) .$$

by incumbents. In fact, net entry contributes to positive growth, although the magnitude is small. Among incumbents, there is some reallocation of market shares towards more productive firms. However, it is far from enough to overcome the pronounced fall in within-firm productivity and the cross-term.[37]

In contrast, the increase in TFP experienced during 2009-2013 is largely driven by net entry, in particular, by unproductive firms exiting the sample. The size of the effect is remarkable, especially given that small and medium firms are under-represented in the sample. Delving deeper into the characteristics of exiting firms shows that during the 2009-2013 episode, firms that exit the market were, on average, bigger in terms of labor market share (7.01% versus 2.78) and three times as unproductive in relation to incumbents (27.16% versus 9.17%) than their 1992-1993 counterparts. Moreover, the annualized exit rate more than doubled from 4.47% to 9.19%.[38]

Back to Table 1.1, the contribution of incumbents, although half as important, is also remarkable. It is still the case that average productivity of incumbents is procyclical, yet the positive effect of the between and cross terms dominate overall. The increase in resource reallocation and a stronger correlation between productivity changes and market share, together with the positive contribution of exiting firms, is consistent with a cleansing effect of the 2009-13 sudden stop which is absent in the 1992-93 episode. The cleansing hypothesis, as discussed by Caballero et al. (1994), argues that crises are periods of accelerated productivity-enhancing reallocations, especially as resources are freed by the exit of unproductive firms. I turn to formally testing the firm-level implications of this interpretation in what follows.

The cleansing hypothesis According to the literature, there is a tight connection between firm exit, input growth and productivity: models of firm dynamics predict that exit is more likely among low productivity firms whereas high productivity firms are expected to grow by more every period. The cleansing hypothesis suggests that recessions accelerate these dynamics. One should therefore observe a stronger correlation between survival, employment and capital growth and productivity levels during crises. To test whether this is the case for the two sudden stop episodes considered, I adjust the empirical specification proposed by Foster et al. (2016) and Dias and Marques (2018) and run the following set of regressions:

$$y_{it} = \lambda + \beta\, tfp_{it} + \delta\, ss^1_t + \gamma\, ss^1_t * tfp_{it} + \mu\, ss^2_t + \theta\, ss^2_t * tfp_{it} + \epsilon_{it}\,,$$

[37]Note that finding procyclical firm-level productivity is not surprising, especially, given that I have no feasible way of controlling for variable capacity utilization.

[38]The corresponding averages for the entire sample are the following: the annualized exit rate is 7.71%, the labor share of exiting firms is 6.43% and the difference in TFP between exiting firms and incumbents is 14.09%.

where y_{it} stands for a set of explanatory variables. It is a dummy variable with value one when a firm reports activity in period t and no activity in period $t+1$ in the exit specification. It is a quantitative variable measuring employment and/or capital growth in the regressions for input growth. The regressor ss_t^1 is a dummy variable for the 1992-93 sudden stop, ss_t^2 is a dummy variable for the 2009-13 sudden stop and tfp_{it} captures the log of firm-level productivity.

For the exit specification, the relationship between survival probability and productivity is expected to be positive and, thus, $\beta < 0$. Under the cleansing hypothesis, this correlation should strengthen during a sudden stop episode and one would anticipate $\gamma < 0$ and $\theta < 0$. Note that the sign of parameters δ and μ provide additional insights regarding the interaction terms. They capture the change in exit rate during the sudden stops that is not correlated with productivity. When positive, it suggests that the increase in exit rates during the crises is disproportionately larger for the least productive firms. For the input growth specification, the exact opposite applies.

Results of these regressions are summarized in Table 1.2. The first column shows the relationship between productivity and the probability of exit. Consistent with earlier findings, firms that exit the market tend to feature lower productivity levels. Focusing on the interaction terms, there is evidence of a cleansing effect only during the second episode. Based on the estimates, 2009-2013 is a period of increasing exit rates, especially among the less productive firms. Note that while the coefficients δ and γ have the correct sign, they are smaller in magnitude than μ and θ and, more importantly, not statistically different from zero.

The second and third columns support further the predictions of the cleansing hypothesis for the 2009-13 episode. First, note that there is a positive impact of productivity on labor growth as predicted by the literature. Of greater interest, this correlation is even higher during the second sudden stop. Together with the negative sign of coefficient μ, there is evidence that high productivity levels somewhat shielded firms from shrinking during the crisis years. The fourth and fifth columns show the capital growth specifications for completeness. Results, however, are uninformative with estimated coefficients displaying no statistical significance.

To further understand the quantitative relevance of my results, Figure 1.4 plots the implied differences in exit probability and labor growth between two firms with productivity level one standard deviation above and one standard deviation below the sectoral mean during normal times, the 1992-93 sudden stop and the 2009-13 sudden stop. The difference in exit rates is 3.7% in the baseline scenario and increases during sudden stops. While the increase is minor during the 1992-93 episode, up to 4.3%, the implied difference almost doubles in the latter case, 7.1%. The magnitudes of the difference for labor growth follow a similar pattern. The baseline gap between a high productivity and low productivity firm is only of 0.9%, increasing to 1.1% during the first sudden stop and up to 2.6% over the second. Note that results for labor growth are robust to considering the subsample of continuing firms.

Table 1.2: Reallocation and TFP

	Exit	Labor growth (continuers & exiters)	Labor growth (continuers only)	Capital growth (continuers & exiters)	Capital growth (continuers only)
	(1)	(2)	(3)	(4)	(5)
constant	0.063***	7.619***	7.769***	7.865	6.663
	(0.002)	(0.291)	(0.276)	(18.446)	(20.579)
tfp_{it}	-0.041***	0.980*	1.060**	-11.489	-13.861
	(0.005)	(0.488)	(0.498)	(12.414)	(14.679)
ss_t^1	0.005	-0.582	-0.842	-8.362	-10.654
	(0.005)	(0.886)	(0.883)	(13.137)	(14.147)
$ss_t^1 * tfp_{it}$	-0.005	0.146	0.087	31.244	34.017
	(0.010)	(1.095)	(1.203)	(19.993)	(22.383)
ss_t^2	0.023***	-7.115***	-6.811***	44.912	51.477
	(0.005)	(0.813)	(0.800)	(57.399)	(65.758)
$ss_t^2 * tfp_{it}$	-0.031***	1.637**	1.804**	-22.326	-29.809
	(0.008)	(0.737)	(0.815)	(34.927)	(45.192)
Observations	34,854	30,861	28,275	30,861	28,275
Industry FE	Yes	Yes	Yes	Yes	Yes

Notes: Regression for exit is a linear probability model where exit=1 if the firm reports positive activity in period t and no activity in period $t+1$. Employment and capital growth are measured from period $t-1$ to period t. tfp_{it} is the log firm-level TFP, ss_t^1 is a dummy equal to one for years 1992-1993 and ss_t^2 is a dummy equal to one for years 2009-2013. Standard errors (in parentheses) are clustered by industry; $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Allocative efficiency Finally, I evaluate an additional theoretical channel through which reallocation may contribute to TFP growth - increased allocative efficiency. Consider the Hsieh and Klenow (2009) framework with a final good featuring a CES production function in differentiated intermediates goods that are imperfectly substitutable. Intermediate good producers have standard Cobb-Douglas production technologies, with capital share α, and are subject to firm-specific exogenous wedges that distort (i) output, τ_{it}^y, and (ii) capital relative to labor, τ_{it}^k. The individual intermediate good producer optimization problem delivers the following first-order conditions with respect to labor, l_{it}, and capital, k_{it}:

$$MRPL_{it} = \left(\frac{1-\alpha}{\mu}\right)\left(\frac{p_{it}y_{it}}{l_{it}}\right) = \left(\frac{1}{1-\tau_{it}^y}\right) w_t\ ,$$
$$MRPK_{it} = \left(\frac{\alpha}{\mu}\right)\left(\frac{p_{it}y_{it}}{k_{it}}\right) = \left(\frac{1+\tau_{it}^k}{1-\tau_{it}^y}\right) r_t\ ,$$

where $p_{it}y_{it}$ is firm nominal value added, w_t is the cost of labor, r_t is the cost of capital and μ is the constant markup of price over marginal cost.

Hsieh and Klenow (2009) formally show that aggregate TFP in this economy is highest when resources are allocated optimally. This is achieved only if firms face equal distortions and marginal revenue products above are equalized.[39] Therefore, the degree of dispersion in firm-specific distortions is informative of the degree of misallocation in the economy. As distortions are unobservable in practice, I measure the standard deviation of marginal revenue products as a proxy of allocative efficiency and study its evolution over time. Periods of higher TFP should be associated with

[39] See Appendix 1.11.2 for a brief review of their argument.

Figure 1.4: Differences between high and low productivity firms

(a) Exit rate

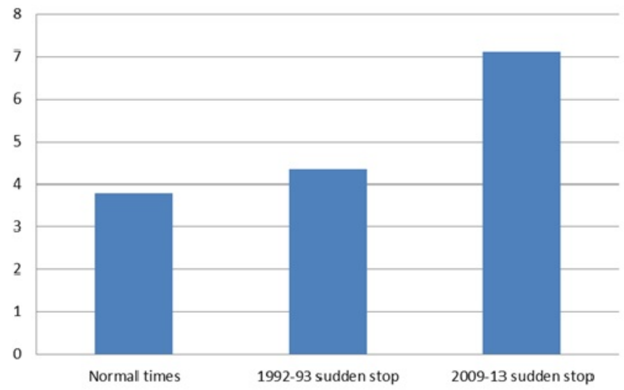

(b) Labor growth (continuers & exiters)

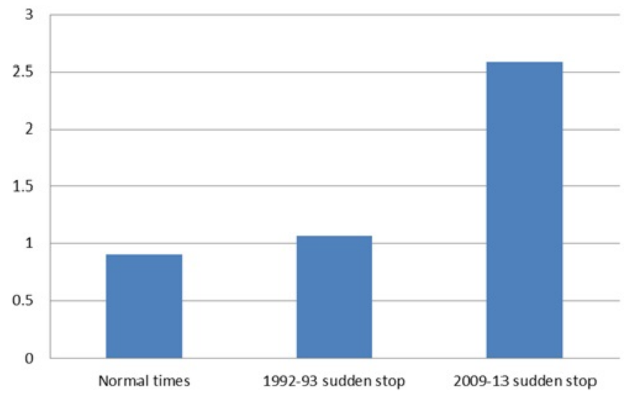

(b) Labor growth (only continuers)

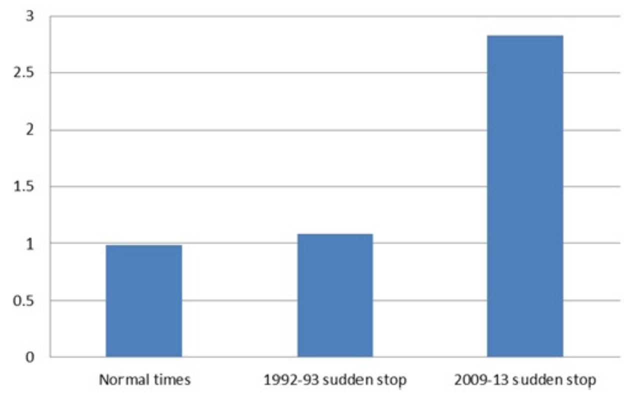

Notes: This figure depicts the predicted difference in probability of exit (panel A, low minus high) and the predicted difference in labor growth rate (panels B and C, high minus low) between a firm one standard deviation above the sectoral mean and a firm one standard deviation below the sectoral mean. Figures are computed from models estimated in Table 1.2.

periods of lower marginal revenue product dispersion and differences in the results for capital and labor can be interpreted as evidence of the different types of wedges that prevail.

I set the capital share to be equal to 0.35 and the constant markup equal to 1.5 as in Gopinath et al. (2017). I first obtain sector-level measures of dispersion in logs which I then aggregate into an economy-wide labor-weighted average using time-invariant weights corresponding to the 2000-2014 labor share average.

Figure 1.5 reports the within-sector standard deviations of marginal revenue products of capital and labor relative to 1990, which is normalized to one. The dispersion of log MRPK is declining over time until the late 1990s when the trend clearly reverses. During the 2000s there is a gradual increase in dispersion, with the more pronounced hikes taking place from 2005 onwards. This is somewhat interrupted during the recent crisis during which dispersion is reduce slightly with the trend reverting back to the pre-crisis level by the end of the sample. The overall description holds for the dispersion of log MRPL too, although the latter depicts much larger volatility.

The increase in the dispersion of log MRPK during the pre-crisis had already been documented by both Gopinath et al. (2017) and García-Santana et al. (2016) using different datasets for the Spanish manufacturing sector. This result justifies their focus on the role of capital misallocation. The difference here is the pattern of log MRPL; while the former papers had reported a relatively flat (or even declining) path, I find that it follows a similar, yet more pronounced, trend to that of log MRPK. According to the Hsieh and Klenow (2009) framework, this should be interpreted as evidence of changing external distortions that affect both factors of production.[40] I argue, however, that internal distortions, such as heterogeneity in price markups, would generate observationally equivalent patterns. Moreover, this is a more realistic interpretation given that the constant markup assumption has been long rebated by the industrial organization literature (see Syverson (2004) and De Loecker and Warzynski (2012) as examples). [41]

In sum, the above findings call for a theory of sudden stops that features heterogeneously productive firms, selection into production and endogenous variables markups. All of these elements, together with the exchange rate dimension, are featured in the theoretical model that I develop in the next section.

1.5 Theoretical model

Consider an infinite-horizon small open economy. Time is discrete and indexed by t. The economy is populated by a representative household that consumes goods

[40] In the model this is represented by an output distortion.

[41] See Fernández et al. (2015) for evidence on the evolution of Spanish price-cost markups for the 1995-2012 period.

Figure 1.5: Evolution of allocative efficiency measures

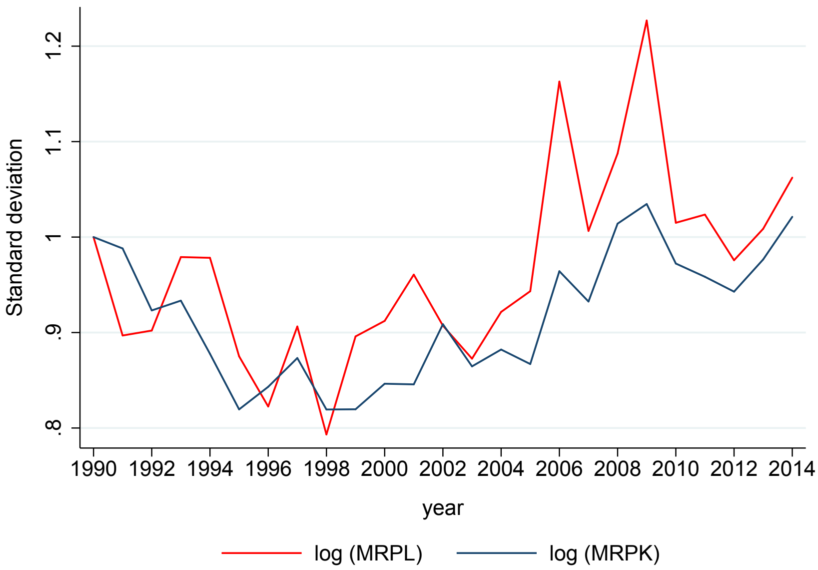

Notes: This figure plots the within-industry dispersion of the marginal revenue products of capital and labor over time. The numbers depicted are relative to 1990, which is normalized to one. Marginal revenue products are measured at the firm-level according to the Hsieh and Klenow (2009) framework. Standard deviations at the sector level are aggregated using time-invariant labor weights.

Source: ESEE data and own calculations.

and leisure and engages in financial transactions with foreign investors. There is also a large number of differentiated firms that produce consumption goods using labor supplied by the households, and a monetary authority that sets the nominal exchange rate as the policy instrument.

1.5.1 A representative household

The representative household derives utility from leisure and the consumption of a set of differentiated goods, indexed by ω, and supplies differentiated types of labor input, indexed by i. The lifetime utility is given by:

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t U \left(q_t(\omega), L_t^i \right) \right], \tag{1.1}$$

where $\mathbb{E}_t$ is the expectation operator conditional on the information set available at time t, β is the discount factor, $q_t(\omega)$ is the consumption level of variety ω and L_t^i is the labor supply of type i. The period utility function is assumed to be:

$$U \left(q_t(\omega), L_t^i \right) = \alpha \int^{N_t} q_t(\omega) \, d\omega - \frac{1}{2} \gamma \int^{N_t} q_t(\omega)^2 \, d\omega - \frac{1}{2} \eta \left(\int^{N_t} q_t(\omega) \, d\omega \right)^2 - \int_0^1 L_t^i \, di \,,$$

where N_t is the number of differentiated varieties available in the economy.

Melitz and Ottaviano (2008) preferences are appealing for three reasons. First, they capture love of variety through γ, which determines the level of product differentiation between consumption goods and is assumed to be strictly positive. As γ increases, consumers place higher weight on the distribution of consumption across varieties. Second, the quadratic form gives rise to a linear demand function which ensures the existence of a choke price and an extensive margin of production even in the absence of fixed costs of production. Third, they generate endogenous variable markups, which capture the effect of market competition on firm sales (the so-called pro-competitive effect) as opposed to standard CES preferences.

Melitz and Ottaviano (2008) preferences also depict a second consumption good, which is homogeneous and assumed to be the numeraire, with a linear production technology that pins down the wage in the economy. As endogenous fluctuations in the wage level are relevant in this analysis, this feature of the original functional form is inconvenient. Moreover, in the context of an internal devaluation, it is also interesting to capture any changes in demand patterns that may arise from movements in wages. My approach is to explicitly model the labor supply decision by assuming preferences that are linear in leisure.[42] The demand parameters α and η therefore measure the substitutability between the consumption of differentiated goods and leisure and are also assumed to be strictly positive.

The budget constraint of the representative agent in terms of domestic currency can be written as:

$$\int^{N_t} p_t(\omega)\, q_t(\omega)\, d\omega + \epsilon_t B_t = \int_0^1 W_t^i L_t^i \; di + \Pi_t + \epsilon_t R_{t-1} B_{t-1}\,, \tag{1.2}$$

where $W_t^i L_t^i$ is the income derived from supplying differentiated labor input i, Π_t is profit received from firms and ϵ_t denotes the nominal exchange rate, defined as units of domestic currency needed to buy one unit of foreign currency.

The representative household can only engage in financial transactions with foreign investors by trading in risk-free foreign denominated bonds B_t, which pay a debt elastic rate of return:

$$R_t = R_t^* + \phi\left(e^{\bar{B}-B_t} - 1\right) + \left(e^{\xi_t - 1} - 1\right)\,, \tag{1.3}$$

where R_t^* is the world interest rate and $\bar{B}$ is the steady state level of debt.[43,44] The only source of uncertainty is ξ_t, which is interpreted as a country risk premium shock,

[42]Given the quasi-linear functional form, there is no income effect for differentiated varieties. However, changes in wages will affect demand through the substitution effect.

[43] This debt-elasticity of the interest rate is assumed to ensure a stationary solution to the model after detrending following Schmitt-Grohé and Uribe (2003).

[44]Households are not allowed to trade in domestic bonds in the baseline model for the sake of simplicity. However, extending the model to include domestic bonds would be trivial as these would be in zero net supply.

similar to that of Garcia-Cicco et al. (2010) and Drechsel and Tenreyro (2017), and assumed to follow an AR(1) process in logs. A sudden stop in the model is a positive realization of ξ_t: an unexpected increase in the cost of international borrowing that forces the domestic economy to deleverage internationally by expanding net exports.

Labor supply is differentiated: there is a unit continuum of labor types which are imperfect substitutes between them. Firms can aggregate labor types according to $L_t = (\int_0^1 L_t^{i\,\frac{\theta-1}{\theta}} di)^{\frac{\theta}{\theta-1}}$, where θ measures the elasticity of substitution. I assume that the representative household supplies all the differentiated labor inputs as in Woodford (2011).[45] Suppose, for example, that each member of the household specializes in one occupation. The representative household has monopoly power to set the wage for each labor type, W_t^i.

Each period the household chooses $q_t(\omega), B_t, L_t^i$ and W_t^i to maximize the expected present discounted value of utility, equation (1.1), subject to the budget constraint, equation (1.2), and the demand for type i labor input, which is given by:

$$L_t^i = \left(\frac{W_t}{W_t^i}\right)^{\theta} L_t \,.$$

Optimality conditions Given quadratic preferences, it may be the case that not all differentiated goods are demanded by the household. However, when a particular good ω is consumed, its inverse demand is determined by:

$$\alpha - \gamma q_t(\omega) - \eta Q_t = \lambda_t p_t(\omega) \ , \tag{1.4}$$

where Q_t is the consumption level over all varieties and λ_t is the time t Lagrangian multiplier. Consumption of a given variety decreases with price, the marginal utility of wealth and total consumption.

The optimal decision for the purchase of the foreign asset, B_t, delivers a standard Euler equation:

$$\lambda_t = \beta R_t E_t \left[\frac{\epsilon_{t+1}}{\epsilon_t}\lambda_{t+1}\right] \,. \tag{1.5}$$

A higher interest rate and expectations of nominal exchange rate depreciation both increase the returns from foreign investment and, thus, encourage consumer savings.

Solving for the optimal wage for labor type i gives:

$$W_t^i = \frac{\theta}{\theta - 1}\frac{1}{\lambda_t} \,. \tag{1.6}$$

Intuitively, higher wages increase household's wealth everything else equal. Given diminishing marginal utility, the Lagrangian multiplier falls. Equation (1.6) also implies that the optimal flexible wage is equalized across labor types *i.e.* $W_t = W_t^i$.

[45] This is equivalent to assuming that each household specializes in the supply of one type of labor input as long as there are equal number of households supplying each type.

Finally, note that the representative household will be willing to satisfy firms' labor demand as long as the real wage covers the marginal rate of substitution between consumption and leisure:

$$\frac{W_t}{P_t} \geq \frac{1}{(\alpha - \eta Q_t) N_t - \gamma Q_t} .$$

1.5.2 Firms

There is a continuum of measure M of domestic firms, each choosing to produce a differentiated variety ω.[46] Labor is the only factor of production and the unit production cost is a concave function in the factor price *i.e.* $C_t = W_t^{\sigma}$, where $0 < \sigma \leq 1$ is the labor income share.[47] Firms only differ in the productivity level z which is drawn from a Pareto distribution $1 - G(z) = \left(\frac{1}{z}\right)^k$ with shape parameter k and minimum productivity level equal to one.

The main focus of the paper is the short-run and, as such, cross-country reallocation of firms is not allowed. [48] This implies that the number of potentially active firms in the economy, M, is fixed and there is no free entry condition. Firms only choose whether to produce or not in each period based on the profitability for the corresponding period.

Firms can sell their varieties in both the domestic and the export market. Markets are segmented and selling abroad requires incurring a per-unit trade cost $\tau > 1$. While domestic demand for variety z, $q_t^H(z)$, is given by equation (1.4), the foreign demand for a domestic variety z, $q_t^{*F}(z)$, is given by:

$$q_t^{*F}(z) = A - B p_t^{*F}(z) , \tag{1.7}$$

where A and B are exogenous given a small-open economy setting. In the spirit of Demidova and Rodriguez-Clare (2009), Appendix 1.11.3 shows that this small open economy is a special case of the two economy framework where the share of potentially active firms in Home, $n = \frac{M}{M+M^*}$ approaches zero.[49]

[46] The same is true for the foreign economy: there is a continuum of measure M^* potentially active foreign firms.

[47] To rationalize this functional form, suppose there is a second factor of production, land, which is inelastically supplied by households and the production function is Cobb-Douglas in land and labor. If the rental price of land is assumed to be constant, the unit production cost is given by $C_t = \left(\frac{W_t}{\sigma}\right)^{\sigma} \left(\frac{\kappa}{1-\sigma}\right)^{1-\sigma}$.

[48] Note that this is only true for the baseline set-up. In one of the extensions, I allow for firm entry and exit and study long-run implications instead.

[49] In the limit z^{*F} is unaffected by changes in Home, the term A includes the price index, the number of consumed varieties and the marginal utility of wealth in Foreign while the term B is proportional to the marginal utility of wealth in Foreign.

Optimality conditions The profit maximization problem delivers the following set of first-order conditions:

$$q_t^H(z) = \max\left\{\frac{\lambda_t}{\gamma}\left[p_t^H(z) - \frac{W_t^\sigma}{z}\right], 0\right\},$$
$$q_t^F(z) = \max\left\{\frac{\lambda_t}{\gamma}\left[p_t^F - \frac{\tau\epsilon_t\,(W_t^*)^\sigma}{z}\right], 0\right\},$$
$$q_t^{*F}(z) = \max\left\{B\left[p_t^{*F}(z) - \frac{\tau W_t^\sigma}{\epsilon_t z}\right], 0\right\},$$

where the expressions for domestically-consumed domestically-produced, henceforth domestic goods, $q_t^H(z)$, and exported goods, $q_t^{*F}(z)$, are given by the optimization of domestic firms while the expression for imported goods, $q_t^F(z)$, results from the optimization of foreign firms. Note that the corresponding prices are also derived from the above expressions.

The labor demand for a domestic firm with productivity level z is given by:

$$L_t(z) = \frac{\sigma}{W_t^{1-\sigma}}\frac{q_t(z)}{z}, \tag{1.8}$$

where $q_t(z)$ will be either $q_t^H(z)$ or $q_t^{*F}(z)$ depending on whether the labor input hired will be used to serve the domestic or the export market.

1.5.3 Aggregation and market clearing

I aggregate firm-level variables and impose market clearing conditions as the building blocks to define the competitive equilibrium.

Productivity thresholds Given that firm-level productivity follows a Pareto distribution, the aggregate productivity level for a given market is summarized by a productivity threshold.[50] This is simply the productivity level of the marginal firm that is indifferent between producing or not for a specific market.

On the supply side, the zero profit condition holds for the marginal firm: it optimally sets its price equal to its marginal cost. On the demand side, the linearity of consumer's demand gives rise to the existence of a choke price. This is the maximum price that can be charged for a given variety at which demand is driven down to zero. By combining these two conditions, the equilibrium thresholds can be expressed as:

$$z_t^H = \frac{\gamma + \eta N_t}{\alpha\gamma\frac{1}{\lambda_t} + \eta P_t} W_t^\sigma, \tag{1.9}$$
$$z_t^F = \frac{\gamma + \eta N_t}{\alpha\gamma\frac{1}{\lambda_t} + \eta P_t} \tau\epsilon_t (W_t^*)^\sigma, \tag{1.10}$$
$$z_t^{*F} = \frac{B}{A}\frac{\tau W_t^\sigma}{\epsilon_t}, \tag{1.11}$$

[50] See Section 1.6 for the formal proof.

where z_t^H is the productivity threshold for domestic firms serving the domestic market, z_t^F is the importer threshold and z_t^{*F} is the exporter threshold. Given the small open economy set-up, the productivity threshold for foreign firms serving the foreign market, z_t^{*H} is exogenously determined and it is irrelevant for the analysis.

Number of firms The number of active firms in the domestic market, N_t is the sum of domestic firms that serve the domestic market, N_t^H, plus the number of foreign importers, N_t^F. Given the number of existing firms in both markets, M and M^*, and the Pareto distribution assumption, the number of active firms is given by:

$$N_t = M\left(\frac{1}{z_t^H}\right)^k + M^*\left(\frac{1}{z_t^F}\right)^k , \tag{1.12}$$

where $\left(\frac{1}{z^H}\right)^k$ is the probability that an incumbent has a productivity level above the cutoff and, thus, generates positive profits. Note that because each firm specializes in a particular variety, N_t is also the number of differentiated varieties available for consumption in the small open economy.

Price level The aggregate price level is given by the sum of prices of all goods consumed domestically, that is, prices of domestically produced goods consumed domestically and import prices:

$$P_t = N_t^H \int_{z_t^H} p_t^H(z)\frac{g(z)}{1-G(z_t^H)}\ dz + N_t^F \int_{z_t^F} p_t^F(z)\frac{g(z)}{1-G(z_t^F)}\ dz\,.$$

Combined with the optimal price expressions that result from the firm's maximization problem and the number of active firms in equilibrium, given by equation (1.12), the above expression is considerably simplified to read:

$$P_t = \frac{2k+1}{2k+2}\frac{W_t^\sigma N_t}{z_t^H}\,. \tag{1.13}$$

The aggregate price level is determined by the number of firms and the average effective marginal cost. The former follows by definition, the latter from the individual firm's optimization problem. Firms charge higher prices whenever their cost of production increase. This is the case when the wage level is high but also when the individual productivity level is low. As the average productivity level in the economy depends positively on the domestic threshold, the aggregate price level decreases in z_t^H.

Wage level I introduce nominal rigidities in the form of sticky information, as in Mankiw and Reis (2002), in the wage setting process. The representative household updates its information set for each labor type it supplies with a probability μ. The aggregate wage is then given by:

$$log\, W_t = \mu \sum_{s=0}^{\infty} (1-\mu)^s \mathbb{E}_{t-s} \left\{ log\, W_t^i \right\} = log\, \frac{\theta}{\theta - 1} - \mu \sum_{s=0}^{\infty} (1-\mu)^s \mathbb{E}_{t-s} \left\{ log\, \lambda_t \right\} . \tag{1.14}$$

A labor type that last updated its information set s periods ago chooses its wage today to be equal to its s-periods-ago expectation of today's flexible wage. Thus, the aggregate wage is a weighted average of the current and all past expectations of today's desired wage. Expectations farther in the past are given less weight because the share of labor types that are stuck with old information decays over time.

Labor market clearing To ensure that the labor market clears in equilibrium, aggregate labor demand must equal aggregate labor supply. To aggregate domestic individual labor demand given by equation (1.8), I sum across all active domestic firms using the Pareto distribution assumption. Labor market clearing then boils down to:

$$L_t = \frac{k}{(k+1)(k+2)} \frac{\sigma}{W_t^{1-2\sigma}} M \left[\frac{\lambda_t}{\gamma} \left(z_t^H \right)^{-(k+2)} + B \frac{\tau^2}{\epsilon_t} \left(z_t^{*F} \right)^{-(k+2)} \right] . \tag{1.15}$$

The balance of payments condition Combining some of the equilibrium conditions above, together with the domestic firms' aggregate profit equation and the consumer's budget constraint gives the aggregate resource constraint of the economy, which, in an open-economy setting, is simply the balance of payments condition. In other words, it states that the current account must be equal to the capital account in equilibrium:

$$EX_t - IM_t = \epsilon_t \left(B_t - R_t B_{t-1} \right) , \tag{1.16}$$

where EM_t and IM_t, the total export and import revenues in domestic currency terms, are given by:

$$IM_t = \int^{N_t^F} p_t^F(\omega)\, q_t^F(\omega)\, d\omega = \frac{1}{k+2} M^* \frac{\lambda_t}{\gamma} \frac{\left(\tau \epsilon_t \left(W_t^* \right)^{\sigma} \right)^2}{2} \left(z_t^F \right)^{-(k+2)} , \tag{1.17}$$

$$EX_t = \epsilon_t \int^{N_t^{*F}} p_t^{*F}(\omega)\, q_t^{*F}(\omega)\, d\omega = \frac{1}{k+2} M \frac{B}{2} \frac{\left(\tau W_t^{\sigma} \right)^2}{\epsilon_t} \left(z_t^{*F} \right)^{-(k+2)} . \tag{1.18}$$

1.5.4 Exchange rate policy

To pin down the nominal variables of the model, I need to determine exchange rate policy. Suppose the central bank implements monetary policy by setting the nominal exchange rate. I consider two exchange rate regimes characterized by different targeting rules. First, consider a currency union. This is equivalent to assuming that the central bank can perfectly commit to a currency peg in which $\epsilon_t = 1$ at every period t.

Second, assume a policy of strict zero wage inflation targeting. This rule simply offsets all the distortions originating from nominal rigidities in the economy by implementing the flexible wage equilibrium, which is given by equation (1.6). Any movements in the real exchange rate will translate one-to-one into movements in the nominal exchange rate. This is the equivalent to a floating arrangement in this framework.[51]

1.5.5 Equilibrium

I am now ready to define a rational expectations equilibrium as a set of stochastic processes $\{z_t^H, z_t^F, z_t^{*F}, IM_t, EX_t, L_t, N_t, B_t, R_t, P_t, \lambda_t, W_t\}_{t=0}^{\infty}$ satisfying equations (1.3), (1.5), and (1.9)-(1.18) given the exogenous process $\{\xi_t\}_{t=0}^{\infty}$ and the central bank's policy $\{\epsilon_t\}_{t=0}^{\infty}$. The foreign wage, W_t^*, is normalized to one.

Appendix 1.11.3 discusses the existence and uniqueness of the non-stochastic steady state.

1.6 Sudden stops and productivity

Before proceeding to the full characterization of the model's solution, it is useful to build some intuition on the potential impact of a sudden stop on productivity. A sudden stop episode forces a real exchange rate depreciation in the domestic economy. This implies a nominal exchange rate depreciation, an internal devaluation or a combination of both depending on the exchange rate policy in place. To ease understanding, I redefine a sudden stop directly as either an exogenous fall in wages, if in the currency union regime, or as an exogenous increase in the nominal exchange rate, if in the floating arrangement. In other words, I disregard the balance of payment condition, given by equation (1.16), for this section and study a version of the model with a partial equilibrium flavour.

1.6.1 Aggregate productivity

The variable of interest is domestic aggregate productivity, which is given by:

$$Z_t^H = N_t^H \int_{z_t^H}^{\infty} \Omega(z)\, z\, \frac{g(z)}{1 - G(z_t^H)}\, dz\,,$$

[51] The exchange rate policy defined here can be easily generalized by assuming a rule such that:

$$(\Pi_t^w)^{\phi_W} (\epsilon_t)^{1-\phi_W} = 1\,, \tag{1.19}$$

where $\Pi_t^w = \frac{W_t}{W_{t-1}}$ is wage inflation and $0 \leq \phi_w \leq 1$ is the weight that the monetary authority puts on wage stabilization. A currency union and a strict wage inflation target are the two extreme versions of this rule, with ϕ_w set equal to zero and one respectively.

where $\Omega(z)$ is the weight used in the aggregation. It must satisfy:

$$N_t^H \int_{z_t^H}^{\infty} \Omega(z)\ \frac{g(z)}{1-G(z_t^H)}\ dz = 1\,.$$

Aggregate productivity is normally computed as: (i) the unweighted average, $\Omega(z) = \frac{1}{N_t^H}$; (ii) the output-weighted average, $\Omega(z) = \frac{q(z)}{Q_t^H}$; or (iii) the revenue-weighted average, $\Omega(z) = \frac{r(z)}{R_t^H}$. [52] The following Lemma establishes that z_t^H is the key statistic for measuring aggregate productivity independent of the weights used in the aggregation.

Lemma 1. *Domestic aggregate productivity, Z_t^H, is an increasing function of the domestic productivity threshold, z_t^H.*

Proof. See Appendix 1.11.4 □

In other words, changes in productivity in this model are governed by firms' entry and exit dynamics. This is in contrast to alternatives in the literature that either model productivity as an exogenous shock to the economy, allow for variable capacity utilization or consider R&D decisions.

Note further that, given Lemma 1, the terms (domestic) aggregate productivity and (domestic) productivity threshold, z_t^H, are used interchangeably for the rest of this section.

1.6.2 Pro-competitive, cost and demand channels

The productivity threshold is determined by the number of firms in the market, the cost of production and the level of consumer demand; all three are potentially subject to change during a sudden stop episode. Let $\hat{X}_t$ define the log deviation of X_t and $\bar{X}$ be its value at steady state.

Proposition 1. *In equilibrium:*

$$\hat{z}_t^H = \underbrace{\frac{1}{2k+2}\frac{\eta}{\alpha\gamma}\frac{\bar{N}\bar{\lambda}\bar{W}^{\sigma}}{\bar{z}^H}\hat{N}_t}_{\textit{Pro-competitive}} + \underbrace{\sigma\hat{W}_t}_{\textit{Cost}} + \underbrace{\hat{\lambda}_t}_{\textit{Demand}}\,.$$

[52] Q_t^H is total domestic output given by:

$$Q_t^H = N_t^H \int_{z_t^H}^{\infty} q(z)\ \frac{g(z)}{1-G(z_t^H)}\ dz\,,$$

and R_t^H is total domestic revenue given by:

$$R_t^H = N_t^H \int_{z_t^H}^{\infty} r(z)\ \frac{g(z)}{1-G(z_t^H)}\ dz\,.$$

Proof. See Appendix 1.11.4 □

The intuition follows next. In the first place, a larger number of active firms in the market, $\hat{N}_t > 0$, implies greater competition. Given the preferences considered, enhanced competition lowers individual firm demand. This forces less productive firms out of the market as profit margins are reduced at every level of productivity. This *pro-competitive effect* was first introduced by Melitz and Ottaviano (2008) that only considers competition in the goods market.

Second, a higher aggregate wage, $\hat{W}_t > 0$, lowers the firm's profit margin by increasing the costs to all firms. Again, a higher productivity level is then required to remain profitable and select into production, therefore, aggregate productivity increases. This is what I denote the *cost effect*, which is the underlying mechanism in the canonical Melitz (2003) model that focuses on competition in the labor market.

Finally, higher aggregate demand from consumers, $\hat{\lambda}_t < 0$, raises individual firm demand at all productivity levels and loosens the minimum productivity requirement. Less productive firms have a higher chance of entering or surviving in the market. This final channel, a novelty of this model, is referred to as the *demand effect.*[53].

1.6.3 Sudden stops and productivity

The following proposition considers the effect of a sudden stop, defined as explained above, on productivity under the two alternative exchange rate regimes.

Proposition 2. *Given a sudden stop,*

1. *In a floating arrangement, only the pro-competitive channel operates and productivity falls:*

$$\hat{N}_t < 0,\ \hat{W}_t = 0 \quad \textit{and} \quad \hat{\lambda}_t = 0 \quad \textit{so that} \quad \hat{z}_t^H < 0\,.$$

2. *In a currency union, all three channel operate and the change in productivity is ambiguous:*

$$\hat{N}_t < 0,\ \hat{W}_t < 0 \quad \textit{and} \quad \hat{\lambda}_t > 0 \quad \textit{so that} \quad \hat{z}_t^H \gtrless 0\,.$$

Proof. See Appendix 1.11.4 □

First, suppose that the nominal exchange rate depreciates one-to-one with the real exchange rate, *i.e.* ϵ_t increases. Under this assumption, the cost and the demand effect are muted as the wage level remains unchanged. There is a fall, however, in the

[53]There is an implicit demand effect in the baseline Melitz (2003) model too. However, the assumption of fixed production costs introduces an additional fixed cost channel (on top of the variable cost channel here considered) that exactly offsets the demand effect.

active number of firms in the domestic economy as the number of importers declines with the loss of competitiveness of foreign firms. There is an unambiguous fall in productivity as a result of this negative pro-competitive effect.

Suppose instead that the aggregate wage adjusts completely: W_t falls while the nominal exchange rate remains unchanged. Under this alternative scenario, the negative pro-competitive effect prevails as there is still a decline in importing firms. The change in wages, in addition, leads to a negative cost effect, production of goods is cheaper, and a negative demand effect, households consume less.[54] In other words, all three channels are operating.

1.6.4 Increasing TFP in a currency union

The change in productivity after a sudden stop is ambiguous in the currency union and depends on parameter values. It is possible, nonetheless, to show under which parameterization, the demand effect dominates and productivity increases.

Corollary 1. *Following a sudden stop in a currency union, a sufficient condition for* $\hat{z}_t^H > 0$ *is that* $1 > \mu\,\sigma(1+k)$.

Proof. See Appendix 1.11.4 □

There are three key parameters for this condition to hold: the share of labor income, σ, the degree of wage rigidities, μ, and the shape parameter of the productivity distribution, k. The share of labor income governs the mapping between the wage level and the unit production cost. As σ increases, labor represents a greater share of the optimal input bundle and falling wages cheapen production costs by more. This reinforces the cost effect of a sudden stop. In the Melitz (2003) model, the cost channel is at its strongest featuring a linear production function which is linear in labor, $\sigma = 1$.

The degree of wage rigidities determines the size of the demand effect. A sudden stop here is simply an improvement in the domestic economy's competitiveness through an exogenous decline in the wage level. As the level of wage stickiness increases and fewer labor-types are allowed to adjust, the decline in labor-specific wages, W_t^i, that is required to achieve the desired overall wage adjustment is larger. This leads to a larger decrease in today's consumer wealth and, thus, a stronger demand effect of a sudden stop.

The shape parameter measures the concentration of firms at the lower end of the productivity distribution. This represents the inverse of the dispersion in firm-level productivity. As firms only differ in their productivity levels, if k increases, they become more homogeneous. This strengthens the pro-competitive channel by tightening the link between the number of firms and the degree of market power among

[54] Recall that a negative demand effect is represented by a positive change in $\hat{\lambda}_t$.

domestic competitors. The less unique a firm is, the more increased competition will lower its individual demand.

Two questions remain unanswered. First, is the above requirement satisfied under reasonable parameterization? Second, do these results hold in the fully fledged model? In what follows I discuss how to calibrate and solve for the general equilibrium version of the model.

1.7 Taking the model to the data

As the model cannot be solved analytically, I next explore its properties by generating impulse response functions, focusing exclusively on a risk premium shock and studying the role of specific parameters in shaping the TFP result.However, to take the model to the data, I first need to redefine one assumption and calibrate parameters using Spanish data.

1.7.1 Number of existing firms

The baseline model described in Section 1.5 is augmented to better suit the analysis that follows. In particular, the number of existing firms is allowed to vary. While this modification improves the predictive performance of the model in general equilibrium, it does not change the main conclusions derived in the previous sections as shown in full detail in Appendix 1.11.5.

The pool of potentially active firms, M, which is assumed to be constant in the benchmark case, now responds to a sudden stop to circumvent the production boom that the model would otherwise generate.[55] This feature of the baseline model is common to many other papers in the sudden stop literature. Kehoe and Ruhl (2009) show that standard models that abstract from financial frictions are unable to reproduce observed decreases in output. The literature has considered featuring imported intermediate goods, labor frictions, variable capacity utilization, Greenwood et al. (1988) preferences, and exogenous TFP declines. Given the new extensive margin introduced in the model, I instead assume a law of motion for the number of existing firms such that in log deviation terms:

$$\hat{M}_t = -\rho\hat{\epsilon}_t - (1-\rho)\hat{W}_t \,.$$

The interpretation is the following: additional labor, domestic or foreign, is required to set up a new firm and, thus, the pool of potentially active firms depends negatively on the cost of labor input, either the domestic wage, ϵ_t, or the foreign wage in domestic currency units, ϵ_t. The parameter ρ measures the degree of complementarity between domestic and foreign labor in setting up new firms.

[55] Note this is true independently of the exchange rate policy that is implemented.

This assumption captures, in essence, some of the implications of the long-run version of the baseline model. Further details of the full extension together with results are available in Section 1.8.3. In short, the long-run version features a fixed input requirement in the form of capital for the production of any differentiated variety. Capital is produced under perfect competition and accumulated through an investment decision. This follows closely Ottaviano (2012) in putting Melitz and Ottaviano (2008) in a DSGE framework. The only difference is that the fixed input requirement is assumed to be a combination of domestic and foreign capital and, therefore, the number of potentially active firms depends on the price of both types of capital *i.e.* domestic wages and the exchange rate. The law of motion above builds on this relationship, however, it does not take into account the dynamic optimization problem that the long-run version entails.

1.7.2 Calibration

Table 1.3 provides a summary of the parameters of the model, their baseline values and the source or the empirical target. The first set of parameters are standard and, thus, values are set in line with the literature and, when possible, consistent with Spanish statistics taking the 2002-2008 period as a reference. The time period of the model is a quarter. Accordingly, the discount factor β is chosen to be 0.99. The output elasticity parameter σ is set to 0.64, roughly the average labor share and within the range that is common in the literature. For the elasticity of substitution for labor types and the index of wage rigidities, values are taken from Galí and Monacelli (2016) which are based on empirical studies on European countries conducted by the OECD. In terms of trade costs, τ is equal to 1.3 following Melitz and Ghironi (2005) and many others. The steady state level of debt, $\bar{B}$, is assumed to be zero, such that trade is balanced in steady state. Regarding the preference parameters, α, γ and η, I borrow the values used in Ottaviano (2012), all equal to 10.

The ESEE firm-level data presented in Section 2 is then used to estimate the shape parameter of the Pareto distribution, following the approach proposed by Del Gatto et al. (2006). Given the observed cumulative distribution, $G(z)$, I run the following regression for every year and industry:

$$ln\left(1 - G\left(z\right)\right) = \beta_0 + \beta_1 ln\left(z\right) + \eta$$

where, assuming a Pareto distribution, the slope coefficient, β_1 provides a consistent estimator for k. For the 2002-2008 period, k is estimated to be, on average, equal to 1.9, close to Del Gatto et al. (2006)'s result of 2 for a combination of European countries in the year 2000. In addition, the regression R^2, which is equal to 0.7, confirms that the Pareto distribution is a reasonable assumption in this setting.

The above estimation provides an additional coefficient, β_0, that maps one-to-one to the realized distribution's cutoff, $\bar{z}^H$. I use the corresponding 2002-2008 average as

Table 1.3: Calibration

	Parameter	Value	Calibration target/source
β	Discount factor	0.99	Annual real return on bonds is 4%
μ	Index of wage rigidity	0.2	Gali and Monacelli (2016)
θ	Elasticity of substitution (labor)	4.3	Gali and Monacelli (2016)
τ	Iceberg trade cost	1.3	Ghironi and Melitz (2005)
γ	Preference parameter	10	Ottaviano (2012)
α	Preference parameter	10	Ottaviano (2012)
η	Preference parameter	10	Ottaviano (2012)
$\bar{B}$	Steady state level of debt	0	Steady state trade balance
σ	Labor share	0.64	National Accounts Spain
n	Relative size of SOE	0.12	Business Demographic Statistics
k	Shape productivity parameter	1.9	Estimated from ESEE data
A	Foreign demand parameter	0.01	Domestic productivity cutoff (1.55)
B	Foreign demand parameter	0.33	Share of exporting firms (63.6%)
$\bar{M}$	Number of total firms	173	Active domestic firms (75.86)
ϕ	Risk premium parameter	3.2	Output volatility (3%)

a first moment target in two different ways. On the one hand, I combine it with the 2002-2008 average number of firms in the ESEE sample to back up the value of $\bar{M}$ given that the number of potentially active firms is unobservable. The corresponding expression is given by $\bar{M} = \left(\bar{z}^H\right)^k \bar{N}^H$.

On the other hand, I use $\bar{z}^H$ to determine the value of the foreign demand parameters, A and B. To do so I proceed in three steps. First, I set the relative size of the domestic economy, n, to match the 12% share of all Euroarea manufacturing firms that Spanish firms represent according to Eurostat's Business Demography Statistics. Next, I take the average 2002-08 propensity to export as an additional first moment target which combined with $\bar{z}^H$ pins down $\bar{z}^{*F}$ as $\frac{\bar{N}^{*F}}{\bar{N}^H} = \left(\frac{\bar{z}^H}{\bar{z}^{*F}}\right)^k$. Third, I back up the wage level that is consistent with the estimated cutoff using a combination of equilibrium conditions (1.9), (1.10), (1.13) and (1.12) in steady state. Parameter values for A and B then follow naturally using equation (1.11) and the trade balance condition.

The risk premium parameter, ϕ, is a theoretical shortcut to ensure stationarity in small open economy frameworks. In the current setting, it measures the severity of the current account reversal given a one standard deviation shock. Thus, I choose its value such that the second theoretical moment of output during a sudden stop exactly matches its empirical counterpart, 3% for the 2009-2013 period. Finally, as there is no obvious candidate value for the degree of complementarity between foreign and domestic labor in setting up new firms, I consider an intermediate case,

$\rho = 0.5$.

1.7.3 Impulse Responses Functions

Figure 1.6 summarizes the model response of key macroeconomic variables to a sudden stop. All variables, but the current account, are expressed in log deviations from steady state. The current account is expressed in levels as trade balance is assumed to hold before the realization of the shock.

Figure 1.6: Macroeconomic effects of a sudden stop

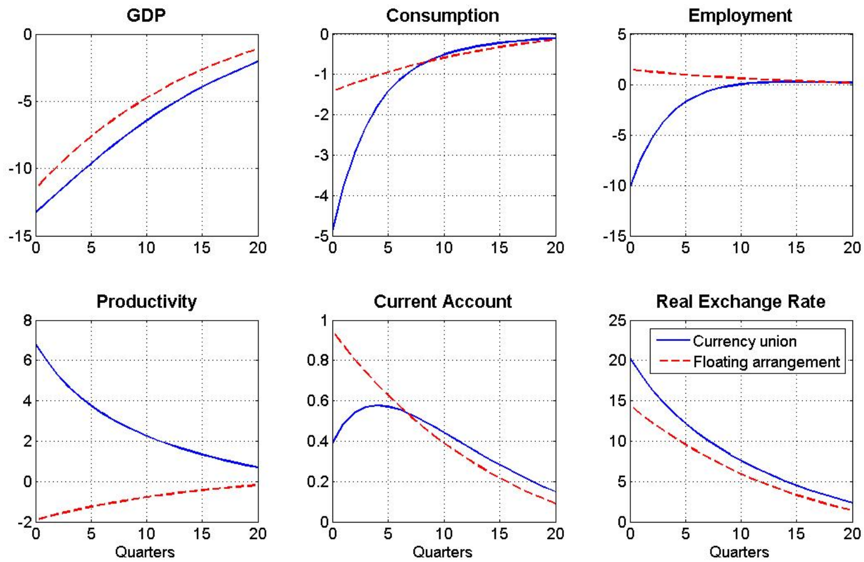

Notes. These figures plot the IRFs of key macroeconomic variables to a one standard deviation shock to the country-specific risk premium. All variables but the current account are expressed in log deviations from steady state. The current account, assumed to be zero in steady state, is expressed in levels.

As expected, a sudden stop is characterized by a depreciation of the real exchange rate and a current account surplus. The model is able to predict a slight delay in the adjustment within a currency union. This is entirely driven by nominal rigidities as the model disregards additional policy instruments available within a currency union, such as public capital inflows, that might directly cushion the adjustment in the data.

The path of TFP diverges across regimes. On the one hand, under the baseline calibration, the negative effect of a lower aggregate demand offsets the positive effect of lower production costs and fewer competing firms on the domestic productivity cutoff and, thus, TFP improves in the currency union. On the other hand, productivity falls unambiguously in the floating regime. I study the sensitivity of these results to alternative parameter values in the following section.

GDP and consumption are both measured in units of foreign currency to ease comparison with Figures 1.1 and 1.2. The model correctly predicts a fall in both variables and under both regimes. Moreover, for a similar GDP decline, the fall in consumption is larger in the currency union, consistent with the aggregate data. The response of employment does not fully match the data: while the event study in Section 1.3 suggests employment in a floating arrangement is unchanged or slightly decreasing, the model predicts a minor increase. On the other hand, there is a clear decline in employment within a currency union both in the model and in the data.

Impulse response functions for all other endogenous variables can be found in Figure 1.10. The current account surplus is explained by a simultaneous increase in export and decline in imports. In the currency union, there is an immediate decline in the price index while wages fall in a staggered fashion. In the floating regime, the exchange rate depreciates on impact with wages and prices remaining unchanged. In both regimes, the number of firms and, thus, the number of varieties falls with the shock.

1.7.4 Sensitivity of the TFP fact

The analytical results of Section 1.6 point to three structural parameters as the main determinants of the overall response of TFP: the degree of wage rigidities, ω, the share of labor income, σ, and the shape parameter of the productivity distribution, k. I next embed the analysis within the general equilibrium framework.

The upper left graph of Figure 1.7 plots the immediate impact of TFP, in log deviations from steady state, for both the currency union and the floating arrangement regimes for $0.1 \leq \omega \leq 0.9$. By definition, under the floating arrangement wages are stabilized completely and, thus, there is no effect of wage frictions whatsoever. For the currency union, nevertheless, higher wage flexibility (higher ω) leads to a smaller increase in TFP.[56]

The upper right graph of Figure 1.7 decomposes the effect of a sudden stop on TFP into the demand, pro-competitive and cost effects for the currency union as defined in Proposition 1.[57] As wages become more flexible, the unit production cost falls by more. The opposite is true for the demand effect: when more labor types are allowed to adjust their wages, the labor-specific wage declines by less and, thus, the required increase in marginal utility of wealth is smaller *i.e.* the positive contribution of negative demand shrinks. The magnitude of the pro-competitive effect also varies with the degree of wage flexibility. The intuition relies on second round effects: as wages fall by more and the increase in productivity cutoff is smaller, the reduction in the number of competitor declines and the pro-competitive effect weakens.

[56]While Figure 1.7 depicts the immediate effect of a sudden stop shock on TFP, conclusions remain true if the cumulative effect on TFP is considered. Results available upon request.

[57]Note that for the floating arrangement it is still the case that only the pro-competitive channel operates.

Panels (b) and (c) perform the same exercise for $0.1 \leq \sigma \leq 0.9$ and $1.5 \leq k \leq 2.5$ correspondingly. On the one hand, as the share of labor falls, the drop in wages that is required to regain international competitiveness increases substantially in a currency union. The greater the fall in wages, the stronger the demand effect and the larger the improvement in TFP. Once again, there is little change in the floating arrangement as the adjustment of the exchange rate is not affected by the production structure of the economy. On the other hand, the shape parameter measures the relative number of low-productivity firms; as it increases, there is a higher concentration of firms in the lower end of the productivity scale. According to the results depicted in Figure 1.7, the behavior of TFP is robust, even in quantitative terms, to different parameterizations of k. Although the size of pro-competitive effect increases with k, as anticipated in Section 1.6, in general equilibrium there is an offsetting increase in the demand effect.

1.8 Extensions

This section briefly introduces a number of extensions to the baseline framework and discusses how the previous results are affected.

1.8.1 A model with capital

The analysis has so far abstracted from the role of capital. This is due to two reasons: first, there is already a number of papers (Reis (2013), García-Santana et al. (2016) and Gopinath et al. (2017)) which have extensively studied the role of physical capital in the context of capital flows. Instead, this paper aims at incorporating an alternative yet complementary explanation to the discussion. Second, the firm-level evidence presented in Section 1.4 is supportive of theories that focus on the composite variable input, and not only on capital. Nonetheless, this extension incorporates explicitly pre-installed physical capital as a second factor of production.

The setting is standard: the production function is Cobb-Douglas in labor, L_t, and pre-installed capital, K_{t-1}. Capital goods are owned by the representative consumer and rented to firms in exchange of a rental rate κ_t. The stock of capital accumulates driven by the investment decision by the representative consumer and the rate at which it depreciates, δ. Appendix 1.11.6 formalizes this extension and provides details on the resulting equilibrium conditions.

Figure 1.11 plots the impulse responses of GDP, consumption, employment, productivity, the current account and the real exchange rate index to a sudden stop shock as defined above. The differences in TFP response across regimes is still noticeable. However, the dynamics depicted do not match those in Figure 1.6 completely. For example, in a currency union the reversal of the current account only holds on impact, net exports fall shortly after and stabilize at a negative level. The rest of

Figure 1.7: The TFP fact - robustness to alternative calibrations

(a) The role of wage rigidities

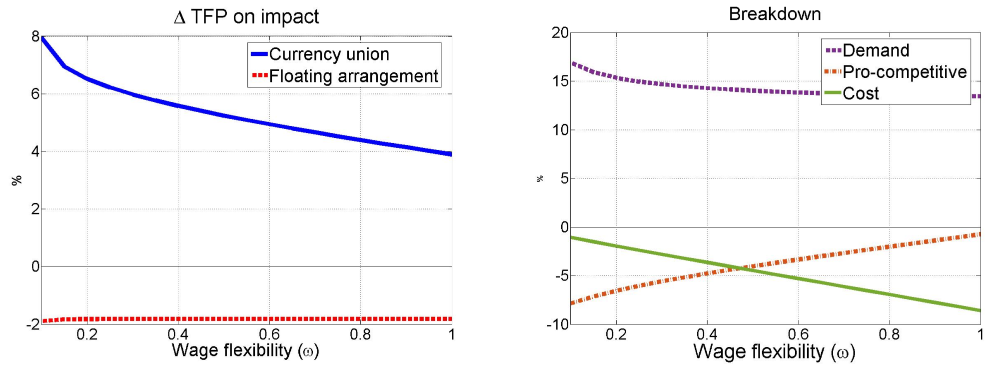

(b) The role of the labor share

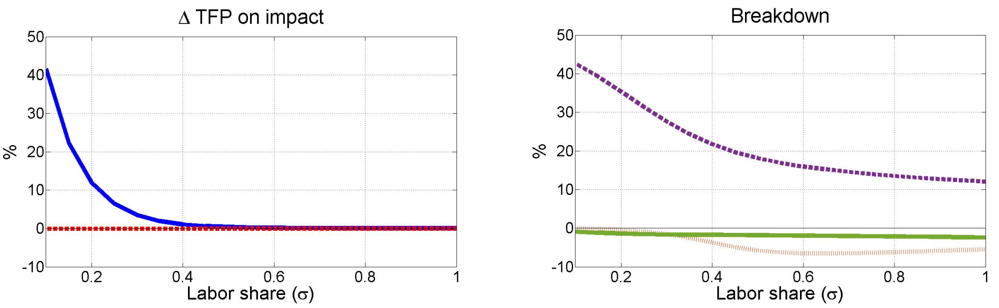

(c) The role of the shape parameter

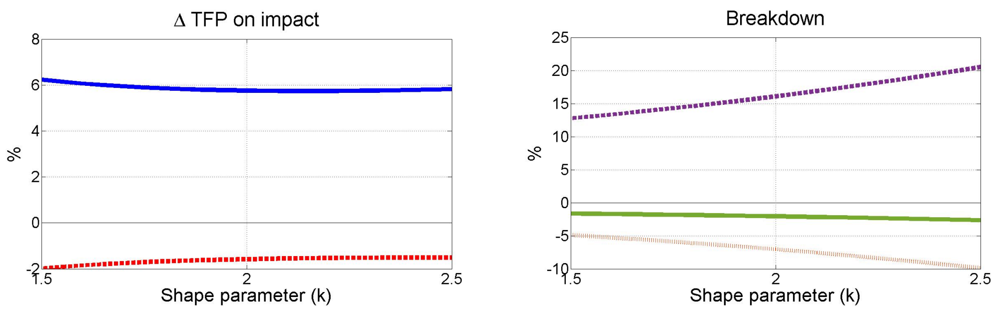

Notes: This figure documents the sensitivity of the TFP fact to different model parameterizations. Panel (a) focuses on different degrees of wage rigidities - higher ω implies lower rigidities. Panel (b) allows for plausible calibrations of the labor share - higher σ implies a larger labor share. Panel (c) explores alternative values of the shape parameter of the Pareto distribution - higher k implies lower dispersion of productivity draws. The first column plots the immediate response of TFP to a one standard deviation shock to the country-specific risk premium in log deviations terms. The second column breaks down the immediate response of TFP into the demand, the pro-competitive and the cost channels as defined by Proposition 1.

macroeconomic variables behave as predicted. In a floating regime, the reversal lasts slightly longer and as long as the current account is positive, the qualitative results of Figure 1.6 hold. For completeness, Figure 1.12 plots the impulse response of the three new variables: investment, price of capital and the unit cost of production.

1.8.2 Extensive versus intensive margin

To account for the decline in firm-level TFP growth observed during the two Spanish sudden stops reported in Section 1.4, I augment the baseline model by assuming that firms' effectiveness in transforming inputs into output depends on (i) an aggregate time-varying component, Z_t; and (ii) a constant firm-specific component, z. The former is represented as an AR(1) process in logs and the latter is drawn from a Pareto distribution as in the baseline setting.

A sudden stop is redefined as a positive realization of the country-specific risk premium, ξ_t, with a simultaneous negative shock to the aggregate component of firm productivity, Z_t. Details of the formalization of this extension are relegated to Appendix 1.11.6. The predicted response of macroeconomic variables is unchanged, at least qualitatively, as depicted by Figure 1.13.

The main difference involves the variable of interest: aggregate domestic productivity. While in the baseline set-up it was sufficient to measure the effect of a sudden stop on the domestic productivity threshold as summarized by Lemma 1; under the new framework, the common component, Z_t, also affects aggregate TFP directly. Figure 1.14 decomposes the effect of a sudden stop on aggregate productivity into the contribution of the productivity threshold (the extensive margin) and that of the common shock (the intensive margin). The right panel fully matches the 2009-13 sudden stop as summarized by Table 1.1. Firm-level TFP is declining over time, however, the exit of unproductive firms completely counteracts the negative effect and productivity increases overall. The left panel is only partly in line with the 1992-93 episode. While Table 1.1 shows a decline in firm-level TFP, the observed contribution of the extensive margin is negligible instead of negative as predicted by the model.

1.8.3 Long-run analysis

This extension studies a long-run version of the baseline model that fully endogenizes the number of existing firms, M_t, in line with Ottaviano (2012). The previous framework is augmented by (i) allowing for investment in capital shares; (ii) introducing a new sector that produces capital; and (iii) imposing a fixed input requirement in terms of capital in the production of differentiated varieties.

In particular, the representative consumer is allowed to buy shares, x_t, of the economy's capital stock, K_t, at price V_t. While capital is assumed to fully depreciate after one period; the investment entitles the representative consumer to a fraction of next period's aggregate firm profit. The consumer budget constraint is correspondingly adjusted to read:

$$\int^{N_t} p_t(\omega)\, q_t(\omega)\, d\omega + \epsilon_t B_t + x_t V_t K_t = \int_0^1 W_t^i L_t^i \; di + x_{t-1}\Pi_t + \epsilon_t R_{t-1} B_{t-1}\,.$$

Capital is supplied under perfect competition by a second sector in the economy. A new unit of capital is produced by combining domestic and foreign units of labor using a Cobb-Douglas production technology: $K_t = \left(l_t^{k,H}\right)^{\rho}\left(l_t^{k,F}\right)^{1-\rho}$. [58] Given the fixed capital requirement, the production of capital determines how many firms will be able to enter the market, $M_t = \frac{K_t}{f_E}$. There is a one-period-time-to-build-lag such that firms that enter at time t, will only be able to produce, provided that they satisfy the corresponding productivity threshold condition, in period $t+1$.

Appendix 1.11.6 describes in greater detail the full equilibrium of this version of the model. It is relevant, however, to highlight one key new optimality condition that emerges from this set-up:

$$M_t = \left(\frac{\rho}{W_t}\right)^{\rho}\left(\frac{1-\rho}{\epsilon_t}\right)^{1-\rho}\beta E_t\left[\frac{\lambda_{t+1}}{\lambda_t}\Pi_{t+1}\right]. \tag{1.20}$$

Intuitively, a lower price of capital encourages investment and increases the number of existing firms. As capital is produced under perfect competition, price is equal to marginal cost and, thus, a function of the price of both types of labor. The price of foreign labor is equal to the foreign wage, which is normalized to one, in domestic currency units *i.e.* the nominal exchange rate. This relationship is captured in reduced form by the law of motion proposed in Section 1.7. In addition, the number of existing firms is also dependent on the discounted expected profits, as profits represent the return on capital investment. This inter-temporal dimension is missing in the previous analysis, however, solving for this long-run version of the model shows that the main conclusions derived above hold.

Figure 1.15 plots the impulse responses of the same variables as the original Figure 1.6, following a sudden stop. The predictions are qualitatively unchanged. The shape of responses is slightly changed because of the delay in adjustment caused by the new timing assumption. The only remarkable difference refers to the relative ordering of GDP and the real exchange rate: in this version, GDP falls by more in the floating arrangement while the real depreciation is less pronounced in the currency union. The opposite is true in the baseline results.

[58] I deviated from Ottaviano (2012) in two ways. First, I introduce foreign labor in the production of capital to ensure a direct role for the nominal exchange rate in firm entry. Second, I consider that while capital fully depreciates, all new units of capital are available for production the following period. The timing is slightly adjusted: investment takes places, next period firms are set-up and capital depreciates.

1.9 Welfare

Once that the model has proved to correctly capture the macroeconomic dynamics that follow a sudden stop episode under the two alternative policy regimes in Section 1.7, it can then be used to conduct normative analysis.

1.9.1 Output loss

I evaluate the performance of exchange rate regimes by comparing the cumulative real output loss resulting from a sudden stop. The reason is two-fold: first, preferences in the model are non-standard, rendering utility-based measures of welfare controversial. Second, output can be easily measured in the empirical data and it is explicitly targeted by policymakers all around the world.

In greater detail, I compute the discounted sum of percentage deviations of realized output from its steady state level following a one standard deviation shock to the country-specific risk premium:

$$\textit{output loss} = \sum_{t=0}^{\infty} \beta^t \left(\frac{Y_t - \bar{Y}}{\bar{Y}} \right),$$

Figure 1.8 plots the ratio of the cumulative output loss of a sudden stop under a currency union to that under a floating arrangement for different degrees of nominal rigidities, with $1 - \omega = 0$ representing a world of perfectly flexible wages. If the ratio is above one, the floating arrangement generates a lower output loss than the currency union and viceversa. Results of the model are summarized by the blue solid line. As can be seen, the desirability of the floating arrangement is increasing in wage stickiness. However, for a wide range of nominal rigidity levels, a currency union performs better in output loss terms.

The second result might feel counter-intuitive since I have argued that in the currency union productivity increases through a welfare-diminishing mechanism: a fall in domestic demand. It is key, however, to understand that a fall in wages and a rise in nominal exchange rates have opposing effects on the number of existing firms, M_t. While the former increases M_t by cheapening domestic labor, the latter decreases M_t by making foreign labor more expensive. In a currency union, the increase in M_t can partly cushion the fall in the number of active domestic firms, $N_t^d = \frac{M_t}{(z_t^H)^k}$, resulting from the increased productivity requirement. If this effect is big enough, the cumulative loss of output is actually smaller than in a floating arrangement as the negative impact of the sudden stop on N_t^d is reduced.

To give a sense of how far the calibrated economy is from the indifference point, where the ratio is exactly one, the shaded area in Figure 1.8 displays the range of values of wage stickiness that have been used by the literature. While the baseline calibration of ω is purposely conservative, $1-\omega = 0.8$, micro-evidence from the ECB

Wage Dynamic Networks has found that the quarterly frequency of unconditional wage adjustments ranges between 20% and 35% *i.e.* $1-\omega \in (0.65, 0.8)$. [59] According to this, which policy regime generates lower output losses after a sudden stop remains an unsettled question.

The Friedman view How do these findings compare to the inherited wisdom on the fixed versus floating debate? The standard case for flexible exchange rates, as first proposed by Friedman (1953), can be summarized in two claims. First, in a world of perfectly flexible wages, a nominal depreciation and an internal devaluation would lead to the same economic outlook. Second, it is the pervasiveness of nominal rigidities, however, that justifies the desirability of floating: when prices do not adjust, quantities do, thus, the lack of wage adjustment in fixed exchange rate regimes leads to suboptimal unemployment and output loss.

To capture the predictions of this traditional view, I modify the baseline model in two ways: shutting down entry and exit firm dynamics and imposing linear production in labor. The former requires featuring a fixed number of homogeneous firms in terms of productivity.[60] The latter consists of setting the labor share, σ, equal to one. The resulting output loss ratio is depicted by the red dashed line in Figure 1.8, an upward-sloping convex curve that starts at exactly the indifference point. Consistent with the two claims above, a floating arrangement always performs better; with the exception of a perfectly flexible wage world, in which the floating arrangement is equivalent to a currency union.

In short, the normative implications of the model match the second claim of the standard argument but disagree regarding the first: it is still the case that increases in wage stickiness are relatively more harmful within currency unions; but accounting for firm dynamics, makes a currency union more desirable if wages are sufficiently flexible.

1.9.2 Other welfare-relevant measures

Findings are robust to considering an alternative performance measure: consumption loss. Figure 1.16 plots the ratio of the cumulative consumption loss of a sudden stop under a currency union v.s. a floating arrangement for different degrees of nominal rigidities. The cumulative consumption loss is computed as explained above. The shape of the two curves remains unchanged: both are increasing and convex with the plot representing the Friedman view starting at exactly one. The main difference, however, lies on the relative steepness. Wage rigidities are increasingly more harmful in terms of consumption relative to output losses for the Friedman view while the

[59] See Druant et al. (2009) and Le Bihan et al. (2012).

[60] A more detailed description of this version of the model together with the corresponding equilibrium conditions can be found in Appendix 1.11.6.

Figure 1.8: Currency union versus floating arrangement

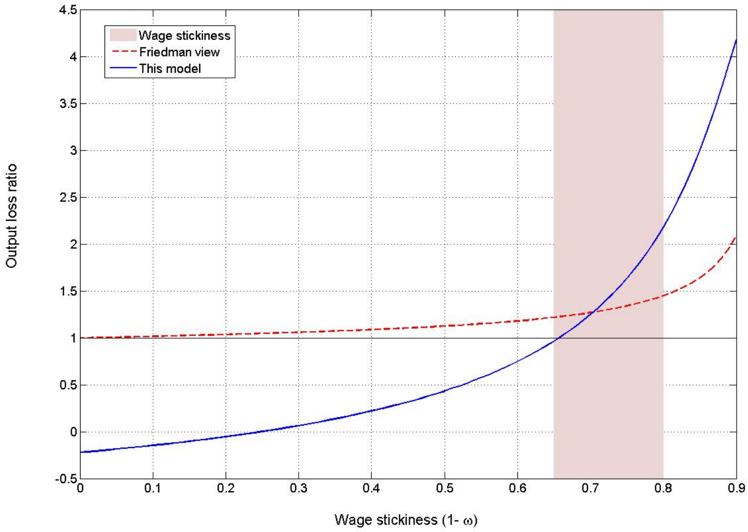

Notes: This figure plots the ratio of the output loss under a currency union to the output loss under a floating arrangement following a sudden stop. The output loss is calculated as the discounted sum of output log deviations after a one standard deviation shock to the country-specific risk premium. The blue solid line refers to the baseline model. The red dashed line refers to a version of the model with no firm dynamics (firms are homogeneous and the number of firms is constant) and linear production in labor ($\sigma = 1$). The shaded area shows the range of plausible values for the wage rigidity parameter as discussed by the literature.

opposite is true for the baseline model. In addition, the indifference point in Figure 1.16 takes place at a higher level of wage stickiness.

Finally, for completeness, I also compute a utility-based measure of welfare as it is standard in the literature. As there is no closed-form representation of the welfare function, I evaluate welfare losses numerically. In particular, I compute the fraction of labor, λ_L, that equates the conditional expectation of future utility along the equilibrium as of time zero to its value in the non-stochastic steady state.[61],. The interpretation is the following: the welfare loss associated to a sudden stop under a given exchange rate regime is the extra amount of labor that the consumer is willing to supply (amount of leisure given up) in steady state to remain indifferent between a world with and without a sudden stop episode. λ_L is implicitly given by the expression:

[61] As in Schmitt-Grohé and Uribe (2007), I consider the conditional rather than the unconditional expectation because different policy regimes tend to have different stochastic steady states. Note that although this strategy computes the constrained policy rule associated with a particular initial state of the economy, this is precisely the state that is of interest for my analysis: the non-stochastic steady state.

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t U \left(q_t^r \left(\omega \right), L_t^r z \right) \right] = U \left(q \left(\omega \right), \left(1 + \lambda_L \right) L \right) \,, \tag{1.21}$$

where $q_t^r(\omega)$ and L_t^r represent the optimal consumption and labor supply paths associated to a particular exchange rate policy r and the right-hand side term measures utility at the non-stochastic steady state. The left-hand side of equation (1.21) is approximated to second-order.

Figure 1.17 shows the corresponding welfare loss ratio, which follows closely the pattern depicted by Figure 1.16. This should not be surprising provided that consumption enters the utility function and, thus, determines the welfare measure directly.

1.10 Conclusion

This paper revisits a classical question in International Macroeconomics: how does exchange rate policy affect macroeconomic performance after a shock? While the literature has commonly praised the advantages of exchange rate flexibility, it has often overlooked the response of productivity. I study the question anew in the context of a sudden stop, emphasizing the divergence in TFP patterns that emerges across exchange rate regimes in the aggregate data and relating them to observed differences in firm dynamics at the micro-level.

The empirical analysis of the paper delivers two main findings. First, TFP systematically collapses under a flexible exchange rate arrangement while it improves, albeit moderately, within a currency union. Second, the difference in productivity growth is largely explained by the reallocation of resources from unproductive exiting firms to productive survivors. While this cleansing effect is quantitatively noticeable after an internal devaluation, it is absent during a nominal depreciation in the cases considered here.

I develop a model that is able to rationalize these empirical facts by endogenizing productivity and incorporating demand effects. The model features three key elements: firm selection, variable markups and elastic labor supply in a small open economy DSGE setting. When jointly combined, productivity is determined by the number of firms (pro-competitive mechanism), the marginal utility of wealth (demand mechanism) and the unit cost of production (cost mechanism). The effect of a sudden stop on productivity works through the combination of each of these channels and depends directly on the degree of currency appreciation vis-à-vis wage devaluation.

Simulations of the model show how accounting for firm dynamics changes the relative macroeconomic performance of exchange rate regimes after a shock. While increased wage flexibility is still desirable within a currency union, greater firm heterogeneity and lower labor income shares also contribute towards smaller output losses after a sudden stop. Importantly, regimes perform differently, even in the extreme case of perfect flexibility.

1.11 Appendices

1.11.1 Aggregate data appendix

Identifying sudden stops: algorithm

The following algorithm combines elements of Calvo et al. (2004) and Cavallo and Frankel (2008).

- Use IMF Balance of Payment annual data for all available countries in the period 1990-2015.
- Drop (i) small countries - in terms of population (below 1 million inhabitants) and in terms of wealth (below 1 billion USD); (ii) countries with incomplete time series.
- Compute year-to-year changes in the financial account.
- Compute rolling averages and standard deviations of the change in the financial account with a window length equal to ten years. Check that at least 60% of the observations in the window are available, otherwise set to missing.
- Identify reversal episodes as subsequent country-year observations that show reductions in the financial surplus half a standard deviation above the mean change as calculated in the previous step. Classify the first and last country-year observation as the start and end of each episode.
- Filter to keep reversal episodes that contain at least one country-year observation with a reduction in the financial surplus one standard deviation above the mean change.
- Filter again to keep reversal episodes that are accompanied by a fall in GDP per capita during the same year or the year that follows immediately after.
- Filter again to keep reversal episodes that are accompanied by a fall in the current account deficit during the same year or the year that follows immediately after. Surviving episodes are classified as sudden stops.

Note that two further refinements are made. First, one year episodes starting in 2009 are dropped from the final sample as they simply capture the global trade collapse that followed the burst of the 2008 financial crisis instead of a country-specific reversal of capital flows. Second, I collapse adjacent sudden stops into the same episode if the gap among the end of the former and the start of the latter is only one year.

1.11.2 Firm-level data appendix

Production function estimation

This appendix reviews the Ackerberg et al. (2015) correction to the proxy approach to production function estimation. I augment it to account for attrition as first proposed by Olley and Pakes (1996).

Consider the standard model,

$$y_{it} = \alpha + \beta_s^k k_{it} + \beta_s^l l_{it} + \omega_{it} + \epsilon_{it} \,, \tag{1.22}$$

where y_{it} is value added, k_{it} is capital and l_{it} is labor input. ω_{it} is unobserved firm-level TFP and modelled as a Markov chain, $\omega_{it} = g\left(\omega_{it-1}\right) + \xi_t$.

Under the assumptions:

1. There exists an observable input or choice variable $m_{it} = f_t(k_{it}, l_{it}, \omega_{it})$ such that f_t is strictly monotonic in ω_{it}.
2. ω_{it} is the only econometric unobservable in the mapping above.

The production function, equation (1.22), can be rewritten as:

$$y_{it} = \alpha + \beta_s^k k_{it} + \beta_s^l l_{it} + f_t^{-1}(k_{it}, l_{it}, m_{it}) + \epsilon_{it} \,,$$

where all regressors are now observable.

First stage As opposed to the standard proxy approach (Olley and Pakes (1996), Levinsohn and Petrin (2003)), allowing for labor dynamics with functional dependence prevents me from identifying the labor coefficient, β^l, in the first stage. Instead, in the first stage I am only able to remove the shock ϵ_{it} from the dependent variable y_{it} by treating f_t^{-1} non-parametrically and recover $\hat{\Phi}_{it}$ from:

$$y_{it} = \Phi_{it}(k_{it}, l_{it}, m_{it}) + \epsilon_{it} \,.$$

Second stage A firm will continue to operate provided its productivity level exceeds the lower bound: $\chi_{it} = 1$ if $\omega_{it} \geq \underline{\omega}_{it}$, where χ_i is a survival indicator variable. I estimate the survival probability, $\hat{P}_{it}$, by fitting a probit model on capital, labor and the proxy variable:

$$P_{it} \equiv Pr\{\chi_t = 1 \,|\, \underline{\omega}_{it}, I_{t-1}\} = h_t(k_{it-1}, l_{it-1}, m_{it-1}) \,,$$

where I_{t-1} is the information set at time $t-1$.

Third stage Given guesses for β^k and β^l, it is possible to obtain the residuals

$$\hat{\omega}_{it} = \hat{\Phi}_{it} - \beta^k k_{it} - \beta^l l_{it} \,,$$

and, exploiting the Markov chain assumption on ω_{it}, obtain the corresponding residual $\hat{\xi}_{it}$ by simply regressing $\hat{\omega}_{it}$ on $\hat{\omega}_{it-1}$ and $\hat{P}_{it}$. β_k and β_l are estimated using the following GMM criterion function:

$$\frac{1}{N}\frac{1}{T}\sum_i\sum_t \begin{pmatrix} \hat{\xi}_{it}k_{it} \\ \hat{\xi}_{it}l_{it-1} \end{pmatrix} = 0\,.$$

Allocative efficiency

This appendix summarizes the Hsieh and Klenow (2017) argument that resource misallocation can hinder aggregate productivity.

As explained in the main text, consider a framework with a final good featuring a CES production function in differentiated intermediates goods that are imperfectly substitutable. Intermediate good producers have standard Cobb-Douglas production technologies, with capital share α, and are subject to firm-specific exogenous wedges that distort (i) output, τ^y_{it}, and (ii) capital relative to labor, τ^k_{it}. The individual intermediate good producer optimization problem delivers the following first-order conditions with respect to labor, l_{it}, and capital, k_{it}:

$$MRPL_{it} = \left(\frac{1-\alpha}{\mu}\right)\left(\frac{P_{it}Y_{it}}{L_{it}}\right) = \left(\frac{1}{1-\tau^y_{it}}\right)W_t\,, \tag{1.23}$$

$$MRPK_{it} = \left(\frac{\alpha}{\mu}\right)\left(\frac{P_{it}Y_{it}}{K_{it}}\right) = \left(\frac{1+\tau^k_{it}}{1-\tau^y_{it}}\right)R_t\,, \tag{1.24}$$

where $P_{it}Y_{it}$ is firm nominal value added, W_t is the cost of labor, R_t is the cost of capital and μ is the constant markup of price over marginal cost.

Define physical and revenue productivities at the firm-level as

$$TFPQ_{it} \equiv A_{it} = \frac{Y_{it}}{K^{\alpha}_{it}L^{1-\alpha}_{it}}\,, \tag{1.25}$$

and

$$TFPR_{it} \equiv P_{it}A_{it} = \frac{P_{it}Y_{it}}{K^{\alpha}_{it}L^{1-\alpha}_{it}}\,. \tag{1.26}$$

By substituting equations (1.23) and (1.24) into equation (1.26),

$$TFPR_{it} = \mu\left(\frac{MRPK_{it}}{\alpha}\right)^{\alpha}\left(\frac{MRPL_{it}}{1-\alpha}\right)^{1-\alpha} = \mu\left(\frac{R_t}{\alpha}\right)^{\alpha}\left(\frac{W_t}{1-\alpha}\right)^{1-\alpha}\frac{\left(1+\tau^k_i\right)^{\alpha}}{1-\tau^y_i}\,,$$

it follows that optimal allocation of labor and capital ensures that firms with higher TFPQ expand production such that they charge lower prices than more unproductive firms and TFPR is equalized across plants. In other words, dispersion in TFPR is solely driven by the presence of firm-specific distortions in this model. Such distortions can lower aggregate TFP by the following expression:

$$TFP_t = \left[\sum_{i=1} \left(A_{it} \frac{\overline{TFPR}_t}{TFPR_{it}} \right)^{\sigma-1} \right]^{\frac{1}{1-\sigma}} ,$$

where $\overline{TFPR}_t$ is the revenue weighted average TFPR.

1.11.3 Model appendix

A model of two large countries: the limit case

This appendix shows that the assumptions required to treat Home as a small open economy can be derived from the steady state version of a model with two countries which are symmetric in everything except size *i.e.* Home is assumed to be small relative to Foreign. In particular, if the two countries are endowed with n and $n-1$ shares of the world's total number of potentially active firms, $\bar{M}$,

$$M = n\bar{M}, \quad M^* = (1-n)\bar{M}, \quad n \in [0,1],$$

then the limit case to be considered is one in which $n \to 0$. The productivity cutoffs of this model would be given by the steady state versions of equations (1.9) and (1.10) together with:

$$z^{*F} = \frac{\gamma + \eta N}{\frac{\alpha\gamma}{\lambda} + \eta P} \tau\epsilon (W^*)^{\sigma} , \tag{1.27}$$

$$z^{*H} = \frac{\gamma + \eta N^*}{\frac{\alpha\gamma}{\lambda^*} + \eta P^*} (W^*)^{\sigma} , \tag{1.28}$$

The number of active firms in Home and Foreign is given by equation (1.12) and

$$N^* = (1-n)\bar{M}^*(z^{*H})^{-k} + n\bar{M}(z^F)^{-k} , \tag{1.29}$$

while the aggregate price level is summarized by equation (1.13) and

$$P^* = \frac{2k+1}{2k+2} \frac{(W^*)^{\sigma} N^*}{z^{*H}} . \tag{1.30}$$

Finally, the balance of payments condition in a zero trade balance steady state can be rewritten as

$$\frac{n}{1-n} = \frac{\lambda}{\lambda^*} \left(\frac{W^*}{W} \right)^{2\sigma} \epsilon^3 \left(\frac{z^{*F}}{z^F} \right)^{(k+2)} , \tag{1.31}$$

To summarize, for a given n, the equilibrium in the model with two countries can be described by Equations (1.9), (1.10), (1.12), (1.13), (1.27)-(1.66) with nine unknown variables $\{z^H, z^F, z^{*H}, z^{*H}, N, N^*, P, P^*, W\}$, taking foreign labor input as the numeraire ($W^* = 1$).

This system, however, can be further collapsed into three equations in three unknowns, namely, z^H, z^{*H} and W:

$$\alpha\gamma\frac{1-\theta}{\theta}z^H W = W^\sigma\left[\gamma+\frac{\eta}{2k+2}\left(\frac{1}{z^H}\right)^k\bar{M}\left(n+(1-n)\left(\frac{W^\sigma}{\tau\epsilon}\right)^k\right)\right], \tag{1.32}$$

$$\alpha\gamma\frac{1-\theta}{\theta}z^{*H} = \left[\gamma+\frac{\eta}{2k+2}\left(\frac{1}{z^{*H}}\right)^k\bar{M}\left((1-n)+n\left(\frac{\epsilon}{\tau W^\sigma}\right)^k\right)\right], \tag{1.33}$$

$$\frac{n}{1-n} = \frac{W^{2\sigma(k+1)-1}}{\epsilon^{2k+1}}\left(\frac{z^{*H}}{z^H}\right)^{(k+2)}. \tag{1.34}$$

As $n \to 0$, Equation (1.38) simplifies to

$$\alpha\gamma\frac{1-\theta}{\theta}z^{*H} = \left[\gamma+\frac{\eta}{2k+2}\left(\frac{1}{z^{*H}}\right)^k\bar{M}\right],$$

which solves for z^{*H} as a function only of parameters. I have, thus, proved the first assumption: the foreign domestic productivity cutoff is not affected by changes at Home for n small enough.

Note that due to the Pareto distribution assumption, z^{*H}, cannot fall below one, the minimum value for productivity. Therefore, I need distinguish between two different cases. Suppose

$$\alpha\gamma\frac{1-\theta}{\theta} < \gamma+\frac{\eta}{2k+2}\bar{M}, \tag{1.35}$$

then the solution to the above equation is larger than one. Once, I have solved for z^{*H}, the foreign demand for the domestic variety is given by

$$q^{*F}(z) = \frac{1}{\gamma+\eta N^*}\left(\alpha+\frac{\eta}{\gamma}\frac{\theta}{1-\theta}P^*\right)-\frac{\theta}{1-\theta}\frac{1}{\gamma}p^{*F}(z), \tag{1.36}$$

where $N^* = \bar{M}\left(z^{*H}\right)^{-k}$ and P^* is a function of z^{*H} as given by Equation (1.30), and, thus, constant.

Suppose, instead, the opposite is true, and the inequality given by Equation (1.35) does not hold. In such a case, z^{*H} remains at one so that all foreign firms produce, $N^* = \bar{M}$. This also means, that the choke price for Foreign is not binding[62] and a new equation for the aggregate price level in Foreign is required. In particular, the new price level is given by

$$P^* = \left(\frac{2}{\bar{M}}-\frac{\eta}{\gamma+\eta N^*}\right)^{-1}\left[\frac{\alpha\gamma\frac{1-\theta}{\theta}}{\gamma+\eta N^*}+\frac{1}{b}\frac{k}{k+1}\right].$$

The rest of the argument follows: foreign demand for the domestic variety is given by Equation (1.36) which implies that A and B in Equation (1.7) are constants as

[62] The maximum price faced by foreign consumers is actually lower than the choke price they would be willing to pay.

none of the foreign variables *i.e.* z^{*H}, N^* and P^*, are affected by changes in Home.

Equilibrium summary

Endogenous variables: z_t^H, z_t^F, z_t^{*F}, L_t, N_t, B_t, R_t, P_t, λ_t, W_t, ϵ_t

Equilibrium conditions:

$$z_t^H = \frac{\gamma + \eta N_t}{\frac{\alpha\gamma}{\lambda_t} + \eta P_t} W_t^\sigma \,, \tag{1.37}$$

$$z_t^F = \frac{\gamma + \eta N_t}{\frac{\alpha\gamma}{\lambda_t} + \eta P_t} \tau \epsilon_t (W_t^*)^\sigma \,, \tag{1.38}$$

$$z_t^{F*} = \frac{B}{A}\frac{\tau W_t^\sigma}{\epsilon_t} \,, \tag{1.39}$$

$$N_t = M(z_t^H)^{-k} + M^*(z_t^F)^{-k} \,, \tag{1.40}$$

$$P_t = \frac{2k+1}{2k+2}\frac{W_t^\sigma N_t}{z_t^H} \,, \tag{1.41}$$

$$L_t = \frac{k}{(k+1)(k+2)} \sigma W_t^{2\sigma-1} M \left(\frac{\lambda_t}{\gamma} \left(z_t^H\right)^{-(k+2)} + \frac{B\tau^2}{\epsilon_t} \left(z_t^{F*}\right)^{-(k+2)} \right) \,, \tag{1.42}$$

$$1 = \beta R_t \mathbb{E}_t \left(\frac{\epsilon_{t+1}}{\epsilon_t} \frac{\lambda_{t+1}}{\lambda_t} \right) \,, \tag{1.43}$$

$$R_t = R_t^* + \phi \left(e^{\bar{B} - B_t} - 1 \right) + \left(e^{\xi_t - 1} - 1 \right) \,, \tag{1.44}$$

$$MB \frac{(\tau W_t^\sigma)^2}{\epsilon_t} \left(z_t^{F*}\right)^{-(k+2)} - M^* \frac{\lambda_t \left(\tau \epsilon_t (W_t^*)^\sigma\right)^2}{\gamma} \left(z_t^F\right)^{-(k+2)} = 2(k+2)\epsilon_t (B_t - R_{t-1} B_{t-1}) \,, \tag{1.45}$$

$$W_t = \prod_{s=0}^{\infty} \left(\frac{\theta}{\theta - 1} \mathbb{E}_{t-s} \left(\frac{1}{\lambda_t} \right) \right)^{\mu(1-\mu)^s} \,, \tag{1.46}$$

$$\text{monetary policy rule} \,. \tag{1.47}$$

Existence and uniqueness of steady state

This appendix solves for the steady state of the model and shows that it is unique provided $\bar{B} = 0$. To ease notation, I drop all time subscripts. The steady state is summarized by one equation in one unknown, which can be solved numerically provided parameter values.

Start by rewriting the wage equation in steady state as

$$\lambda = \frac{\theta}{\theta - 1} \frac{1}{W} \,. \tag{1.48}$$

Combine (1.37) and (1.41) to get

$$z^H \alpha\gamma = W^\sigma \lambda \left(\gamma + \frac{\eta}{2k+2} N \right) \,. \tag{1.49}$$

Rewrite z^F as a function of z^H, given equations (1.37) and (1.38),

$$z^H = \frac{\tau\epsilon}{W^\sigma} z^H \,, \tag{1.50}$$

and plug into equation (1.40)

$$N = \left(\frac{1}{z^H}\right)^k \left(M + M^* \left(\frac{W^\sigma}{\tau\epsilon}\right)^k\right) .$$

which can now be combined with equation (1.48) and (1.49) such that

$$z^H \alpha\gamma = \frac{\theta}{\theta - 1}\frac{1}{W^{1-\sigma}}\left(\gamma + \frac{\eta}{2k+2}\left(\frac{1}{z^H}\right)^k \left(M + M^* \left(\frac{W^\sigma}{\tau\epsilon}\right)^k\right)\right) . \tag{1.51}$$

Next, note that in steady state the interest rate is given by $R = \frac{1}{\beta}$ and bond holdings are $B = \bar{B}$ (see equations (1.43) and (1.44) respectively). Imposing this on the balance of payment condition, (1.45), together with equations (1.39), (1.48) and (1.50), delivers

$$M \frac{A^{k+2}}{B^{k+1}} \frac{\epsilon^{k+1}}{(\tau W^\sigma)^k} - M^* \frac{\theta}{\theta - 1} \frac{W^{\sigma(k+2)-1}}{\gamma} \frac{\left(z^H\right)^{-(k+2)}}{(\tau\epsilon)^k} = -2\,(k+2)\,\epsilon \frac{(1-\beta)}{\beta} \bar{B} \,. \tag{1.52}$$

Equation (1.52) can be rewritten in terms of z^H and then plugged into equation (1.51). This would deliver a system of one equation in one unknown: if the economy is embedded in a currency union, the exchange rate is equal to one and the unknown is W. If the economy has a floating arrangement, the wage level is equal to the target and the unknown is ϵ. In any case, there exists a steady state equilibrium.

Impose that trade balance holds in equilibrium ($\bar{B} = 0$). Equation (1.52) is simplified to

$$\frac{1}{z^H} = \left[\gamma \frac{\theta - 1}{\theta} \frac{M}{M^*} \frac{A^{k+2}}{B^{k+1}} \frac{\epsilon^{k+2}}{W^{2\sigma(k+1)-1}}\right]^{\frac{1}{k+2}} ,$$

and can now substitute for z^H in equation (1.51) as follows

$$\alpha\gamma \frac{\theta - 1}{\theta} = \left[\gamma \frac{\theta - 1}{\theta} \frac{M}{M^*} \frac{A^{k+2}}{B^{k+1}} \frac{\epsilon^{k+2}}{W^{2\sigma(k+1)-1}}\right]^{\frac{1}{k+2}} \left[\gamma + \frac{\eta}{2k+2}\left[\gamma \frac{\theta - 1}{\theta} \frac{M}{M^*} \frac{A^{k+2}}{B^{k+1}} \frac{\epsilon^{k+2}}{W^{2\sigma(k+1)-1}}\right]^{\frac{k}{k+2}} \left(M + M^* \left(\frac{w^\sigma}{\tau\epsilon}\right)^k\right)\right] .$$

The left hand side is a positive constant. The right hand side is:

1. A monotonically decreasing function in W with positive limit of zero and a negative limit of $+\infty$ in the currency union regime.

2. A monotonically increasing function in ϵ with positive limit of $+\infty$ and a negative limit of zero in the currency union regime.

Thus, in both cases, there exists a unique solution.

1.11.4 Proofs

Proof of Lemma 1

Proof. Unweighted average productivity is given by

$$\tilde{z}_t^H = \int_{z_t^H}^{\infty} z \, \frac{g(z)}{1 - G(z_t^H)} \, dz = \frac{k}{k-1} z_t^H \, .$$

Average productivity weighted by output is given by

$$\hat{z}_t^H = \int_{z_t^H}^{\infty} z \, \frac{q(z)}{q(\tilde{z}_t^H)} \, \frac{g(z)}{1 - G(z_t^H)} \, dz \, .$$

Noting that $\frac{q(z)}{q(\tilde{z}_t^H)} = \frac{z - z_t^H}{\tilde{z}_t^H - z_t^H} \frac{\tilde{z}_t}{z}$, the above expression simplifies to $\hat{z}_t^H = \tilde{z}_t^H$.

Average productivity weighted by revenue is given by

$$\bar{z}_t^H = \int_{z_t^H}^{\infty} z \, \frac{r(z)}{r(\tilde{z}_t^H)} \, \frac{g(z)}{1 - G(z_t^H)} \, dz \, .$$

Noting that $\frac{r(z)}{r(\tilde{z}_t^H)} = \frac{z^2 - (z_t^H)^2}{(\tilde{z}_t^H)^2 - (z_t^H)^2} \frac{(\tilde{z}_t)^2}{z^2}$, the above expression simplifies to $\bar{z}_t^H = \frac{2k^3}{(2k-1)(k^2-1)} z_t^H$. □

Proof of Proposition 1

Proof. By combining equations (1.9) and (1.13), the domestic productivity threshold can be rewritten as

$$z_t^H = \frac{\lambda_t W_t^\sigma}{\alpha \gamma} \left[\gamma + \frac{\eta}{2k+1} N_t \right] \, . \tag{1.53}$$

To derive the expression in Proposition 1 log-linearize equation (1.53) around its steady state. □

Proof of Proposition 2

Proof. To see this formally, combine equations (1.9), (1.10), and (1.12) to rewrite the equilibrium number of active firms in the domestic market as

$$N_t = \left(\frac{1}{z_t^H} \right)^k \left[M + M^* \left(\frac{W_t^\sigma}{\tau \left(W_t^*\right)^\sigma \epsilon_t} \right)^k \right] ,$$

and combine with the expression for z_t^H above, equation (1.53), to get

$$z_t^H - \frac{\eta}{2k+2} \frac{\lambda_t W_t^\sigma}{\alpha \gamma} \left(\frac{1}{z_t^H} \right)^k \left[M + M^* \left(\frac{W_t^\sigma}{\tau \left(W_t^*\right)^\sigma \epsilon_t} \right)^k \right] = \frac{\lambda_t W_t^\sigma}{\alpha} , \tag{1.54}$$

from here it is straightforward to see that there is a negative relationship between z_t^H and ϵ_t *i.e.* the left-hand side of equation (1.54) is increasing in both z_t^H and ϵ_t. It then follows that $\zeta_{z_t^H,\epsilon_t} = \frac{\partial z_t^H}{\partial \epsilon_t}\frac{\epsilon_t}{z_t^H} < 0$.

The relationship between z_t^H and W_t is less obvious. The right-hand side of equation (1.54) is decreasing in wages as $\lambda_t W_t^{\sigma} \propto \frac{1}{W_t^{\frac{1}{\mu}-\sigma}}$ by Lemma 2. The left-hand side, however, depends on parameter values and, thus, $\zeta_{z_t^H,W_t} = \frac{\partial z_t^H}{\partial W_t}\frac{W_t}{z_t^H} \gtrless 0$. □

Lemma 2. *There is a negative relationship between the marginal utility of income and the wage level.*

Proof. In steady state, wages are equalized across labor types and equation (1.14) can be rewritten as

$$\lambda_t = \frac{1-\theta}{\theta}\frac{1}{W_t}.$$

During the dynamics, the negative relationship still holds as

$$\lambda_t \propto \frac{1-\theta}{\theta}\frac{1}{W_t^{\frac{1}{\mu}}}.$$

□

Proof of Corollary 1

Proof. Suppose $\mu\sigma < \frac{1}{1+k}$, then the left-hand side of equation (1.54) is increasing in wages. Thus, there is an unambiguous negative relationship between z_t^H and W_t that ensures $\zeta_{z_t^H,W_t} = \frac{\partial z_t^H}{\partial W_t}\frac{W_t}{z_t^H} < 0$. □

1.11.5 Auxiliary assumption

This appendix discusses how the auxiliary assumption introduced in Section 1.7.1, the law of motion for existing firms, changes results derived in Section 1.6.

The law of motion can be written in levels as

$$M_t = \frac{\bar{M}}{\epsilon_t^{\rho}\left(\frac{W_t}{\bar{W}}\right)^{1-\rho}}. \tag{1.55}$$

The definition of aggregate productivity and the equilibrium condition for the domestic productivity threshold remain unchanged. Therefore, Lemma 1 and Proposition 1 still hold.

The proof of Proposition 2 needs to be adjusted slightly to read:

Proof. To see this formally, combine equations (1.9), (1.10), (1.12) and (1.55) to rewrite the equilibrium number of active firms in the domestic market as

$$N_t = \left(\frac{1}{z_t^H}\right)^k \frac{\bar{M}}{\epsilon_t^\rho \left(\frac{W_t}{\bar{W}}\right)^{1-\rho}} \left[1 + \frac{1-n}{n}\left(\frac{W_t^\sigma}{\tau\left(W_t^*\right)^\sigma \epsilon_t}\right)^k\right],$$

and combine with the expression for z_t^H above, equation (1.53), to get

$$z_t^H - \frac{\eta}{2k+2}\frac{\lambda_t W_t^\sigma}{\alpha\gamma}\left(\frac{1}{z_t^H}\right)^k \frac{\bar{M}}{\epsilon_t^\rho \left(\frac{W_t}{\bar{W}}\right)^{1-\rho}} \left[1 + \frac{1-n}{n}\left(\frac{W_t^\sigma}{\tau\left(W_t^*\right)^\sigma \epsilon_t}\right)^k\right] = \frac{\lambda_t W_t^\sigma}{\alpha}, \tag{1.56}$$

from here it is straightforward to see that there is a negative relationship between z_t^H and ϵ_t *i.e.* the left-hand side of equation (1.54) is increasing in both z_t^H and ϵ_t. It then follows that $\zeta_{z_t^H,\epsilon_t} = \frac{\partial z_t^H}{\partial \epsilon_t}\frac{\epsilon_t}{z_t^H} < 0$.

The relationship between z_t^H and W_t is less obvious. The right-hand side of equation (1.54) is decreasing in wages as $\lambda_t W_t^\sigma \propto \frac{1}{W_t^{\frac{1}{\mu}-\sigma}}$ by Lemma 2. The left-hand side, however, depends on parameter value and, thus, $\zeta_{z_t^H,W_t} = \frac{\partial z_t^H}{\partial W_t}\frac{W_t}{z_t^H} \gtrless 0$. □

Given the modified proof of Proposition 2, the corresponding Corollary should now read

Corollary 2. *Following a sudden stop in a currency union, a sufficient condition for $\hat{z}_t^H > 0$ is that $1 + \mu\,(1-\rho) > \mu\,\sigma(1+k)$.*

Proof. Suppose $(1+k)\,\mu\sigma < \mu(1-\rho)$, then the left-hand side of equation (1.56) is increasing in wages. Thus, there is an unambiguous negative relationship between z_t^H and W_t that ensures $\zeta_{z_t^H,W_t} = \frac{\partial z_t^H}{\partial W_t}\frac{W_t}{z_t^H} < 0$. □

However, as $\mu > 0$ and $1-\rho > 0$, it is the case that the sufficient condition in Corollary 1 is more restrictive than that in Corollary 2 *i.e.* $1 + \mu\,(1-\rho) > 1 > \mu\,\sigma(1+k)$.

1.11.6 Extensions to the model

A model with capital

This appendix describes a version of the baseline model that features pre-installed capital as the second input in the production of differentiated varieties. In particular, the unit cost of production is now given by

$$c_t = \left(\frac{W_t}{\sigma}\right)^\sigma \left(\frac{\kappa_t}{1-\kappa}\right)^{1-\sigma}, \tag{1.57}$$

where κ_t is the rental price of capital.

The evolution of the capital stock, K_t, is determined by investment by the representative consumer and depreciation. The budget constraint is accordingly modified to read:

$$\int^{N_t} p_t\left(\omega\right) q_t\left(\omega\right) d\omega + \epsilon_t B_t + V_t K_t = \int_0^1 W_t^i L_t^i \, di + \Pi_t + \epsilon_t R_{t-1} B_{t-1} + V_t (1-\delta) K_{t-1} + \kappa_t K_{t-1} \, ,$$

where V_t is the price of capital and δ is the depreciation rate.

There is an additional equilibrium condition governing the optimal capital choice:

$$V_t = \beta E_t \left[\frac{\lambda_{t+1}}{\lambda_t} \left((1-\delta) V_{t+1} + \kappa_{t+1} \right) \right] \, , \tag{1.58}$$

and a transversality condition that ensures the absence of bubbles in equilibrium:

$$\lim_{T \to \infty} \beta^T E_t \lambda_{t+T} K_{t+T} = 0 \, . \tag{1.59}$$

In addition, capital market clearing ensures that capital supplied by the representative household is equated to the aggregate demand by firms:

$$K_{t-1} = \frac{M_t}{(k+1)(k+2)} \frac{(1-\sigma) c_t^2}{\kappa_t} \left[\frac{\lambda_t}{\gamma} (z_t^H)^{-(k+2)} + \frac{B \tau^2}{\epsilon_t} (z_t^{*F})^{-(k+2)} \right] \, . \tag{1.60}$$

The rational expectations equilibrium of this extension is the set of stochastic processes $\{z_t^H, z_t^F, z_t^{*F}, IM_t, EX_t, L_t, N_t, B_t, R_t, P_t, \lambda_t, W_t, \kappa_t, c_t, V_t, M_t, K_{t-1}\}_{t=0}^{\infty}$ satisfying equations (1.3), (1.5), (1.12), (1.14), (1.57)-(1.60) and

$$\frac{z_t^H}{c_t} \left(\frac{\alpha \gamma}{\lambda_t} + \eta P_t \right) = \gamma + \eta N_t \, ,$$

$$\frac{z_t^F}{\tau \epsilon_t c_t^*} \left(\frac{\alpha \gamma}{\lambda_t} + \eta P_t \right) = \gamma + \eta N_t \, ,$$

$$z_t^{*F} = \frac{B}{A} \frac{\tau c_t}{\epsilon_t} \, ,$$

$$M_t = \frac{1}{\epsilon_t^{\rho} c_t^{1-\rho}} \, ,$$

$$P_t = \frac{2k+1}{2k+2} \frac{c_t N_t}{z_t^H} \, ,$$

$$L_t = \frac{M_t}{(k+1)(k+2)} \frac{\sigma c_t^2}{W_t} \left[\frac{\lambda_t}{\gamma} (z_t^H)^{-(k+2)} + \frac{B \tau^2}{\epsilon_t} (z_t^{*F})^{-(k+2)} \right] \, ,$$

$$EX_t - IM_t = \epsilon_t (B_t - R_{t-1} B_{t-1}) + V_t (K_t - (1-\delta) K_{t-1}) \, ,$$

$$IM_t = \frac{1}{k+2} M_t^* \frac{\lambda_t}{\gamma} \frac{(\tau \epsilon_t c_t^*)^2}{2} \left(z_t^F \right)^{-(k+2)} \, ,$$

$$EX_t = \frac{1}{k+2} M_t \frac{B}{2} \frac{(\tau c_t)^2}{\epsilon_t} \left(z_t^{*F}\right)^{-(k+2)} ,$$

given the exogenous process $\{\xi_t\}_{t=0}^{\infty}$ and the central bank's policy $\{\epsilon_t\}_{t=0}^{\infty}$. The foreign unit production cost, c_t^*, is normalized to one.

This extension of the model is parametrized following the same principles as the baseline frame-work. This implies setting, $A = 1.123$, $B = 0.540$ and $\phi = 0.012$; all other values are unchanged. In addition, the depreciation rate, δ, is set to 2.5% following Melitz and Ghironi (2005) among many others.

Extensive versus intensive margin

This appendix describes a version of the baseline model that accounts for common shocks to firm-level productivity. Suppose the efficiency of a firm in transforming the input bundle into output is described by the composite $Z_t z$, where Z_t is a new time-varying common component and z is the firm-specific productivity level drawn from a Pareto distribution as described in Section 1.5.

The definition of the rational expectations equilibrium has to be adjusted slightly. In particular, equilibrium conditions (1.9), (1.11), (1.13), (1.15) and (1.18) are substituted by:

$$z_t^H = \frac{\gamma + \eta N_t}{\frac{\alpha\gamma}{\lambda_t} + \eta P_t} \frac{W_t^\sigma}{Z_t} ,$$

$$z_t^{*F} = \frac{B}{A} \frac{\tau W_t^\sigma}{\epsilon_t Z_t} ,$$

$$P_t = \frac{2k+1}{k+1} \frac{W_t^\sigma}{Z_t} \frac{N_t}{z_t^H} ,$$

$$L_t = \frac{\sigma k}{(k+1)(k+2)} \frac{M}{W_t} \left(\frac{W_t}{Z_t}\right)^2 \left[\frac{\lambda_t}{\gamma} \left(z_t^H\right)^{-(k+2)} + B \frac{\tau^2}{\epsilon_t} \left(z_t^{*F}\right)^{-(k+2)}\right] ,$$

$$EX_t = \frac{1}{k+2} M \frac{B}{2} \frac{(\tau W_t^\sigma)^2}{Z_t^2 \epsilon_t} \left(z_t^{*F}\right)^{-(k+2)} .$$

and Z_t is assumed to be an exogenous AR(1) process with auto-correlation coefficient equal to 0.99.

In addition, aggregate productivity will now be determined by both the productivity cutoff and the shock to the common component. In particular, unweighted and average productivity weighted by output are both given by: $\frac{k}{k-1} Z_t z_t^H$. Average productivity weighted by revenue is given by: $\frac{2k^3}{(2k-1)(k^2-1)} Z_t z_t^H$.

Long-run analysis

This appendix describes a long-run version of the baseline model where the number of existing firms, M_t, is endogenous. The set-up follows Ottaviano (2012) in putting Melitz and Ottaviano (2008) in a DSGE framework. The key innovation is the introduction of capital which is supplied by a second sector, accumulated by consumers and required for the set-up of firms producing the differentiated varieties. In what follows, I highlight how these assumptions and new implications fit into the set-up presented in Section 1.5.

The representative household As explained in the main text, the representative consumer is allowed to buy shares, x_t, of the economy's capital stock, K_t, at price, V_t. While capital is assumed to fully depreciate after one period; the investment entitles the representative consumer to a fraction of next period's aggregate firm profit. The consumer budget constraint is correspondingly adjusted to read:

$$\int^{N_t} p_t(\omega)\, q_t(\omega)\, d\omega + \epsilon_t B_t + x_t V_t K_t = \int_0^1 W_t^i L_t^i \, di + x_{t-1}\Pi_t + \epsilon_t R_{t-1} B_{t-1}\,.$$

Regarding the household's optimization problem, there is an additional optimality condition describing the purchase of capital shares. In particular:

$$1 = \beta E_t \left[\frac{\lambda_{t+1}}{\lambda_t} \frac{\Pi_{t+1}}{V_t K_t} \right] .$$

Capital investment is encouraged when the price of capital is low or when expected future returns are high. Given risk aversion, returns are adjusted by the stochastic discount factor: returns are more desirable whenever the marginal utility of income is higher.

Production of capital Capital is produced under perfect competition using a Cobb-Douglas technology that combines units of domestic labor, $l_t^{k,H}$ and foreign labor, $l_t^{k,F}$: $K_t = \left(l_t^{k,H}\right)^{\rho} \left(l_t^{k,F}\right)^{1-\rho}$.

Producers of capital choose labor inputs such that costs are minimized. For this analysis, only the demand for domestic labor is relevant,

$$l_t^{k,H} = \left(\frac{\rho}{1-\rho} \frac{\epsilon_t}{W_t} \right)^{1-\rho} K_t. \tag{1.61}$$

Production of differentiated varieties I assume that f_E units of capital are required for a firm to produce a differentiated variety. The timing is such that the fixed entry cost is due one period before the firm is able to start production. This implies that the realization of the firm's productivity draw is still unknown. The resulting free-entry condition pins down the number of firms that will be potentially active in period $t+1$, denoted by M_t:

$$M_t = \frac{K_t}{f_E}\,. \tag{1.62}$$

Aggregation and market clearing The number of active firms in the domestic market, N_t, has to be modified to account for the new timing assumption. In particular, the number of firms at time t will depend on the number of firms that paid the fixed capital requirement in period $t-1$ such that:

$$N_t = M_{t-1}\left(\frac{1}{z_t^H}\right)^k + M^*\left(\frac{1}{z_t^F}\right)^k . \tag{1.63}$$

Aggregate labor demand is augmented to include the domestic labor input used in the production of capital as given by equation (1.61), such that the labor market clearing condition now reads:

$$L_t = \frac{k}{(k+1)(k+2)}\frac{\sigma}{W_t^{1-2\sigma}}M_{t-1}\left[\frac{\lambda_t}{\gamma}\left(z_t^H\right)^{-(k+2)} + B\frac{\tau^2}{\epsilon_t}\left(z_t^{*F}\right)^{-(k+2)}\right] + \left(\frac{\rho}{1-\rho}\frac{\epsilon_t}{W_t}\right)^{1-\rho} f_E M_t\,, \tag{1.64}$$

where the free market condition, equation (1.62), is used to substitute for capital.

Given the capital investment decision, aggregate profit is now a variable of interest. It is computed by summing profits of domestic and export sales. More precisely,

$$\Pi_t = \frac{k}{2(k+1)(k+2)}W_t^{2\sigma}M_{t-1}\left[\frac{\lambda_t}{\gamma}\left(z_t^H\right)^{-(k+2)} + B\frac{\tau^2}{\epsilon_t}\left(z_t^{*F}\right)^{-(k+2)}\right] . \tag{1.65}$$

A new market clearing condition for capital ensures that demand by consumers is equated to supply by producers. Given the perfect competition assumption, this simply implies that the price of capital is equal to its marginal cost. Formally,

$$V_t = \left(\frac{W_t}{\rho}\right)^\rho\left(\frac{\epsilon_t W_t^*}{1-\rho}\right)^{1-\rho} .$$

As the consumer's budget constraint has been modified, the resulting balance of payment condition is:

$$EX_t - IM_t = \epsilon_t\left(B_t - R_t B_{t-1}\right) + \left(\frac{W_t}{\rho}\right)^\rho\left(\frac{\epsilon_t W_t^*}{1-\rho}\right)^{1-\rho}(1-\rho)f_e M_t\,, \tag{1.66}$$

where EM_t and IM_t, the total export and import revenues in domestic currency terms, are given by:

$$EX_t = \frac{1}{k+2}M_{t-1}\frac{B}{2}\frac{\left(\tau W_t^\sigma\right)^2}{\epsilon_t}\left(z_t^{*F}\right)^{-(k+2)} , \tag{1.67}$$

and equation (1.17) respectively. Note that the above balance of payment condition is derived by imposing that, in equilibrium, capital shares add up to one.

Solving the model The rational expectations equilibrium of this extension is the set of stochastic processes $\{z_t^H, z_t^F, z_t^{*F}, IM_t, EX_t, L_t, N_t, B_t, R_t, P_t, \lambda_t, W_t, M_{t-1}, \Pi_t\}_{t=0}^\infty$ satisfying equations (1.3), (1.5), (1.9)-(1.11), (1.13), (1.14), (1.17), (1.20), (1.63)-(1.67) given the exogenous process $\{\xi_t\}_{t=0}^\infty$ and the central bank's policy $\{\epsilon_t\}_{t=0}^\infty$.

The foreign wage, W_t^*, is normalized to one.

This extension of the model is parametrized following the same principles as the baseline framework. The additional parameter f_E is calibrated such that the economy starts at the same steady state as the baseline.

The Friedman view

This appendix describes a version of the baseline model with no firm dynamics. Suppose there is a fixed number of firms, all of which feature the same productivity level. While the solutions to the household's and the (representative) firm's optimization problems are unchanged, the choke price is no longer a relevant variable, provided that there is positive production in all markets. In other words, the equilibrium cannot be written in terms of productivity thresholds.

Instead, the rational expectations equilibrium is now defined as the set of stochastic processes $\left\{q_t^H, q_t^F, q_t^{*F}, p_t^F, p_t^{*F}, L_t, B_t, R_t, \lambda_t, W_t, \epsilon_t\right\}_{t=0}^{\infty}$ satisfying (1.3), (1.5), (1.14),

$$q_t^H = \frac{1}{2\gamma + \eta(N_t^H + N_t^F)}\left(\alpha - \frac{\lambda_t W_t^{\sigma}}{Z_t}\right),$$

$$q_t^F = \frac{1}{2\gamma + \eta(N_t^H + N_t^F)}\left(\alpha - \frac{\lambda_t \tau \epsilon_t}{Z_t}\right),$$

$$q_t^{*F} = \frac{1}{2}\left(A - B\frac{\tau W_t^{\sigma}}{\epsilon_t Z_t}\right),$$

$$L_t = \frac{\sigma}{W_t^{1-\sigma}}\frac{N_t^H}{Z_t}\left(q_t^H + \tau q_t^{*F}\right),$$

$$p_t^F = \frac{1}{2\gamma + \eta(N_t^H + N_t^F)}\left(\frac{\alpha\gamma}{\lambda_t} - \left(\gamma + \eta(N_t^H + N_t^F)\right)\frac{\tau\epsilon_t}{Z_t}\right),$$

$$p_t^{*F} = \frac{1}{2B}\left(A + B\frac{\tau W_t^{\sigma}}{\epsilon_t Z_t}\right),$$

$$\epsilon_t N_t^H p_t^{*F} q_t^{*F} - N_t^F p_t^F q_t^F = \epsilon_t\left(B_t - R_{t-1}B_{t-1}\right),$$

given the exogenous processes $\left\{Z_t, N_t^H, N_t^F, \xi_t\right\}_{t=0}^{\infty}$ and the central bank's policy $\{\epsilon_t\}_{t=0}^{\infty}$.

I need to re-calibrate some of the parameters before solving this version of the model. I set $Z_t = 1$, $N_t^H = n\bar{M}$ and $N_t^F = (1-n)\bar{M}$ at all times where n and $\bar{M}$ are calibrated as in the baseline model. The debt elasticity of interest rate, $\phi = 2.1$, is set again to match output volatility.

Figure 1.9: Effect of a sudden stop in the distribution of TFP

(a) 1992-93 episode

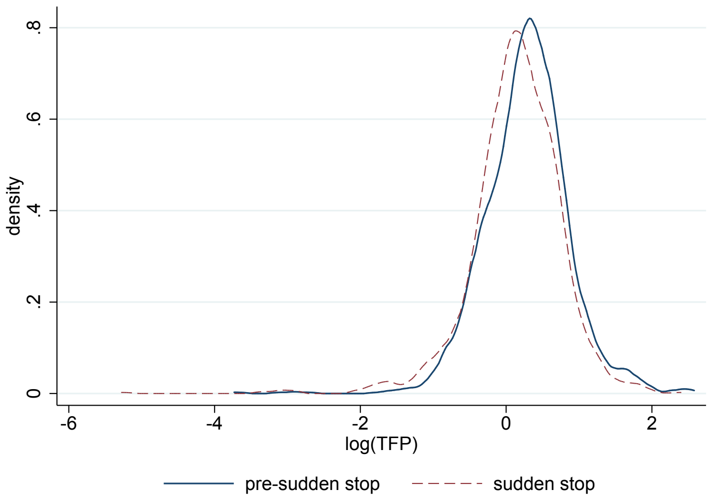

(b) 2009-13 episode

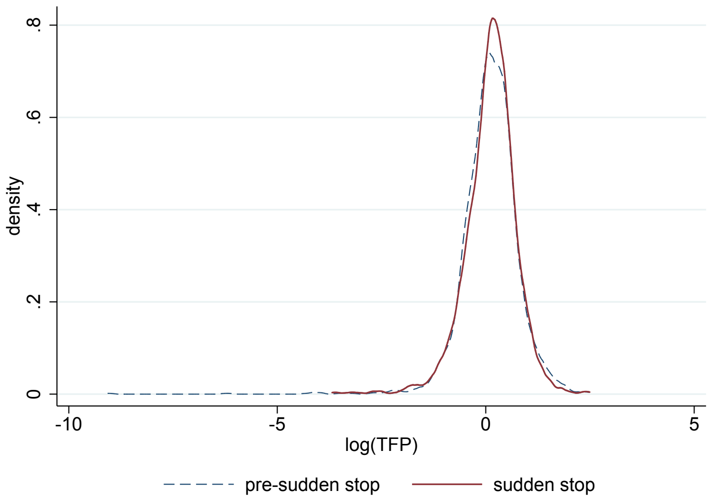

Notes: This graph plots kernel density estimates for firm-level log (TFP) before and after a sudden stop. Panel (a) refers to the 1992-93 episode, while Panel (b) focuses on the 2009-13 episode. The corresponding base and end years are 1991 and 1993 for the first episode; 2009 and 2013 for the second episode.

Source: ESEE data and own calculations.

Figure 1.10: Impulse response functions

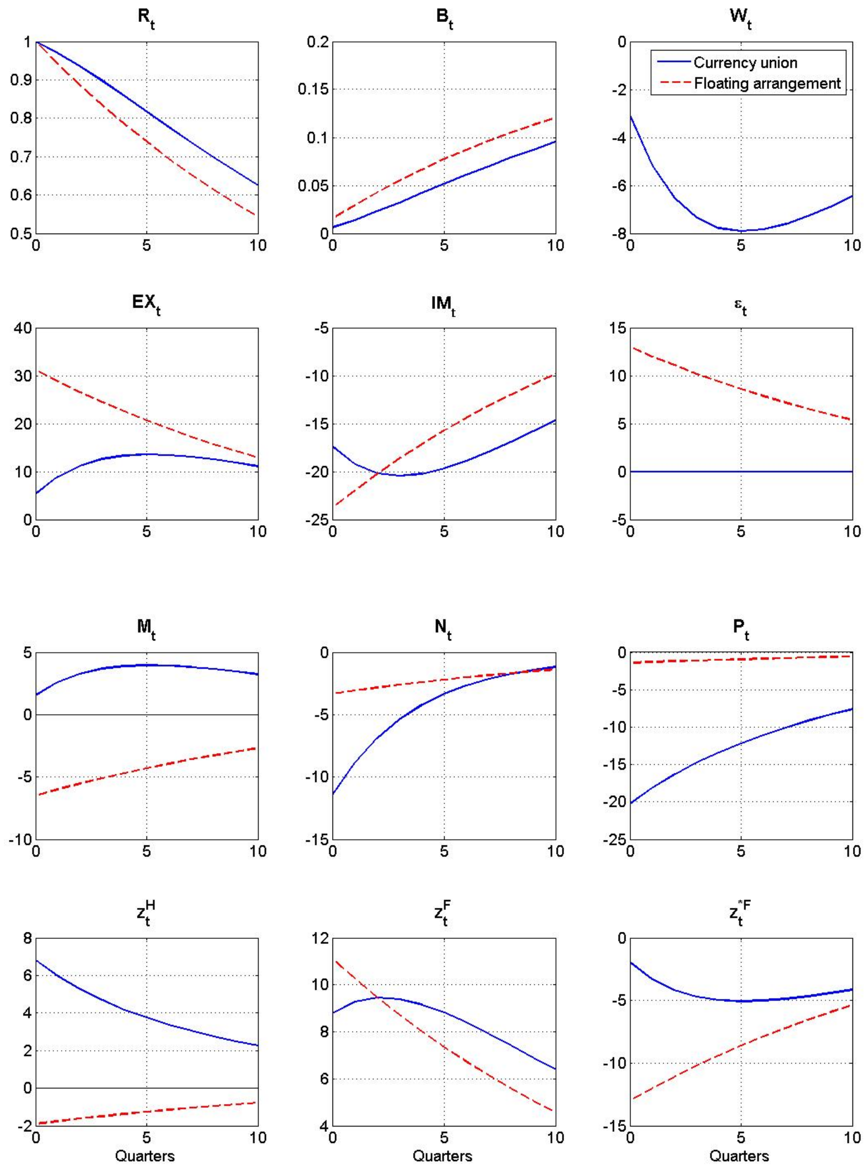

Notes: These figures plot the IRFs of all endogenous variables to a one standard deviation shock to the country-specific risk premium. All variables but bonds, B_t, are expressed in log deviations from steady state. The holding of bonds, assumed to be zero in steady state, is expressed in levels.

Figure 1.11: A model with capital - Macroeconomic effects of a sudden stop

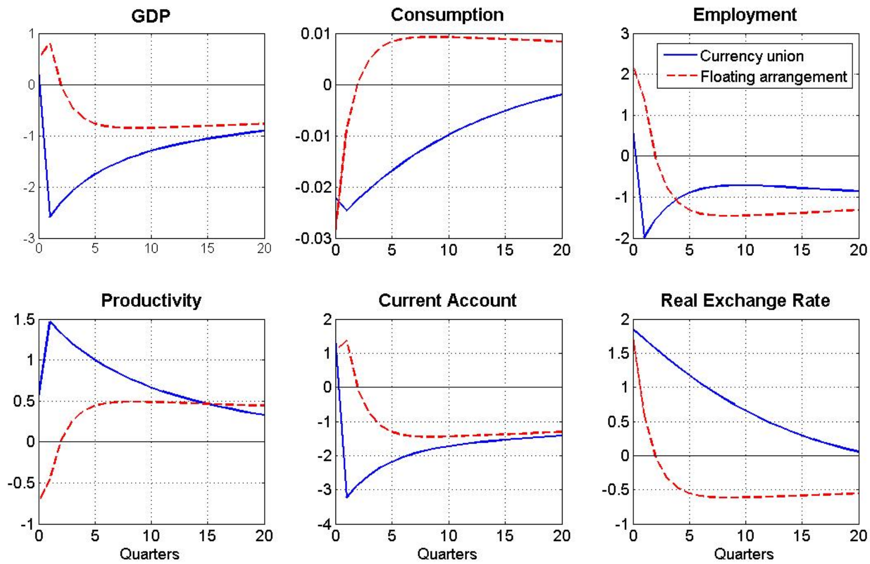

Notes: These figures plot the IRFs of key macroeconomic variables to a one standard deviation positive shock to the country-specific risk premium in a version of the model featuring pre-installed physical capital as described in Appendix 1.11.6. All variables but the current account are expressed in log deviations from steady state. The current account, assumed to be zero in steady state, is expressed in levels.

Figure 1.12: A model with capital - other variables

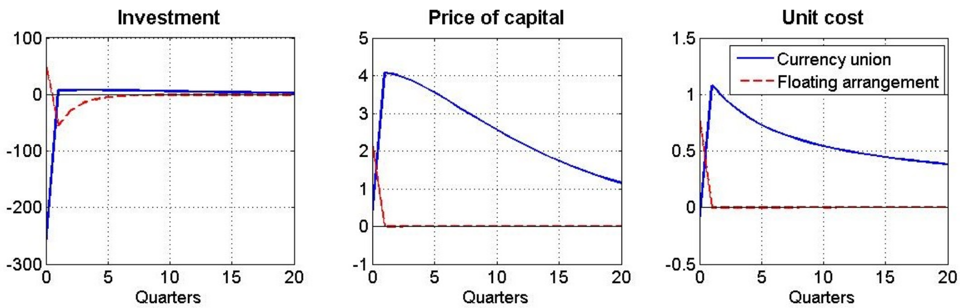

Notes: These figures plot the IRFs of investment, the unit cost of production and the rental rate of capital to a one standard deviation positive shock to the country-specific risk premium in a version of the model featuring pre-installed physical capital as described in Appendix 1.11.6. All variables are expressed in log deviations from steady state.

Figure 1.13: Extensive versus intensive margin - Macroeconomic effects of a sudden stop

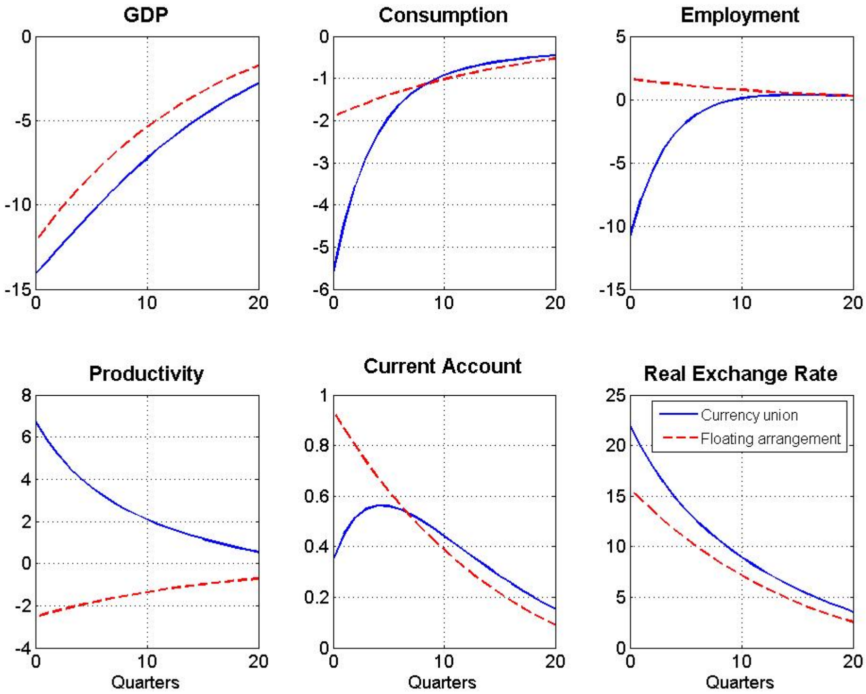

Notes: These figures plot the IRFs of key macroeconomic variables to a one standard deviation positive shock to the country-specific risk premium and a one standard deviation negative shock to the aggregate component of firm TFP. All variables but the current account are expressed in log deviations from steady state. The current account, assumed to be zero in steady state, is expressed in levels.

Figure 1.14: Extensive versus intensive margin - TFP decomposition

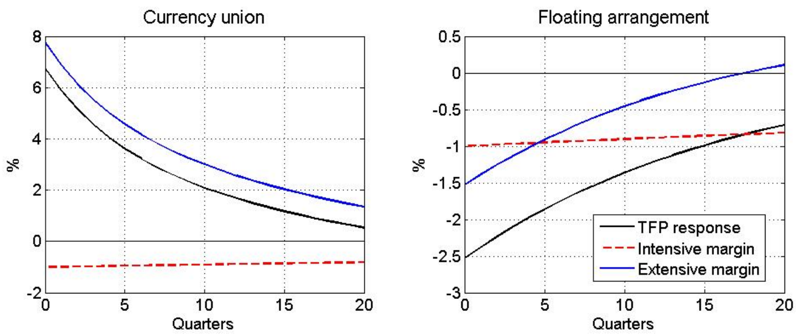

Notes: These figures decompose the overall response of TFP to a sudden stop as described above into the change in the aggregate component of firm-level productivity (the intensive margin) and the change in the productivity threshold (extensive margin). All variables are expressed in log deviations from steady state.

Figure 1.15: Long-run analysis - Macroeconomic effects of a sudden stop

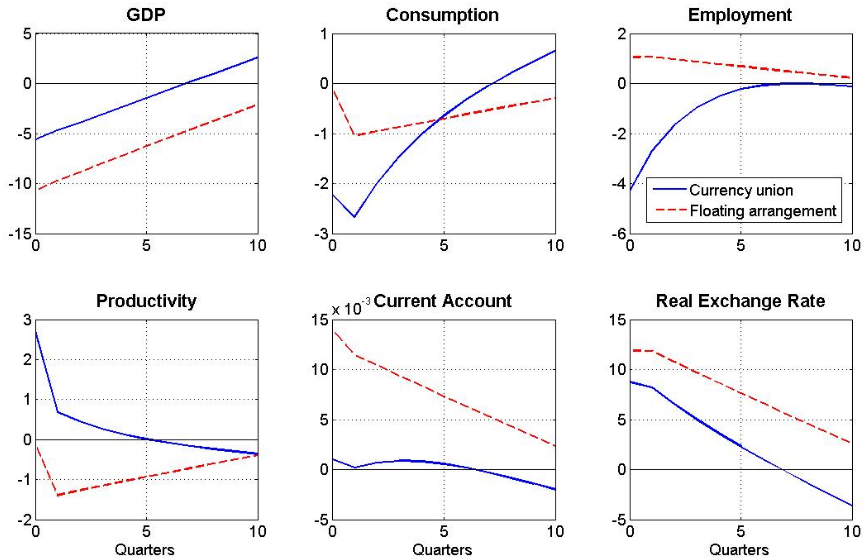

Notes: These figures plot the IRFs of key macroeconomic variables to a one standard deviation positive shock to the country-specific risk premium in the long run version of the model as described in Appendix 1.11.6. All variables but the current account are expressed in log deviations from steady state. The current account, assumed to be zero in steady state, is expressed in levels.

Figure 1.16: Currency union versus floating arrangement - consumption Loss

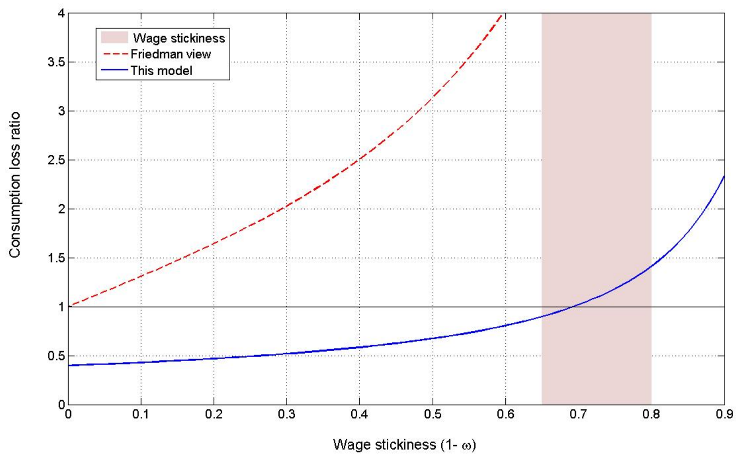

Notes: This figure plots the ratio of the consumption loss under a currency union to the consumption loss under a floating arrangement following a sudden stop. The consumption loss is calculated as the discounted sum of consumption log deviations after a one standard deviation shock to the country-specific risk premium. The blue solid line refers to the baseline model. The red dashed line refers to a version of the model with no firm dynamics (firms are homogeneous and the number of firms is constant) and linear production in labor ($\sigma = 1$). The shaded area shows the range of plausible values for the wage rigidity parameter as discussed by the literature.

Figure 1.17: Currency union versus floating arrangement - welfare Loss

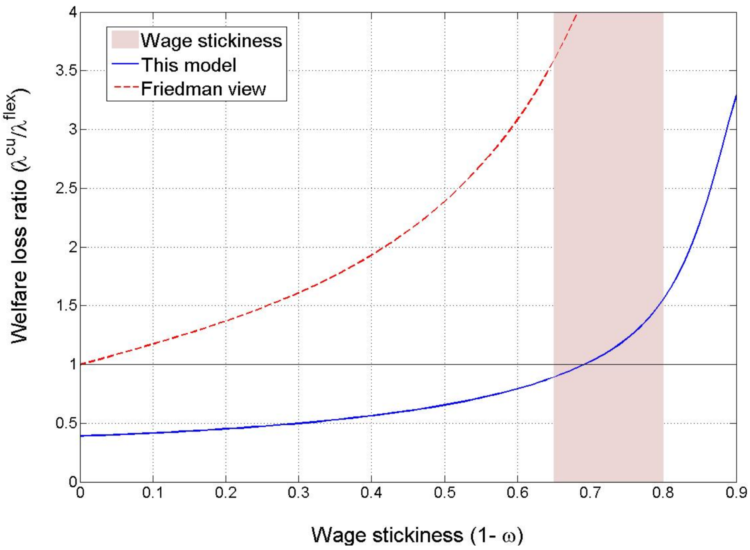

Notes: This figure plots the ratio of the welfare loss under a currency union to the welfare loss under a floating arrangement following a sudden stop. The welfare loss is calculated as the share of steady state employment that makes the representative consumer indifferent between the steady state and the dynamic equilibrium path after a one standard deviation shock to the country-specific risk premium. The blue solid line refers to the baseline model. The red dashed line refers to a version of the model with no firm dynamics (firms are homogeneous and the number of firms is constant) and linear production in labor ($\sigma = 1$). The shaded area shows the range of plausible values for the wage rigidity parameter as discussed by the literature.

Table 1.4: List of sudden stops

country	start year	end year	exchange rate	country	start year	end year	exchange rate
Albania	1991	1992	4	Macedonia FYR	2009	2010	2
Argentina	1995	1995	2	Malaysia	1998	1998	4
Argentina	1999	2002	4	Mali	1991	1991	1
Argentina	2014	2014	3	Mexico	1995	1995	4
Belarus	2014	2015	3	Moldova	1998	2003	3
Brazil	2015	2015	4	Moldova	2012	2013	3
Bulgaria	1991	1991	4	Morocco	1996	1996	3
Bulgaria	2009	2010	2	New Zealand	2004	2010	4
Chile	1999	1999	3	Nicaragua	1991	1991	2
Chile	2009	2010	4	Oman	1999	2000	2
Colombia	1998	1999	3	Oman	2010	2010	2
Croatia	1997	2002	2	Philippines	1998	1998	4
Croatia	2009	2010	2	Poland	1990	1990	4
Cyprus	2011	2011	1	Portugal	2001	2003	1
Czech Rep.	1997	2002	3	Portugal	2009	2013	1
Czech Rep.	2008	2008	3	Romania	1999	1999	4
Czech Rep.	2011	2013	3	Russia	1998	2002	3
Ecuador	1999	2000	0	Rwanda	1994	1994	4
Estonia	1996	2001	2	Saudi Arabia	1992	1992	2
Estonia	2008	2009	2	Saudi Arabia	1999	2000	2
Ethiopia	1991	1991	3	Senegal	1994	1994	1
Ethiopia	2003	2003	3	Sierra Leone	1996	1996	4
Finland	1991	1993	3	Slovak Republic	1997	2002	3
Finland	2013	2013	1	South Africa	2008	2008	4
France	1991	1993	2	Spain	1993	1993	3
Gabon	1999	1999	1	Spain	2009	2013	1
Greece	1993	1993	2	Sri Lanka	2001	2001	3
Greece	2009	2013	1	Sudan	2010	2010	3
Haiti	2003	2003	4	Sweden	1991	1991	3
Haiti	2009	2010	3	Thailand	1997	1998	4
Indonesia	1998	1998	4	Turkey	1994	1994	4
Iran	1992	1995	4	Turkey	2001	2001	4
Ireland	2009	2014	1	Ukraine	1998	2003	2
Israel	2001	2001	3	Ukraine	2014	2015	4
Italy	1993	1994	3	United Kingdom	1990	1991	3
Italy	2011	2014	1	United States	2007	2007	4
Kenya	1991	1992	4	Uruguay	2001	2001	3
Korea	1997	1998	4	Venezuela	1994	1994	4
Latvia	2008	2009	3	Venezuela	1999	2000	3
Lithuania	1997	2002	2	Yemen Rep. of	2009	2014	3

Table 1.5: Moments of the distribution

	1992-93 episode		2009-13 episode	
	pre-sudden stop	sudden stop	pre-sudden stop	sudden stop
mean	0.28	0.14	0.11	0.12
mode	0.29	0.17	0.14	0.16
sd	0.58	0.62	0.69	0.62
skewness	-0.40	-1.24	-2.37	-0.89
kurtosis	7.04	10.42	27.92	7.13
min	-3.73	-5.28	-9.07	-3.68
max	2.58	2.40	2.49	2.49

Notes: This table summarizes moments of the distribution of firm-level log (TFP) before and after a sudden stop. The first two columns refer to the 1992-93 episode, while the last two focus on the 2009-13 episode. The corresponding base and end years are 1991 and 1993 for the first episode; 2009 and 2013 for the second episode.

Source: ESEE data and own calculations.

1.12 Bibliography

Ackerberg, D. A., K. Caves, and G. Frazer (2015): “Identification Properties of Recent Production Function Estimators,” *Econometrica*, 83, 2411–2451.

Almunia, M., P. Antras, D. Lopez-Rodriguez, and E. Morales (2018a): “Venting Out: Exports during a Domestic Slump,” .

Almunia, M., D. López-Rodríguez, and E. Moral-Benito (2018b): “Evaluating the Macro-Representativeness of a Firm-Level Database: an Application for the Spanish Economy,” .

Basu, S. and J. Fernald (2001): *Why Is Productivity Procyclical? Why Do We Care?*, University of Chicago Press, 225–302.

Behrens, K. and Y. Murata (2006): “Gains from Trade and Efficiency under Monopolistic Competition: a Variable Elasticity Case,” .

Benigno, G. and L. Fornaro (2014): “The Financial Resource Curse,” *The Scandinavian Journal of Economics*, 116, 58–86.

Bernard, A. B., J. Eaton, J. B. Jensen, and S. Kortum (2003): “Plants and Productivity in International Trade,” *American Economic Review*, 93, 1268–1290.

Bilbiie, F., F. Ghironi, and M. Melitz (2008): “Monetary Policy and Business Cycles with Endogenous Entry and Product Variety,” *NBER Macroeconomics Annual*, 22, 299–353.

Bilbiie, F. O., I. Fujiwara, and F. Ghironi (2014): “Optimal Monetary Policy with Endogenous Entry and Product Variety,” *Journal of Monetary Economics*, 64, 1–20.

Bilbiie, F. O., F. Ghironi, and M. J. Melitz (2012): “Endogenous Entry, Product Variety, and Business Cycles,” *Journal of Political Economy*, 120, 304–345.

Braggion, F., L. J. Christiano, and J. Roldos (2009): “Optimal Monetary Policy in a ‘Sudden Stop’,” *Journal of Monetary Economics*, 56, 582–595.

Caballero, R. J., M. L. Hammour, et al. (1994): “The Cleansing Effect of Recessions,” *American Economic Review*, 84, 1350–1368.

Calvo, G., A. Izquierdo, and L.-F. Mejía (2004): “On the Empirics of Sudden Stops: the Relevance of Balance-Sheet Effects,” .

Calvo, G. A. (1998): “Capital Flows and Capital-Market Crises: the Simple Economics of Sudden Stops,” *Journal of Applied Economics*, 1, 35–54.

Calvo, G. A. and E. Talvi (2005): "Sudden Stop, Financial Factors and Economic Collpase in Latin America: Learning from Argentina and Chile," .

Cavallo, E. A. and J. A. Frankel (2008): "Does Openness to Trade Make Countries More Vulnerable to Sudden Stops, or Less? Using Gravity to Establish Causality," *Journal of International Money and Finance*, 27, 1430–1452.

Christiano, L. J., C. Gust, and J. Roldos (2004): "Monetary Policy in a Financial Crisis," *Journal of Economic Theory*, 119, 64–103.

Cúrdia, V. (2007): "Monetary Policy under Sudden Stops," .

De Loecker, J. and F. Warzynski (2012): "Markups and Firm-Level Export Status," *American Economic Review*, 102, 2437–71.

Del Gatto, M., G. Mion, and G. Ottaviano (2006): "Trade Integration, Firm Selection and the Costs of Non-Europe," .

Demidova, S. and A. Rodriguez-Clare (2009): "Trade Policy under Firm-Level Heterogeneity in a Small Economy," *Journal of International Economics*, 78, 100–112.

Dias, D. A. and C. R. Marques (2018): "Every Cloud Has a Silver Lining: Micro-Level Evidence on the Cleansing Effects of the Portuguese Financial Crisis," .

Dias, D. A., C. R. Marques, and C. Richmond (2016): "Misallocation and Productivity in the Lead Up to the Eurozone Crisis," *Journal of Macroeconomics*, 49, 46–70.

Drechsel, T. and S. Tenreyro (2017): "Commodity Booms and Busts in Emerging Economies," *Journal of International Economics*, 112, 200–218.

Druant, M., S. Fabiani, G. Kezdi, A. Lamo, F. Martins, and R. Sabbatini (2009): "How Are Firms' Wages and Prices Linked: Survey Evidence in Europe," .

Farhi, E., G. Gopinath, and O. Itskhoki (2013): "Fiscal Devaluations," *Review of Economic Studies*, 81, 725–760.

Feenstra, R. C. (2003): "A Homothetic Utility Function for Monopolistic Competition Models, without Constant Price Elasticity," *Economics Letters*, 78, 79–86.

Fernald, J. G. and J. C. Wang (2016): "Why has the cyclicality of productivity changed? What does it mean?" *Annual Review of Economics*, 8, 465–496.

Fernández, C., A. Lacuesta, J. M. Montero, and A. Urtasun (2015): "Heterogeneity of Markups at the Firm Level and Changes during the Great Recession: the Case of Spain," .

Fornaro, L. (2015): "Financial Crises and Exchange Rate Policy," *Journal of International Economics*, 95, 202–215.

Foster, L., C. Grim, and J. Haltiwanger (2016): "Reallocation in the Great Recession: Cleansing or Not?" *Journal of Labor Economics*, 34, S293–S331.

Foster, L., J. C. Haltiwanger, and C. J. Krizan (2001): *Aggregate Productivity Growth: Lessons from Microeconomic Evidence*, University of Chicago Press, 303–372.

Friedman, M. (1953): *Essays in Positive Economics*, University of Chicago Press.

Galí, J. and T. Monacelli (2016): "Understanding the Gains from Wage Flexibility: the Exchange Rate Connection," *American Economic Review*, 106, 3829–3868.

Galí, J. and T. Van Rens (2014): "The vanishing procyclicality of labor productivity," .

Gandhi, A., S. Navarro, and D. Rivers (2016): "On the Identification of Production Functions: How Heterogeneous is Productivity?" .

Garcia-Cicco, J., R. Pancrazi, and M. Uribe (2010): "Real Business Cycles in Emerging Countries?" *American Economic Review*, 100, 2510–31.

García-Santana, M., E. Moral-Benito, J. Pijoan-Mas, and R. Ramos (2016): "Growing like Spain: 1995-2007," .

Gopinath, G., Ş. Kalemli-Özcan, L. Karabarbounis, and C. Villegas-Sanchez (2017): "Capital Allocation and Productivity in South Europe," *The Quarterly Journal of Economics*, 132, 1915–1967.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988): "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review*, 402–417.

Griliches, Z. and H. Regev (1995): "Firm Productivity in Israeli Industry 1979–1988," *Journal of Econometrics*, 65, 175–203.

Gros, D. and C. Alcidi (2015): "Country Adjustment to a 'Sudden Stop': Does the Euro Make a Difference?" *International Economics and Economic Policy*, 12, 5–20.

Guidotti, P. E., F. Sturzenegger, A. Villar, J. de Gregorio, and I. Goldfajn (2004): "On the Consequences of Sudden Stops," *Economia*, 4, 171–214.

Hsieh, C.-T. and P. J. Klenow (2009): "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 124, 1403–1448.

——— (2017): "The Reallocation Myth," .

Ilzetzki, E., C. M. Reinhart, and K. S. Rogoff (2017): "Exchange Arrangements Entering the 21st Century: Which Anchor Will Hold?" .

Kehoe, T. J. and K. J. Ruhl (2009): "Sudden Stops, Sectoral Reallocations, and the Real Exchange Rate," *Journal of Development Economics*, 89, 235–249.

Klein, M. W. and J. C. Shambaugh (2008): "The dynamics of exchange rate regimes: Fixes, floats, and flips," *Journal of International Economics*, 75, 70–92.

Le Bihan, H., J. Montornès, and T. Heckel (2012): "Sticky Wages: Evidence from Quarterly Microeconomic Data," *American Economic Journal: Macroeconomics*, 4, 1–32.

Levinsohn, J. and A. Petrin (2003): "Estimating Production Functions Using Inputs to Control for Unobservables," *The Review of Economic Studies*, 70, 317–341.

Mankiw, N. G. and R. Reis (2002): "Sticky Information versus Sticky Prices: a Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 117, 1295–1328.

Melitz, M. and F. Ghironi (2005): "International Trade and Macroeconomic Dynamics with Heterogeneous Firms," *The Quarterly Journal of Economics*, 120, 865–915.

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695–1725.

Melitz, M. J. and G. I. Ottaviano (2008): "Market Size, Trade, and Productivity," *The Review of Economic Studies*, 75, 295–316.

Melitz, M. J. and S. Polanec (2015): "Dynamic Olley-Pakes Productivity Decomposition with Entry and Exit," *The Rand Journal of Economics*, 46, 362–375.

Melitz, M. J. and S. J. Redding (2014): "Heterogeneous Firms and Trade," *Handbook of International Economics,*, 4, 1–54.

Mendoza, E. G. (2006): "Endogenous Sudden Stops in a Business Cycle Model with Collateral Constraints: a Fisherian Deflation of Tobin's Q," .

——— (2010): "Sudden Stops, Financial Crises, and Leverage," *American Economic Review*, 100, 1941–66.

Merler, S. and J. Pisani-Ferry (2012): "Sudden Stops in the Euro Area," *Review of Economics and Institutions*, 3.

MEZA, F. AND E. QUINTIN (2007): "Factor Utilization and the Real Impact of Financial Crises," *The BE Journal of Macroeconomics*, 7, 1–41.

NEUMEYER, P. A. AND F. PERRI (2005): "Business Cycles in Emerging Economies: the Role of Interest Rates," *Journal of Monetary Economics*, 52, 345–380.

OLLEY, G. S. AND A. PAKES (1996): "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64, 1263–1297.

OTTAVIANO, G. (2012): "Firm Heterogeneity, Endogenous Entry, and the Business Cycle," .

OTTAVIANO, G., T. TABUCHI, AND J.-F. THISSE (2002): "Agglomeration and Trade Revisited," *International Economic Review*, 409–435.

REIS, R. (2013): "The Portuguese Slump and Crash and the Euro Crisis," *Brookings Papers on Economic Activity*, 2013, 143–210.

SCHMITT-GROHÉ, S. AND M. URIBE (2003): "Closing Small Open Economy Models," *Journal of International Economics*, 61, 163–185.

——— (2007): "Optimal Simple and Implementable Monetary and Fiscal RFsules," *Journal of Monetary Economics*, 54, 1702–1725.

——— (2016): "Downward Nominal Wage Rigidity, Currency Pegs, and Involuntary Unemployment," *Journal of Political Economy*, 124, 1466–1514.

SCHUMPETER, J. A. ET AL. (1939): *Business cycles*, McGraw-Hill New York.

SHAMBAUGH, J. C. (2004): "The Effect of Fixed Exchange Rates on Monetary Policy," *The Quarterly Journal of Economics*, 119, 301–352.

SYVERSON, C. (2004): "Market Structure and Productivity: A Concrete Example," *Journal of Political Economy*, 112, 1181–1222.

WOODFORD, M. (2011): *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.

WOOLDRIDGE, J. M. (2009): "On Estimating Firm-Level Production Functions Using Proxy Variables to Control for Unobservables," *Economics Letters*, 104, 112–114.

Chapter 2

How Do Central Banks Control Inflation? A Guide for the Perplexed

2.1 Introduction

How do central banks keep inflation on target? How do we prevent hyperinflations setting in, with all their tragic consequences for human welfare? And, if we fear that the economy is stuck in secular stagnation, in a liquidity trap, or in a fiscal crisis, is the central bank powerless to prevent exploding deflation or inflation? These are crucial questions that a student fresh off a macroeconomics class should be able to answer using the current state of knowledge. Yet, our experience over the years is that students are flummoxed by these issues. Undergraduates mostly retain that central banks print money and money means more inflation. They are then thoroughly confused when they realize that most central banks barely mention money in their speeches, that they do not actually choose how much money to print, and that the monetary base has increased many-fold with no appreciable inflation in the last decade. Graduate students learn about the setting of interest rates and about inflation-output trade-offs mapped out by the Phillips curve, and perhaps even about the welfare costs of inflation and positive interest rates. But as soon as you ask them to reconcile the Fisher principle—higher interest rates are associated with higher inflation one to one—the Taylor rule—increasing interest rates more than one-to-one in response to inflation keeps inflation constant—and the empirical evidence—raising interest rates lowers inflation—you are likely to get an incoherent answer. And, if you ask what role does fiscal policy play in all this, almost surely expect a blank stare.

The reason for this state of affairs is that there are few, if any, accessible entry points to this literature. Reading the journal articles, a student gets consumed by debates about the role of transversality conditions, the monetarist or fiscal theories of the price level, or the microfoundations of money, and finds it difficult to see that there is actually a well established theory of how the prive level is pinned down. The goal of this article is to provide a pedagogical survey of the literature on price-level determination. The approach that we will take is to highlight the common features of different view points by using a single neoclassical general-equilibrium model of the economy, noting that different theories simply focus on different equations and markets within the same model. We will focus on the central bank as the agency with a mandate to deliver a value for the price level, but without neglecting the interaction with the fiscal authorities.

This is not an article about the optimal way to conduct monetary policy or about how to trade off variability in inflation or real activity. It is a survey of the *process* of monetary policy, that is the way in which given a target for inflation, the central bank goes about delivering. Letting P_t^* denote this target for the price level, the question is how to choose central bank actions to deliver an actual price level P_t as close as possible to the target. To make the question stark, we take this target as given. While it is exogenous to the model in this paper, the target may well be stochastic, have a unit root, or depend on other variables, as in the dominant interpretations. Whether by setting some policy tool, the central bank can ensure that in equilibrium there is a single P_t that is as close as possible to P_t^*, involves two questions. The first is a *determinacy* question, on whether a process can uniquely deliver a price level using its tools. The second is an *effectiveness* question, on whether that process leads to a lower deviation between the actual and the target price level.

The paper is structured as follows. Section 2.2 sets up the canonical dynamic model to show that classical analysis of supply and demand does not pin down the price level. It introduces a basic description of what is a central bank and its actions, which will be expanded as the paper moves along.

Sections 2.3-2.6 discuss four classes of processes to control the price level. Each of them focuses on a different equilibrium condition in a particular market, and a different policy tool of the menu available to the central bank. While they may appear distinct, we show that they all refer to the same model, but refer to different strategies by the central bank.

The first of these, in section 2.3 focuses on financial markets, the arbitrage conditions that describe equilibrium in them, and consequently on the central bank's setting of the interest on the deposits that banks hold at the central bank. The second approach in section 2.4 looks at equilibrium in the market for currency, and the policy tool is the monopoly power of the central bank to print money. Section 2.5 describes a third approach, where the need for fiscal resources by the government plays central role, and the policy tool is the size of the balance sheet of the central bank. Next comes section 2.6 where the key market is that for different goods, so the equilibrium condition is the law of one price, and the central bank policy tool is the definition of the unit of account. Finally, section 2.7 considers a fifth approach that allows for the study of nominal rigidities.

Section 2.8 provides some tentative conclusions. Future revisions of this draft will include also a hybrid processes for inflation control, and a empirical discussion of the effectiveness of each approach based on the size of the control errors they give rise to.

2.2 Inflation in equilibrium

At the core of most dynamic macroeconomic models is an Euler equation of the form:

$$\mathbb{E}_t \left[M_{t+1}(1 + R_t) \right] = 1. \tag{2.1}$$

The operator $\mathbb{E}_t(.)$ captures the private-sector expectations of the future as of date t. These need not be rational; we only assume that they are a valid expectations operator and satisfy the law of iterated expectations. R_t is the promised return at date t on a real investment that pays off at date $t + 1$, while M_{t+1} is a stochastic discount factor.

The economic intuition behind this equation is that M_{t+1} reveals how many goods next period the private agents would require in exchange for one unit of goods today. In other words, M_{t+1} is the marginal rate of substitution between consumption today and tomorrow. Since $1+R_t$ is the opportunity cost of consuming one more unit today in terms of consumption foregone tomorrow, then the equation above is the usual optimality condition that agents must be indifferent between consuming or saving one more unit.

An alternative, but equivalent, investment intuition is that for there to not be arbitrage profits, it must be that the risk and time adjusted net return on any investment is zero. The stochastic discount-factor gives the adjustment factor for time and risk. If investors are risk neutral then M_{t+1} would be equal to a constant β that captures solely impatience, and the equation states that the real return is approximately equal to the rate of time preference $-\ln(\beta)$.

A strong assumption for all of this paper is the following: *M_{t+1} is exogenous.* This assumption is commonly known as the classical dichotomy. It states that no matter what the price level happens to be, real tradeoffs are unchanged. This allows the paper to focus on determining the price level.

An *equilibrium* in this economy is then a solution for $\{R_t, P_t\}_{t=0}^{\infty}$ such that given an exogenous $\{M_{t+1}\}_{t=0}^{\infty}$, equation (2.1) holds.

2.2.1 General equilibrium micro-foundations

To see the central role of the Euler equation, consider a simple exchange economy economy populated by many private agents that have the same time-separable preferences over a single consumption good $\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t)$ and that can trade bonds

K_t with each other subject to a budget constraint: $P_t C_t + P_t K_{t+1} \leq P_t Y_t + P_t K_t (1 + R_{t-1})$. There is no capital or any storage technology and the bonds are in zero net supply, so market clearing imposes that consumption is equal to aggregate output $C_t = Y_t$. Aggregate output is a random endowment that just falls from the sky. Optimal behavior in this economy is then entirely described by equation (2.1) where the stochastic discount factor is equal to: $M_{t+1} = \beta U'(Y_{t+1})/U'(Y_t)$, which is exogenous.

2.2.2 Price level (in)determinacy

In equilibrium, the real interest rate is given by:

$$R_t = \mathbb{E}_t \left[M_{t+1} \right]^{-1} - 1, \tag{2.2}$$

at all dates. But nothing pins down the price level. Any sequence $\{P_t\}_{t=0}^{\infty}$ is consistent with the Euler equation holding and the economy being in equilibrium.

More formally, let s_t be the state of the world at date $t \geq 1$, and let $s^t = (P_0, s_1, ..., s_t)$ be the history of states until date t. Define then inflation as $\Pi(s^t) \equiv P_t/P_{t-1}$, the increase in the price level. A nominal equilibrium is then an initial value P_0 and a function $\Pi(s^t)$ for all dates $t \geq 1$.

The level of inflation is unique or determinate in equilibrium if:

1. There is a unique scalar P_0 in equilibrium.
2. If $\Pi'(s^t)$ and $\Pi''(s^t)$ both satisfy equilibrium conditions, then $\Pi'(s^t) = \Pi''(s^t)$.

The first condition requires that even if the entire future of inflation is pinned down from today onward, still one must know what today's price level is. What ultimately matters is what is a dollar's worth in real goods at any given date. Without pinning down the initial value of the dollar, then for a given inflation path, the actual price level P_t could be any number.

The second condition states that, while if states of the world are different, the inflation may be different, for a given state of the world, inflation must be unique. If, in spite of all fundamental features of the world being the same, inflation can be different, then the central bank has failed to pin down inflation.

The result that in equilibrium, without a central bank, then inflation is indeterminate holds in *any* classical model. The reason dates back to Hume (1752): dollars are just a unit of account with which the prices of goods are determined. If people started denominating prices in cents instead of dollars nothing would change. There is no demand or supply that pins down that 100 cents equals one dollar. Nothing in classical economics pins down the price level or inflation, in the same way that nothing in it determines whether measurements should be in inches or centimeters.

2.2.3 Nominal bonds do not solve indeterminacy

In the simple model, neither consumption nor savings are denominated in dollars. Perhaps this is what leads the value of the dollar to be indeterminate. To see if it is so, introduce nominal bonds. If these bonds exist, then there is a no-arbitrage relation between real and nominal bonds:

$$\mathbb{E}_t \left[M_{t+1} \left(1 + R_t - \frac{1 + I_t}{\Pi_{t+1}} \right) \right] = 0, \tag{2.3}$$

This states that, once adjusted by the stochastic discount factor, saving in real or nominal bonds must yield the same expected return. It is often called the Fisher equation. It can easily be micro-founded by letting the representative agent choose whether to save in nominal or real bonds, in which case this is the indifference condition between the two that holds at the optimum.

Introducing nominal bonds does not solve price level indeterminacy. While this gives an additional equilibrium condition, there is also an additional endogenous variable, I_t, that must be pinned down. Using the Fisher equation to pin down what I_t will be, given Π_t, one is still left with no equations to pin down neither the initial price level nor inflation thereafter.

2.2.4 Introducing a central bank

Modern central banks perform many roles. We start by describing a crucial yet minimal such role and throughout the paper build it up with more of the features observed in the world.

In modern digital economies, people use electronic means of payment like debit or credit cards to settle their transactions. Because for any given transaction the seller may have an account in bank A, and the buyer an account with bank B, there must be a settlement whereby bank A collects payment for bank B. The central bank is the clearing house where the payments between banks take place. These payments are made using another digital mean of payment, often named reserves, which are nothing but liabilities of the central bank towards banks.[1] Because reserves are the ultimate form of payment, they are the unit of account. That is, because reserves are denominated in dollars, and people choose to denominate their prices in dollars as well. The price of a good is simply how many units of reserves must be given to obtain that good.

The current stock of reserves is just a list of entries in a spreadsheet at the central bank, one for each bank. Because the central bank has sole control over the spreadsheet, it can freely choose how to remunerate these reserves, that is at what rate does it increase the number in each entry in the spreadsheet. In this way,

[1] For people that prefer to settle some transactions without using a bank (a minority today) the central bank also issues banknotes, a particular durable good, and commits to exchange them for reserves one-to-one at all times. We will discuss currency in section 2.4.

reserves are just like a nominal bond, but one for which I_t is *chosen* by the central bank. Because of its role as the controller of the spreadsheet for the unit of account, the central bank has one tool that it can use for policy: the interest rate at which it remunerates the reserves $\{I_t^v\}_{t=0}^{\infty}$.

To understand the power of doing so, just note that now the Fisher equation provides an equation to solve for inflation given that now R_t, M_t, as well as I_t^v are exogenous. This equation is:

$$\mathbb{E}_t\left(\frac{M_{t+1}}{\Pi_{t+1}}\right) = \frac{1}{1+I_t^v}. \tag{2.4}$$

Expected (SDF-adjusted) inflation is now determined.

However, *actual* inflation is not pinned down. The mean of the random variable $\Pi(s^{t+1})$ may be pinned down, but its distribution is not. There is an infinite number of possible inflations at different sets of the world that satisfy this equation, violating the second requirement for determinacy of the price level. As thoroughly discussed by Nakajima and Polemarchakis (2005), the actual price level may turn out to be arbitrarily far from this expectation, as long as people's future expectations are sufficiently uncertain.

Moreover, the first requirement for uniqueness of inflation is also not satisfied. If there is no uncertainty in the economy, the expectations operator disappears from the equation above, and indeed there is a single Π_{t+1} at each date. However, there is no other condition for pin down P_0. If people expect higher prices in the future, the price level will simply jump up today, keeping inflation form date 1 onward the same, but the real value of money indeterminate. This indeterminacy was made famous by Sargent and Wallace (1975).

The next four sections show that the central bank can go further though and indeed fully pin down inflation. Each of them is a different (mutually exclusive) policy option by the central bank. Namely, the central bank can choose to control: the remuneration it offers on its reserves, the quantity of reserves it issues, the quantity of banknotes it issues promising to exchange them for reserves, and what to buy with these reserves. This gives rise to the four different approaches that we discuss from now onward.

2.2.5 Fiscal policy

Before doing so though, one must complete the description of the model. No model of economic policy is complete without specifying how the central bank interacts with the fiscal authority. As part of their actions, central banks may earn income and rebate dividends to the fiscal authorities. For now, we assume that the fiscal authority issues no bonds and has some exogenous spending programs, so it endogenously picks taxes in order to pay for its spending minus the dividends it receives from the central bank. This is sometimes called a Ricardian policy, and it implies that the treasury fiscally backs the central back, accepting whatever dividends it receives even if they are negative.

2.3 No arbitrage approach: remunerating reserves

Banks can choose to hold reserves or real investments. Imagine that the expected price level tomorrow is higher. Then, the return on reserves would be smaller than that on real investments. The dollars that the central bank pays to the holders of the reserves can afford fewer goods than the ones that holding real investment today would give access to. Banks would want to hold zero reserves and invest all of their resources in real terms if this persisted, which would not be an equilibrium given the positive supply of reserves. Rather, as they demand fewer reserves, their value must rise. Because the real value of the reserves is $1/P_t$ since, after all, they are the unit of account, this means the price level must rise. A higher price level means that the real return on reserves is now higher, rising until the point where people are indifferent again.

The forces of no-arbitrage applied to central bank reserves therefore move the equilibrium price level today. The equation capturing these forces is the Fisher equation.[2] The policy tool at the disposal of the central bank is its ability to choose how to remunerate these reserves. The process for monetary policy in this case consists of specifying how to remunerate reserves appropriately in order to pin down inflation on target. We start with this policy strategy because it is the one followed by most central banks today.

2.3.1 A simple but unused strategy: real payments on reserves

To begin though, we start with a policy process that no central bank that we know of uses, but which makes the forces of arbitrage very transparent at pinning down inflation. Imagine the central bank promises to remunerate reserve holders with a payment in real goods of X_t, as suggested by Hall and Reis (2017). Governments have issued indexed bonds for a long time across the world, and so could (and maybe should) central banks, which is what promising a real payment amounts to. The nominal return on reserves in dollars would then be $I^v_{t,t+1} = (1 + X_t)P_{t+1}$.

Rearranging equation (2.4) and using the result in equation (2.2) then delivers:

$$\mathbb{E}_t \left[M_{t+1} \left(1 + R_t - \frac{(1 + X_t)P_{t+1}P_t}{P_{t+1}} \right) \right] = 0 \;\Leftrightarrow\; P_t = \frac{1 + R_t}{1 + X_t}\,. \tag{2.5}$$

[2]This equation assumes that reserves are a pure financial asset that provides no payment or liquidity services, so they are valued purely for their returns. Reis (2016) estimates that the market for reserves has been saturated in the United States since approximately 2011, once reserves started exceeding $1 trillion.

Since X_t is exogenously chosen by policy, and R_t is exogenously pinned down by real forces, then the above equation delivers a determinate price level.

The central bank can then follow a process for this payment: $1+X_t = (1+\hat{R}_t)/P_t^*$, given its current estimate of the current real interest rate $\hat{R}_t$, and the target for the price level. To clarify the notation used, $\mathbb{E}_t(P_{t+j})$ is the public's expectation at t of what the price level will be at date $t+j$, while $\hat{P}_{t+j}$ is the central bank's expectation at t. We do not assume rational expectations, so these may not be the same. Also, we use small letters to denote log-linear deviations from a constant steady state, so $r_t = \ln((1+R_t)\beta)$.

Give this process, the actual price level is:

$$p_t = p_t^* + \varepsilon_t \text{ with } \varepsilon_t = r_t - \hat{r}_t \,. \tag{2.6}$$

Insofar as the central bank makes mistakes assessing the state of the economy, in the form of the current real interest rate, the actual price level will deviate from its target.

The intuition for how the price level is pinned down is the following. The real return on any investment is pinned down exogenously by the stochastic discount factor. If the central bank promises a real payment on reserves, then arbitrage pins down how many goods reserves are worth today. This is the economic force behind the Fisher equation: since real bonds and reserves both deliver the same payment tomorrow, they must be worth the same today. But, since reserves are denominated in dollars, not goods, then this pins down the price level today.

2.3.2 A Wicksellian rule

Most central banks instead announce a nominal interest rate on their reserves. Now reserves and bonds are equivalent assets: each is denominated in dollars and promises a risk-free payment in dollars.[3] Therefore, by no arbitrage it must be that $I_t^v = I_t$.

It is convenient in this section to work with a log-linearized version of the Fisher equation around a steady-state point where inflation is constant $P_{t+1}/P_t = \Pi$. The Fisher equation then becomes:

$$i_t - \mathbb{E}_t(\Delta p_{t+1}) = r_t. \tag{2.7}$$

We already saw that simply setting i_t^v, and thus i_t as it wishes, does not pin down inflation. However, imagine the central bank sets the interest rate according to a feedback rule. In particular, consider a Wicksellian interest-rate process as described

[3] Again this assumes that the market for reserves is saturated. Otherwise, the interest on reserves would only put a floor on short-term interest rates, since lending to the central bank carries the least default risk in the economy and the maximum liquidity. The short-term interest rate can still be set by the central bank but now only by both setting the interest on reserves and controlling the amount of reserves through constant open-market operations.

in Woodford (2011a):

$$i_t = \hat{r}_t + \hat{p}^*_{t+1} - p^*_t + \phi(p_t - p^*_t), \tag{2.8}$$

where $\phi > 0$.

Combining the process with the Fisher equation (2.7) gives a difference equation for the price level:

$$(1+\phi)p_t - p^*_t = (r_t - \hat{r}_t) + \mathbb{E}_t(p^*_{t+1} - \hat{p}^*_{t+1}). \tag{2.9}$$

Iterating forward and imposing the terminal condition:

$$\lim_{T\to\infty} (1+\phi)^{-T}\, \mathbb{E}_t\left(p_{t+T} - p^*_{t+T}\right) = 0\,, \tag{2.10}$$

then pins down the price level at the target $p_t = p^*_t + \varepsilon_t$, where the error is given by:

$$\varepsilon_t = \sum_{s=0}^{\infty} (1+\phi)^{-s-1}\, \mathbb{E}_t\left[r_{t+s} - \hat{r}_{t+s} + p^*_{t+1+s} - \hat{p}^*_{t+1+s}\right]. \tag{2.11}$$

Inflation is now determinate. For a realization of shocks, there is a single inflation rate, and the initial price level is pinned down as well. If $\varepsilon_t = 0$, then this feedback process is observationally equivalent to a peg. Yet, the mere presence of $\phi > 0$ solves the two degrees of indeterminacy discussed in the previous section.

Under this process for monetary policy, expectations about future policy are as important as current policy. With the payment on reserves rule, only the current promised payment next period pinned down prices, with this feedback rule, it is the iteration of interest rates until the infinite future that determines the price level. Announced forward guidance by the central bank about changes in the target for prices in the future p^*_{t+1+s} or in the expected state of the economy in the future $\hat{r}_{t+s}$ both affect prices right away today.[4]

2.3.3 A Taylor rule

The more common feedback rule, which many central banks admit to follow is that proposed by Taylor (1993):

$$i_t = \hat{r}_t + \hat{p}^*_{t+1} - p^*_t + \phi(\Delta p_t - \Delta p^*_t)\,. \tag{2.12}$$

Interest rates set by the central bank now respond to deviations of inflation from target, as opposed to the price level. The condition on the strength of that response is now $\phi > 1$, the so-called Taylor principle whereby the central bank raises nominal

[4]These two types of announcements have been called Odysseian and Delphic, respectively, and have been quantitatively criticized because they can have large effects right now even when the future is distant.

interest rates more than one-to-one if inflation deviates from target.

Similar steps of algebra now lead to:

$$\Delta p_t = \Delta p_t^* + \sum_{s=0}^{\infty} \phi^{-s-1}\, \mathbb{E}_t \left[r_{t+s} - \hat{r}_{t+s} + p^*_{t+1+s} - \hat{p}^*_{t+1+s} \right]. \tag{2.13}$$

Since this holds for all $t \geq 0$, it also pins down p_0 for any given p_{-1}, which is given by history. Thus both degrees of determinacy are satisfied.

An enormous literature has used this rule as the process for monetary policy following Taylor (1993)'s demonstration that it seemingly fits the U.S. experience well.

The intuition for why the Taylor rule works goes as follows. Imagine that inflation is higher than target at date t by one log unit. Then, the central bank will raise the nominal interest rate by ϕ leading to an increases in expected inflation between t and $t+1$ of ϕ. But this in turn leads the central bank to raise i_{t+1} by ϕ^2 which raises expected inflation between $t+1$ and $t+2$ by that amount. The process continues so inflation keeps on rising exponentially and the feedback rule delivers inflation in T periods to be larger by ϕ^T.

If this is possible, inflation is not pinned down. Subject to this path for future prices, the current price level can be one unit, or any other arbitrary amount higher or lower; p_0 is still indeterminate. Future stochastic inflation is determinate, as the equations for ε_t show, since there is unique map from the exogenous shocks to the endogenous realizations of inflation. The degree of indeterminacy is therefore solely on the initial price level. Feedback rules by themselves did not solve the indeterminacy of the initial price level. What did it was instead the terminal condition.

2.3.4 The elusive terminal condition

For the Taylor rule, the terminal condition was:

$$\lim_{T \to +\infty} \phi^{-T}\, \mathbb{E}_t \left(\Delta p_{t+T} - \Delta p^*_{t+T} \right) = 0. \tag{2.14}$$

It was this condition that ruled out any price level that differs from the target by anything but the errors in ε. A single path for the price level satisfies this condition, so determinacy is achieved.

The Taylor principle that $\phi > 1$ is important because it imposes that the random variable defined by $\mathbb{E}_t \left(\Delta p_{t+T} - \Delta p^*_{t+T} \right)$ belongs to $O(\phi)$. That is, if expected inflation deviates from target, those deviations cannot grow faster than at the rate ϕ. If, for instance, such deviations are bounded, which is coherent with the approximation error in the log-linearization being small, then this condition will hold.

Feedback rules associate any indeterminacy of the initial price level to explosive paths for inflation from then onwards. One can argue that such explosions are implausible and that people would never believe them. In the literature, equilibria in which explosive beliefs are ruled out are called locally bounded equilibria. If one thinks that these are the only plausible equilibria, then this gives rise to the terminal condition above. Among the set of bounded equilibria, inflation will be uniquely pinned down.

Yet, where does this equilibrium restriction or terminal condition come from? It is not a transversality condition. Those apply to the real value of savings, whereas the condition needed here is on a purely nominal variable, the price level. In the same way that optimal behavior imposed no money illusion in the Euler equation, it also imposes no money illusion in transversality conditions. Also, there is no sense in which the economy blows up if this condition does not hold. The unit of account may be exploding, but agents with no money illusion do not care for this at all. All real outcomes and variables continue to be finite.

Cochrane (2011) provides a scathing critique of the use of this bounded assumptions. A large literature has debated alternative assumptions that mostly fit into three lines of argument.

Escape clauses

The idea of an escape clause originates with Obstfeld and Rogoff (1983). The central bank could commit to a feedback process only while inflation does not go on an unbounded explosive path. Otherwise, if inflation exceeds a pre-announced threshold, the central bank would switch to another policy process (a money rule as in Taylor (1996) and Christiano and Rostagno (2001)). These other rules would then be able to pin down the price level at the date of that switch at a finite level.

More formally, the policy strategy is now to follow the feedback interest rate rule while inflation is within some interval $\left[\pi^L, \pi^H\right]$, but if at some date T, inflation π_T falls outside this interval, then choose another policy process that at date T pins down inflation, and so determines uniquely p_T at the desired target p^*_T perhaps with an error. For instance, this might be the payment on reserves rule that we discussed already.

In this case, iterating on the difference equation for the price level in equation (2.9), in the case of the Wicksellian rule, now delivers:

$$p_t = p^*_t + \sum_{s=0}^{T-1} (1+\phi)^{-s-1} \mathbb{E}_t \left[r_{t+s} - \hat{r}_{t+s} + p^*_{t+1+s} - \hat{p}^*_{t+1+s}\right] + (1+\phi)^{-T} \mathbb{E}_t \left(p_{t+T} - p^*_{t+T}\right) \tag{2.15}$$

Because the last term on the right-hand side is now uniquely pinned down by the regime switch, then the price level of the left-hand side is uniquely pinned down as well.

This approach to pinning down the price level is valid and perhaps even realistic. If inflation was rising without bound, no central bank would stick to blindly following a feedback rule that tells it to raise interest rate more and more, even as it sees inflation rising faster and faster. Switching to one of the other policy strategies described in this paper, might well be what would happen, and even if the error in that was very large, as long as it delivers uniqueness, the feedback rule would uniquely pin down the price level in all other periods.

Theoretically though, in the sense of answering: "how does the central bank control inflation?" this is not entirely satisfactory, as it relies on the other policy processes to do it. At the extreme, if either the width of the interval $\left[\pi^L, \pi^H\right]$ goes to zero, or the errors ε_t are large enough, the economy would spend essentially no time with the feedback rule after $t = 1$.

Equilibrium refinements

An alternative approach also relies on regime switches but these are now not providing a terminal price level, but rather making an off-equilibrium threat that ensures the regime switch never happens. Much of this work builds on Bassetto (2005), and includes Atkeson et al. (2010)'s sophisticated and Christiano and Takahashi (2018)'s strategy equilibria.

The policy process is committed to a Taylor rule while inflation stays in a bounded interval $\left[\pi^L, \pi^H\right]$. But if π_T falls outside this interval, there is switch in process in the next period at $T + 1$. This new process is able to uniquely pin down inflation π_{T+1}, but differently from before, this now happens to some level well inside the interval, and in particular to a level such that $\pi_{T+1} < \pi^H - r_t$.

Therefore, from the Fisher equation (2.7) the regime switch pins down $i_T = \pi_{T+1} + r_t < \pi^H$. Yet, at T, the Taylor rule implies that since π_T was larger than π^H, and the Taylor rule coefficient is larger than one, then $i_T > \pi^H$. This is a contradiction. The only way to avoid it is for inflation to never leave the bounded interval. If the width of the interval is large enough such that the size of the exogenous shocks would never send the economy there, then what is ruled out are the explosions that led to indeterminacy with a Taylor rule. Because the feedback rule implies that inflation explodes at rate ϕ^{-1}, then one of the bounds will be reached for sure in finite time for any inflation level but for the one that satisfies the elusive terminal condition that we used before.

Christiano and Takahashi (2018) emphasize that the process to which policy switches must credibly deliver indeed a low inflation from then onwards. This may require for instance that r_t is endogenous to make sure that the key inequality $\pi_{T+1} + r_t < \pi^H$ is reached.

Non-rational expectations

A third argument to justify a terminal condition in equation (2.14) focuses on the expectations operator. What the central bank needs in order to pin down inflation is for people not to start expecting that inflation in an arbitrary far away future is growing (or falling) at an explosive rate. Therefore, assumptions on how expectations of these far away events are formed can deliver determinacy. Woodford (2013) reviews macroeconomic models with non-rational expectations. They involve two related concepts.

First is the concept of a temporary equilibrium, defined as a competitive equilibrium at each date in time that depends on an exogenous set of subjective expectations. One can write the temporary equilibrium as a map $\Gamma(\cdot)$ from the sum of the subjective expectations of future endogenous variables, e_t, to actual outcomes: $\pi_t = \Gamma(e_t)$. For our simple model, combine a generalized version of the Fisher equation (2.7) that reads:

$$\underbrace{\sum_{s=0}^{\infty} \beta^s r_{t+s}}_{\equiv rr_t} = i_t + \underbrace{\sum_{s=0}^{\infty} \beta^s E_t^* \left(\beta i_{t+s+1} - \pi_{t+s+1}\right)}_{\equiv e_t}, \tag{2.16}$$

with the Taylor rule as given by equation (2.12). $\Gamma(\cdot)$ is a linear operator such that:

$$\pi_t = \frac{1}{\phi}\left(rr_t - \bar{\imath}_t\right) - \frac{e_t}{\phi}, \tag{2.17}$$

where $\bar{\imath}_t \equiv \hat{r}_t + \hat{p}^*_{t+1} - p^*_t - \phi\pi^*_t$ is a set of exogenous variables that do not depend on expectations.[5]

Second is the revision process by which expectations are updated. This often requires specifying first a map $\Psi(\cdot)$ from the conjectured set of expectations, e_t, to expectations that would be correct if people acted under the conjectured expectations: $e^*_t = \Psi(e_t)$; as in the "calculation equilibrium" of Evans and Ramey (1992, 1995, 1998) This is simply derived by plugging the temporary equilibrium relation into the definition of subjective expectations,

$$e^*_t = \sum_{s=0}^{\infty} \beta^s E^*_t \left\{ \left(\beta - \frac{1}{\phi}\right) rr_{t+s+1} + \frac{1}{\phi}\bar{\imath}_{t+s+1} - \left(\beta - \frac{1}{\phi}\right) e_{t+s+1} \right\}. \tag{2.18}$$

Note that the non-explosive rational expectations equilibrium is defined as the fixed point of this problem.

Let k be the stage in the revision process so that $e(k)$ are the expectations at this stage. Different expectations models will involve different law of motions. For example, reflective expectations, introduced by García-Schmidt and Woodford

[5] The standard first-order difference equation in inflation that results from combining the baseline Fisher equation and the Taylor rule:

$$\pi_t = \frac{1}{\phi}\left(r_t - \bar{\imath}_t\right) + \frac{1}{\phi} E^*_t \pi_{t+1},$$

is an alternative representation of this temporary equilibrium.

(2019), feature a continuous specification:

$$\frac{\partial e_t(k)}{\partial k} = e_t^*(k) - e_t(k) = \Psi\left(e_t(k)\right) - e_t(k)\,, \tag{2.19}$$

while the k-level thinking literature, as in Farhi and Werning (2018), delivers the discrete counterpart:

$$e_t(k) = \Psi^k\left(e_t(0)\right)\,. \tag{2.20}$$

Both models are closely related and are examples of eduction. Moreover, if the levels of thinking are assumed to follow a Poisson distribution, both models are observationally equivalent.[6] The main difference, however, lies in the interpretation. Reflection is concerned with average beliefs allowing individuals to form expectations individually, whereas in level-k thinking, all agents are assumed to be at the same stage of the iteration process.

A third class of expectations models that fit into a similar setup are models of least squares learning like McCallum (2009) and Evans and McGough (2018), which are instead based on experience. These models assume that agents behave as statisticians when forecasting the future value of variables and adjust expectations as new data becomes available:

$$e_t = \chi\left(\pi_{t-1}, \pi_{t-2}, \dots\right)\,, \tag{2.21}$$

where inflation dynamics are described by $\pi_t = \Gamma\left(\chi\left(\pi_{t-1}, \pi_{t-2}, \dots\right)\right)$. Agents are assumed to estimate the parameters of a system which simultaneously depends on those estimates. The forecasting rule $\chi(\cdot)$ follows from the standard least squares formula.

Next, we show a simple example in which the previously described models of expectations converge to the non-explosive rational expectations equilibrium in the limit.

Proposition 3. *Consider a stationary environment, in which $r_t = 0$ for all t, and the monetary policy rule intercept is fixed such that $\bar{\imath}_t = \bar{\imath}$ for all t. Moreover, expectations of future endogenous variables are constant at all horizons and given by (π_e, i_e).*

(i) Reflective expectations converge to bounded rational expectations asymptotically.

(ii) Level-k thinking converges to bounded rational expectations asymptotically.

(iii) Least-squares learning converges to the bounded rational expectations asymptotically.

[6] See García-Schmidt and Woodford (2015) for the formal argument.

Proof. The non-explosive rational expectations equilibrium is given by $\pi^{RE} = -\frac{1}{\phi-1}\bar{\imath}$ and the corresponding belief sequence:

$$e^{RE} = \frac{1}{1-\beta}\left[\beta i^e - \pi^e\right] = \frac{1}{1-\beta}\left[\beta\bar{\imath} + (\beta\phi_\pi - 1)\pi\right] = \frac{1}{\phi_\pi - 1}\bar{\imath}.$$

To prove part *(i)*, we show that belief revision dynamics represented by (2.19) converges as k grows boundlessly to the belief sequence associated with the non-explosive rational expectations equilibrium. Given the above assumptions equation (2.19) reads:

$$\dot{e} = \frac{1}{1-\beta}\left[\frac{1}{\phi_\pi}\bar{\imath} - \left(1 - \frac{1}{\phi_\pi}\right)e\right], \tag{2.22}$$

which has a stable solution at $\bar{e} = e^{RE}$ provided that $\phi > 1$ i.e. the Taylor principle holds.[7]

To prove part *(ii)*, we show that belief revision dynamics represented by (2.20) converges as k grows boundlessly to the belief sequence associated with the non-explosive rational expectations equilibrium. For a finite k, equation (2.20) can be rewritten as:

$$e(k) = \frac{1}{1-\beta}\frac{1}{\phi}\bar{\imath}\sum_{s=0}^{k-1}\left(\frac{1-\beta\phi}{\phi}\frac{1}{1-\beta}\right) + \left(\frac{1-\beta\phi}{\phi}\frac{1}{1-\beta}\right)^k e(0)\,. \tag{2.23}$$

Thus, as $k \to \infty$, it follows that $e(k) \to \frac{1}{\phi_\pi - 1}\bar{\imath}$ as long as $\phi > 1$ and $e(0)$ is finite.

To prove part *(iii)*, we make use of the e-stability principle that states that the the mapping from the agent's perceived law of motion (PLM), based on the econometric specification, to the actual law of motion (ALM), which results from the temporary equilibrium, governs the stability of equilibria under learning. The corresponding mapping here is:

$$\dot{\theta} = \frac{1}{\phi}(\theta - \bar{\imath}) - \theta\,, \tag{2.24}$$

which solves as $\bar{\theta} = -\frac{1}{\phi-1}\bar{\imath}$, implying $\pi_t = \pi^{RE}$ given the least-squares learning specification.[8] E-stability is guaranteed given $\phi > 1$. □

[7] Recall the stability theorem: let $\frac{dx}{dy} = f(x)$ be an autonomous differential equation and suppose $x(t) = x^*$ is an equilibrium. Then, if $f'(x^*) < 0$, x^* is stable.

[8] To obtain the mapping above, note that the PLM is given by:

$$\pi_t = \theta\,, \tag{2.25}$$

and plugging this into the temporary equilibrium in its first-order difference equation version delivers the ALM:

$$\pi_t = \frac{1}{\phi}(\theta - \bar{\imath})\,. \tag{2.26}$$

2.3.5 Long-term interest rates and forward guidance

Central banks sometimes issue bonds, together with their traditional reserves. Whereas the latter are central bank liabilities that are used as the unit of account and to settle digital transactions that involve banks, the former are standard bonds of a fixed maturity that are then paid off with reserves. In the same way that the central bank chooses the way in which it remunerates reserves, it can potentially choose how to remunerate these bonds. In doing so, the central bank is choosing not just a short-term nominal interest rate, but also longer-term ones. This provides an alternative way to try to control inflation (Adão et al., 2014; Magill and Quinzii, 2014).

If the central bank issues a j period bond and pays I_t^j interest rate on it, then the Euler equation that applies to this new form of investment is:

$$\mathbb{E}_t \left[\frac{M_{t+1} M_{t+2} ... M_{t+j} (1 + I_t^j)}{\Pi_{t+1} \Pi_{t+2} ... \Pi_{t+j}} \right] = 1. \tag{2.27}$$

Since the central bank chooses I_t^j and the M_{t+j} are exogenous, then this provides one more set of equations, one for each date t, with which to try to determine the sequence of inflation rates over time. Increasing the number of equations without increasing the number of unknowns gives hope that perhaps inflation is now determinate. If the central bank issues many such bonds of different maturities, then it creates more equilibrium conditions with which perhaps it can control inflation.

To see this at play, consider the simple case in which there is only uncertainty about M_{t+1} and that this follows a two-state stationary Markov chain with values M_H and M_L and transition matrix with probabilities that are non-negative and satisfy $B_{HH} + B_{HL} = 1$ and $B_{LH} + B_{LL} = 1$. Controlling inflation is determining the two values of inflation Π_H and Π_L uniquely for each of the two states. The Euler equations with respect to the one period reserves and the two period bonds can be written at state s as:

$$(1 + I_s^1) \left(B_{sH} \frac{M_H}{\Pi_H} + B_{sL} \frac{M_L}{\Pi_L} \right) = 1, \tag{2.28}$$

$$(1 + I_s^2) \left(B_{sH} \frac{M_H}{\Pi_H (1 + I_H^1)} + B_{sL} \frac{M_L}{\Pi_L (1 + I_L^1)} \right) = 1. \tag{2.29}$$

These are two equations in two unknowns. Standard linear algebra shoes that as long as $I_H^1 \neq I_L^1$, then there is a unique solution for inflation. The key condition for determinacy is now that the central bank does not set the interest on reserves to be the same across states of the world.

Note that this approach does not pin down P_0. Only the stochastic degree of indeterminacy disappears. Intuitively, not just the mean of inflation but also how it covaries with the stochastic discount factor across two successive periods are now pinned down by arbitrage. Thus, the indeterminacy of inflation across states of the world can be solved as long as the nominal interest rate varies with those states of the world. However, while these interest rates are varying over states, over time they are still pegged in the sense of section 2.2. Thus, the problem of controlling P_0 remains.

No central bank that we are aware of has ever used this monetary policy process. At the same time, central banks since 2010 have “gone long" by using unconventional monetary policies, like quantitative easing, to exert control over long-term interest rates, approximating the behavior just described. The major advanced-economy central banks did not issue bonds, but they bought and sold government bonds in order to exert some control over their interest rate through quantitative easing policies.

Another way to achieve the same goal is to have forward guidance, whereby the central bank uses speeches and announcements to control not just the current interest on reserves, but also what people expect it to be. In fact, similar steps to before, show that if the central bank announces both its current interest rates on reserves, as well as its expected value for tomorrow, this again provides two equation with which to solve for inflation across states.

2.3.6 Global analysis, banknotes, and the zero lower bound

A distinctive feature of the previous section is that it did not log-linearize the equilibrium conditions. The previous log-linearizations in the study of feedback rules were not crucial to the conclusions. To simplify the problem, eliminate uncertainty on the real interest rate so the stochastic discount factor is now constant at β, and there is only possibly indeterminacy with respect to the initial price level.

The Taylor monetary policy process then implies that:

$$\frac{P_{t+1}}{P_t \beta} = \frac{P^*_{t+1}}{P^*_t \beta} \left(\frac{P_t P^*_{t-1}}{P_{t-1} P^*_t} \right)^{\phi}. \tag{2.30}$$

The left hand side is the nominal interest rate given by the Fisher equation, and the right-hand side is the non-linear Taylor rule with $\phi > 1$. Letting $\tilde{\Pi}_{t+1} = (P_{t+1}/P_t)/(P^*_{t+1}/P^*_t)$ be the deviations of gross inflation from target, this simplifies to just $\tilde{\Pi}_{t+1} = \tilde{\Pi}^{\phi}_t$, a nonlinear difference equation. Taking logs though, one gets precisely the same dynamics as in the log-linearized case. If inflation starts on target, when $\tilde{\Pi} = 1$ it stays there forever. If it deviates upwards or downwards, then this leads log inflation to explode to plus or minus infinity. Again a limit condition that rules out these explosions will pin down inflation.[9]

This global analysis shows that with negative deviations of inflation from target, at some point inflation will be so negative that the nominal interest rate must be

[9] It is easy to show that the same with the Wicksellian case: $P_{t+1}/P^*_{t+1} = (P_t/P^*_t)^{1+\phi}$ which as an identical representation to Figure 2.1, but with deviations of the price level from target as opposed to inflation.

negative as well. Yet, central banks together with digital reserves also typically issue physical banknotes. Since both are supposed to serve as the unit of account, they exchange one-to-one. Banknotes though come with the property that they pay no interest. This puts a constraint on the interest on reserves, namely that banks would want to substitute all of their reserves for banknotes if interest rates were to become negative. Banknotes therefore imply a zero lower bound on what the payment of interest on reserves can be.[10]

The ZLB poses a constraint on the process for monetary policy, since from the Fisher equation, non-negative nominal interest rates require that $P_{t+1}/P_t \geq \beta$. With the payment on reserves rule, or the Wicksellian rule, this lower bound makes no difference to the determination of inflation.[11] With a Taylor feedback rule though, the ZLB becomes a serious issue. The Taylor monetary policy process then implies that:

$$1 = \frac{\beta}{\Pi_{t+1}} \max \left\{ \frac{\Pi^*}{\beta} \left(\frac{\Pi_t}{\Pi^*} \right)^{\phi}, 1 \right\} \tag{2.31}$$

This difference equation for inflation is represented in Figure 2.1. As soon as inflation is equal to β, it stays there forever. This is a global steady state equilibrium of the system: a deflation trap. Moreover, note that if P_0 is below target, the system will converge in a finite number of periods to the deflation trap. Since any such deviation leads to this same outcome, consistent with the equilibrium conditions, then the price level is again indeterminate: any initial inflation between P^*_{t+1}/P^*_t and β is consistent with an equilibrium, and is not ruled out by excluding explosive solutions.

This problem arose in the first place because of the presence of banknotes. The next section focuses on banknotes directly as process to control inflation.

2.4 The monetarist approach to inflation control

Banknotes, or currency, are distinct from reserves in four ways. First, they can be freely held by anyone in the economy, not just banks. Second, they are physical but still cost close to zero to be produced by the central bank, who is the monopoly producer of them and their services. Third, they are anonymous as people do not have to declare to government bodies how much currency they have and who they

[10] In reality, the lower bound is below zero because banknotes are physical while reserves are electronic, so holding the former comes with inconvenience and storage costs. Eliminating banknotes, charging a tax on them, or defaulting on the commitment to exchange them one-to-one for reserves are the three ways to eliminate the zero lower bound.

[11] To see this with the payment on reserves process, note that the modified process now is: $1 + x_t = \max \left\{ \frac{1+\hat{R}_t}{P^*_t}, \frac{1}{P_{t+1}} \right\}$. Once the zero lower bound does not bind, then this process pins down the price level on target. When the zero lower bound binds at date t, then $P_{t+1} = P_t$. As long as the zero lower bound does not bind forever, then the price level is determinate, but the central bank is incapable of achieving the desired target as inflation is equal to the negative of the real interest rate during the ZLB episodes.

Figure 2.1: Inflation dynamics at the ZLB

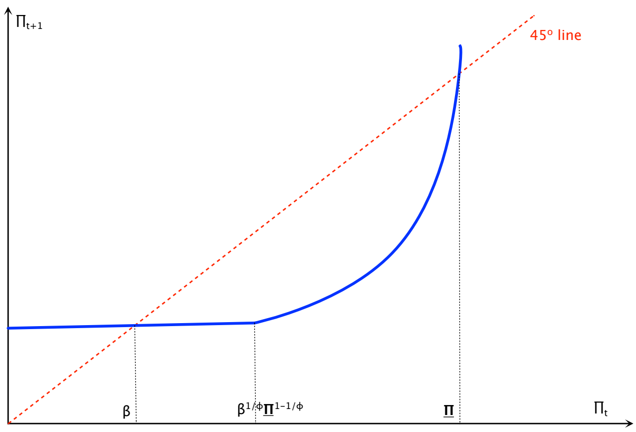

transfer it to. Fourth, they pay no interest.

All combined, these properties lead to there being some demand for money which declines with the nominal interest rate in the economy. They can be held by people with tastes for privacy, or who do not want to use a bank as an intermediary to make their payments, so they may provide some services in facilitating transactions. Because the alternative to banknotes are reserves, the higher is the interest rate they pay, the higher the opportunity cost of holding currency and thus the lower the demand for it.

A particular form of this demand function is:

$$\frac{H_t}{C_t P_t} = e^{u_t} \left(\frac{I_t}{1 + I_t} \right)^{-\tilde{\eta}}, \tag{2.32}$$

where $H_t/P_t \geq 0$ are real money balances, and u_t represents a money demand shock. The income elasticity of the demand for currency is set to one to be consistent with the (rough) balanced growth fact that the left hand-side of this equation (sometimes called the inverse of velocity) does not have a strong trend over decades in the data. The interest-rate elasticity is measure by $\tilde{\eta}$, which we take to be constant, even though attempts to estimate it often find it is not, and is close to zero. More generally, attempts to estimate this equation have been fraught with difficulties since both the measurement of H_t is hard, and all estimates find very poor fits or that the variance of u_t is very large.

The central bank controls the supply of banknotes H_t^S. However, because of the existence of close substitutes to currency produced by the private market, financial innovation implies that the market clearing condition is

$$H_t = e^{v_t} H_t^S, \tag{2.33}$$

where v_t is very volatile due to shocks to payments systems and financial innovation. Moreover, because central banks stand ready to exchange reserves for banknotes one-to-one at all times, even controlling the exact supply of banknotes is hard.

Combining these two equations, for currency supply and demand, together with the Fisher equation and using the log-linearized version such that $h_t = \log(H_t^S/\bar{H})$ and $\eta = \tilde{\eta}/I$, we get the key equation of the monetarist approach:

$$h_t - p_t = c_t - \eta(r_t + \mathbb{E}_t \Delta p_{t+1}) + u_t - v_t. \tag{2.34}$$

2.4.1 Microfoundations and the terminal condition

Relative to the microfoundations in section 2.2, private agents can now transfer funds across time at zero interest rate by holding currency, and their preferences are

$$u(C_t, H_t/P_t) = \left[C_t^{1-1/\tilde{\eta}} + e^{u_t/\tilde{\eta}}(H_t/P_t)^{1-1/\tilde{\eta}}\right]^{\tilde{\eta}/(\tilde{\eta}-1)}. \tag{2.35}$$

The first order condition gives rise immediately to the demand function stated before.

Another optimality condition from this optimization problem is the transversality condition:

$$\lim_{T\to\infty} \beta^T \left(\frac{\partial U(C_T, H_T)}{\partial C_T}\right)\left(\frac{H_T}{P_T}\right) = 0. \tag{2.36}$$

This states that the agents do not want to hold positive stocks of currency until infinity because they could instead convert them into consumption and raise utility.

2.4.2 Money growth processes

The central bank could choose to exogenously set money supply h_t as it wished. This is a different approach to the one in the previous section. For a given h_t that controls inflation, there will follow an endogenous i_t that clears the market for currency. The interest paid on reserves now becomes a secondary tool, adjusting endogenously to support the exogenous choice of currency supply.

An effective rule for this supply is

$$h_t = p_t^* + \hat{c}_t - \eta(\hat{r}_t + \hat{p}_{t+1}^* - p_t^*) + \hat{u}_t - \hat{v}_t, \tag{2.37}$$

which focuses on the central bank's estimates of the shocks to supply and demand for means of payment, together with the real interest rate and the level of consumption as before.

Combining equations (2.34) and (2.37), iterating forward and imposing the transversality condition, (2.36), yields $p_t = p_t^* + \epsilon_t$ uniquely. The central bank is able to control inflation and the associated error is:

$$\epsilon_t = \frac{1}{\eta}\sum_{s=0}^{\infty}\left(\frac{\eta}{1-\eta}\right)^s \mathbb{E}_t[\hat{c}_{t+s} - c_{t+s} - \eta(\hat{r}_{t+s} - r_{t+s}) + \tag{2.38}$$

$$(\hat{u}_{t+s} - u_{t+s}) - (\hat{v}_{t+s} - v_{t+s}) + \eta(p^*_{t+1+s} - \hat{p}^*_{t+s+1})]. \tag{2.39}$$

Monetarism controls inflation. At the same time, the control error is very large, especially when compared to the one with interest-rate rules. The shocks to currency demand and currency supply, u_t and v_t, are very large.

2.4.3 Seignorage processes

When the central bank prints more currency, it buys goods with it and these give rise to a resource flow called seignorage. Since it costs close to nothing to produce currency and there is a downward-sloping demand for it, currency is not a liability of the central bank, but rather a durable good that it produces and sells for its value $1/P_t$. Seignorage is then equal to:

$$S_t = \frac{H_t^S - H_{t-1}^S}{P_t} = C_t\left[e^{u_t - v_t}\left(\frac{I_t}{1+I_t}\right)^{-\tilde{\eta}} - e^{u_{t-1}-v_{t-1}}\left(\frac{I_{t-1}}{1+I_{t-1}}\right)^{-\tilde{\eta}}\frac{P_{t-1}C_{t-1}}{P_tC_t}\right]. \tag{2.40}$$

The second equality used the expression for currency demand. This expression makes clear that seignorage and inflation are tightly linked. Higher expected inflation comes with higher nominal interest rates, which lowers the demand for currency and lowers seignorage. At the same time, a higher unexpected inflation implies that more goods can be bought with the newly printed banknotes, which raises seignorage.

Consider then a policymaker which follows a policy process for seignorage. It is committed to generating some revenue, like a government fiscal agency providing some public service and which is given a target for sales. It will then print more or fewer banknotes as needed to reach this target. We can log-linearize the expression above to end up with a second-order difference equation for inflation, which once combined with the transversality condition provides a unique solution for inflation. A central bank that behaves strictly like a fiscal agent is able to control inflation as explained by Sargent and Wallace (1975).

There are a few unsatisfying features of this process for inflation control. First, since S in steady state is bounded above by Ce^{u-v}, there is a limit that likely moves over time to how much one can use this approach. If the central bank aims to raise too much revenue, above this limit, this will not pin down inflation. Second, seignorage is usually quite small, and shocks to demand and supply of currency make its link to inflation very volatile. Therefore, calibrating it for use is prone to large control errors. Third, seignorage is very easy to raise in the short run. This makes it a tempting source of revenue for desperate policymakers needing to fill a fiscal hole, even if the end result is high inflation. In fact, historically, many hyperinflations were preceded by large fiscal crises, and ended with fiscal reforms (Sargent and Wallace, 1975).

2.5 The budget approach

So far, we have described a central bank that issues reserves, prints currency, and may even also borrow through long-term bonds. With these funds, central banks buy assets and pay for expenses as well as rebate dividends to the fiscal authorities. Letting V_t denote the value of reserves chosen at $t-1$, A_t the real value of assets chosen at $t-1$, E_t be real expenses and D_t be real dividends, then the flow budget constraint of the central bank is:

$$V_{t+1} + (1+R_{t-1})P_t A_t = (1+I_{t-1})V_t + P_t A_{t+1} + P_t(E_t + D_t - S_t). \tag{2.41}$$

Because private agents must voluntarily decide to hold reserves, they will have a transversality condition that imposes a no Ponzi scheme condition on the central bank. This takes the form:

$$\lim_{T\to\infty} \beta^T \left(\frac{\partial U(C_T, H_T)}{\partial C_T}\right)\left(\frac{V_{T+1}}{P_T} - A_{T+1}\right) = 0. \tag{2.42}$$

It states that the liabilities of the central bank must at infinity be covered by its assets so the holders of the reserves always get paid.

Another important condition is the solvency constraint for the central bank. Central banks are part of the government, but they are independent, and this independence reflects itself as a constraint on the flow of dividends between the central bank and the government. If there was no independence, then the central bank's balance sheet could just be consolidated with that of the government, as unrestricted D_t would connect the two. An independent central bank that was fully fiscally backed would have its dividends always match precisely its net income, whether this is positive or negative, so its net worth would always be zero and its budget constraint would again never bind. Instead central banks are often independent and not fully backed, which can be expressed as a constraint on the stream of dividends they must deliver.

We assume a simple case of such a solvency constraint, whereby the present value of dividends from the central bank must be zero, so the central bank is intertemporally solvent. The intertemporal budget constraint of the central bank then is:

$$(1+I_{t-1})\frac{V_t}{P_t} = (1+R_{t-1})A_t + \mathbb{E}_t \sum_{j=0}^{\infty} M_{t+1+j}\left(S_{t+j} - E_{t+j}\right). \tag{2.43}$$

Intuitively, this states that the real liabilities of the central bank equal its real assets plus the present value of revenues minus expenses. Substantively, the economics behind it is the solvency condition for the central bank. Operationally, it points to a new tool that the central bank can use to control inflation: the size of the balance sheet, or the amount of outstanding reserves.

2.5.1 The budget process

Assume for simplicity that $E_t - S_t$ is exogenous i.i.d. mean zero ε_t. Rearranging the budget constraint of the central bank then gives:

$$P_t = \frac{(1 + I_{t-1})V_t}{(1 + R_{t-1})A_t - \varepsilon_t}. \tag{2.44}$$

Since V_t and A_t were chosen at $t-1$, they are exogenous at period t. Therefore, this gives a single equation with a single unknown and uniquely pins down P_t.

If the central bank chooses the size and composition of its balance sheet with a target inflation in mind, then this process will deliver inflation on target. All else equal, the larger is V_t the larger will the price level be. The intuition is simply that if the central bank has a larger debt, then this debt must be worth less, and the inverse of the price level is the real value of the debt. Controls errors will be related to surprises to revenues and expenses as captured by ε_t. If the central bank has further control over its net expenses though, it can at least in principle minimize these errors.

2.5.2 Relation to the FTPL

In this framework, there is no fiscal policy. Nevertheless, the above argument can be easily framed in terms of the Fiscal Theory of the Price Level. Simply assume away central bank independence, so that the budget of the central bank and the government are integrated. Their total nominal liabilities are now the sum of central bank reserves and the market value for the government's multitude of public bonds, netting out the government bonds held by the central bank. In turn, net revenues now include not just the central bank's, but much larger and more relevant, the primary surplus of the government.

The economic logic of price level determinacy is the same, and that is that the larger is the total government debt, then the larger the price level will be in order to bring the real value of that debt in line with the real surpluses. Trickier in this perspective is the coordination between authorities. Now, control errors will arise also from the side of the government and its fiscal surplus. Moreover, attempts by the central bank to expand its balance sheet can be offset by decreases in borrowing from the government. Finally, changes in the maturity of the outstanding public debt affect the variability of inflation, but the central bank does not control the Treasury's choices of how much debt of each maturity to issue.

2.6 Gold standard approach: choosing pegs

So far, we have considered only one good in the economy, and focused on financial choices, the choice of currency, and the choice to hold government liabilities. Now, we focus instead on consumption choices to describe one of the most classic forms of inflation control: pegging the exchange rate of your unit of account to an external good or currency.

With many non-durable goods available for consumption, the optimality condition of the household is that the log of the marginal rate of substitution between any two goods, say good j and good 0 must equal the log of their relative prices:

$$\rho_t(i) = p_t(i) - p_t(0). \tag{2.45}$$

The marginal rate of substitution is determined by the amounts of each good consumed and produced which, just like the stochastic discount factor earlier, are exogenous from the endowments.

In turn, the price level is an index over the price of all goods, which we take as being a simple geometric average:

$$p_t = \sum_{i=0}^{I} \omega_i \, p_t(i) = p_t(0) + \sum_{i=1}^{I} \omega_i \rho_t(i). \tag{2.46}$$

The weights ω_i are non-negative and sum to one, and should be independent of the overall price level since agents do not suffer from money illusion.

This is the new key equation. It is capturing demand for goods given changes in their relative prices. Because that choice is static, no terminal conditions are needed or show up. The policy tool is that the central bank can choose to elect one (or several goods) to peg its unit of account to.

2.6.1 A commodity peg

The tool of the central bank is to decide which one good is to become the anchor for the unit of account. Reserves, or even currency, can now be denominated in terms of units of that good. Taking that good as a gram of gold, say, then the central bank can say that such a gram is worth 100 dollars say. Because the central bank decides the unit of account for its reserves it can simply do this by decree. It just has to announce that from now onwards 100 dollars of reserves will be able to buy one gram of gold. The central bank will issue reserves, and if asked will exchange these for currency, in unlimited amounts to keep the price of one gram equal to 100 dollars.

This clearly uniquely determines the price level. The question is rather how effective this approach is. Consider a policy process then that decides that the arbitrary good 0 will be the commodity to which the price level is pegged to. The central bank would like to adjust this peg in case of relative-price movements, but only has an estimate of them. Its policy process is:

$$p_t(0) = p_t^* - \sum_{i=1}^{I} \omega_i \hat{\rho}_t(i), \tag{2.47}$$

where $\hat{\rho}_t(i)$ is the estimate of the corresponding relative prices. The result is that the price level is on target with control errors given by

$$\varepsilon_t = \sum_{i=1}^{I} \omega_i (\hat{\rho}_t(i) - \rho_t(i)), \tag{2.48}$$

The main problem with this approach is that these errors can be large. Changes in the supply of the key good 0 or in the public's tastes for it become sources of deviations of inflation from target. Moreover, insofar as the good is complementary with other, then changes in their supply will spill over to relative prices again leading to control errors. The ideal good has a stable supply and is not complementary or substitutable with that many other goods. Gold or other precious metals meet these two criteria explaining why they have been used through history.

2.6.2 A peg to a basket

Instead of picking a good, the central bank could choose a wide consumption basket, but having to come up with estimates of what their consumption basket weights will be in any given period $\hat{\omega}_i$. In this case, the policy process now is:

$$\sum_{i=0}^{I} \hat{\omega}_i p_t(i) = p_t^*, \tag{2.49}$$

and the control error is: $\varepsilon_t = \sum_{i=0}^{I} (\hat{\omega}_i - \omega_i) p_t(i)$. While this is more demanding of the central bank and harder to implement, it may lead to smaller errors.

2.6.3 An exchange rate peg

Even more common than pegging to gold, is to peg to a foreign currency like the dollar or the euro. A small open economy, which imports goods from other countries denominated in this dominant foreign currency, will often accumulate large amounts of this currency and stand ready to buy or sell them against its domestic reserves to keep to this peg.

Letting j be a good denominated in this foreign currency, the first-order condition in logs from the consumer is:

$$\rho_t(i, j) \equiv p_t(i) - p_t(j) - e_t, \tag{2.50}$$

where e_t is the log nominal exchange rate. Letting α denote the measure of home bias and p_t^F be the foreign price index, which are both exogenous to the domestic economy, the domestic price level is defined as:

$$p_t = \alpha \sum_{i=0}^{I} \omega_i p_t(i) + (1-\alpha) \sum_{j=0}^{J} \omega_j \left[p_t(j) + e_t\right] = \alpha \sum_{i=0}^{I} \omega_i p_t(i) + (1-\alpha) \left[p_t^F + e_t\right] . \tag{2.51}$$

Since the relative prices can be combined to read $\sum_{j=0}^{J} \omega_j \, \rho_t(i,j) = p_t(i) - p_t^F - e_t$, it follows that the price index becomes:

$$p_t = \alpha \sum_{i=0}^{I} \sum_{j=0}^{J} \omega_i \omega_j \rho_t(i,j) + e_t + p_t^F . \tag{2.52}$$

Pegging the exchange rate is choosing an e_t. This pins down the right-hand side of this equation, and therefore the price level in the domestic economy.

2.7 The Phillips curve approach to inflation control

Central bankers often discuss the process of controlling inflation as a result of their ability to affect real activity. The argument goes as follows: central bank policy has an effect on the level of slack in the economy, meaning the amount of resources in left unused. When firms use labor and capital more intensively (less slack), production costs tend to rise and firms want to raise their relative prices accordingly. As they all do so, the absolute price level rises.

So far, we have used (and maybe abused) the classical dichotomy to separate the control of inflation from the effects of monetary policy on the real economy. In this section, we let go of the assumption that output is exogenous. We replace it with the standard assumption in the new Keynesian literature that monopolistic firms choose the price of the good they produce to increase proportionally with marginal costs and a markup. The price set is subject to nominal rigidities, and the firms stand ready to produce whatever is demanded at that price.

The effect of introducing nominal rigidities and demand-determined output is to deliver an expectations-augmented Phillips curve (in log-linearized variables):

$$\pi_t = \pi_t^e + \kappa m c_t + u_t \tag{2.53}$$

The first term on the right-hand side π_t^e is a measure of expected inflation by the firms who set prices. Different models of nominal rigidities lead to different specifications of this term, and we will discuss a few alternatives in section 2.7.2. The second term is a measure of aggregate real marginal costs mc_t, which is associated with slack in product and labor markets. Several central bank tools can affect slack and the costs faced by firms. Section 2.7.1 discusses credit while section 2.7.3 considers other alternatives. The third and final term u_t is a stationary mean-zero markup shock, capturing changes in the market power of firms.

2.7.1 The New Keynesian Phillips curve and credit tools

The most popular model of nominal rigidities is the sticky price model of Calvo (1983). It assumes that firms set their price every period according to a Poisson arrival, and otherwise have to keep their price unchanged from last period. Being forward looking, when firms choose their price they take into account the probability they will not be able to change it again into the infinite future. Expected future inflation becomes an aggregate sufficient statistic for how expectations of future changes in the price level, marginal costs, and markups will evolve. The result is that $\pi^e = \beta\,\mathbb{E}_t(\pi_{t+1})$ where $\beta < 1$ is the steady-state discount factor.

Among the many tools used by central banks that we have discussed so far, the missing candidate is credit policies. Central banks can have a direct effect on the amount and price of private credit to firms. They do so directly by setting the requirements for minimum reserves that banks must hold at the central bank as a fraction of their deposits (or loans) and by imposing a variety of macroprudential tools, including requirements that credit is bound to be a multiple of net worth, or that a fraction of assets must be held in liquid marketable assets instead of loans. In the recent decade, both the Bank of England and the European Central Bank lent funds to banks at favorable rates under the condition that these funds were lent to firms.[12] We capture this myriad of policies by assuming that the central bank targets the equilibrium real lending rate faced by firms r_t^L.

One component of firms' marginal costs is the cost of having to raise funds externally through loans to pay for investment and working capital. The opportunity cost of doing so is the external finance premium $r_t^L - r_t$, the difference between raising funds externally from banks or internally where the opportunity cost is the real interest rate. Marginal costs can then be written as: $mc_t = \psi(r_t^L - r_t) + c_t$ where c_t includes all the other stochastic determinants of marginal costs that we take to be independent of inflation and ψ is the external finance premium elasticity of marginal costs.

Combining these ingredients gives, as in previous sections, a pair of equations capturing the economic mechanism and the policy tool rule:

$$\pi_t = \beta\,\mathbb{E}_t(\pi_{t+1}) + \kappa\psi(r_t^L - r_t) + \kappa c_t + u_t \tag{2.54}$$

$$r_t^L = \hat{r}_t + (\kappa\psi)^{-1}\left[(1+\beta)p_t^* - p_{t-1} - \beta\hat{p}_{t+1}^* - \kappa\hat{c}_t - \hat{u}_t\right] - p_{t-1} + e_t^l\,. \tag{2.55}$$

In the policy tool, a new source of error emerges, e_t^l. This is the error that arises from the lending rate not being directly set by the central bank, but only targeted

[12] These policies were called the Funding to Lending scheme, and the Targeted Long-term Repurchase Operation, respectively.

through different policies.

Imposing that $\lim_{T\to\infty} \beta^t \, \mathbb{E}_t(...) = 0$ is uncontroversial since this is a combination of forecasting errors that even minimally-rational agents should not expect to explode over time, and real marginal costs which involve real resources and so are subject to the consumer's transversality condition. Therefore, iterating forward, inflation is determinate and under control subject to errors given by:

$$\varepsilon_t = \frac{1}{1+\beta} \sum_{s=0}^{\infty} \left(\frac{1}{1+\beta} \right)^s E_t \big[\beta (p^*_{t+1+s} - \hat{p}^*_{t+1+s}) + \kappa (c_{t+s} - \hat{c}_{t+s}) + \kappa\psi (\hat{r}_{t+s} - r_{t+s}) + (u_{t+s} - \hat{u}_{t+s}) + \kappa\psi e^l_{t+s} \big] . \quad (2.56)$$

A central bank that relies on the Phillips curve to control inflation is subject to three new sources of errors, beyond uncertainty about real interest rates and miscommunication about the target and the policy tool as in our previous analyses. First are errors on markup shocks, which policymakers often refer to as "cost-push" shocks, and were notably used to explain the failure to control U.S. inflation in the 1970s. Separating changes in markups from movements in marginal costs is itself difficult in real time, and often takes economists years or decades to agree on.

Second, are errors on what current and future marginal costs are. Productivity is notably volatile and hard to predict, as is the the marginal cost of hiring an additional worker when there are different firms and households. As a result, central banks end up employing significant resources to measure growth potential, output gaps, natural rates, and other concepts that go into c_t. In addition, the slope of the Phillips curve, both in the sense of the structural parameter κ, and in the sense of the reduced-form correlation between inflation and measures of slack, seems to shift often enough that the policy rule is hard to implement without considerable errors.

Aside from these sources of errors arising from the Phillips curve, there is a third source arising from the credit channel of monetary policy. Because r^L_t is not a direct policy tool, but rather an intermediate target that a myriad of tools try to target, there will be errors e^l_t. The source of the error is similar to the supply and demand shocks for currency in section 2.4: financial innovation in credit markets or any shock to the banking sector will spill over to the level of inflation if the central bank uses this policy tool.

2.7.2 Alternative models of the Phillips curve and expected inflation

The literature has produced different models of nominal rigidities beyond sticky prices.[13] A large class of them focuses on imperfect information as the reason why

[13] And even within sticky price models, the Calvo model above is just the more popular, with several alternatives including adjustment-cost models, fixed-cost state-dependent adjustment models, partial indexation models, and others.

firms do not perfectly adjust their prices to shocks. A simple canonical model in this class has $\pi_t^e = \mathbb{E}_{t-1}(\pi_t)$, which results from some firms receiving information with a one-period delay, and so choosing their price based on their old information.[14] A third class of models assumes instead that agents are backward-looking and imitate econometricians in using past time-series realizations data to make forecasts about the future. The simplest classic model is the adaptive expectations assumption that $\pi_t^e = \pi_{t-1}$.[15]

Starting with the imperfect-information case, the policy rule that controls inflation is now:

$$r_t^L = \hat{r}_t + (\kappa\psi)^{-1} \left(\pi_t^* - \beta\hat{\mathbb{E}}_{t-1}(\pi_t) - \kappa\hat{c}_t - \hat{u}_t\right) + e_t^l , \tag{2.57}$$

The novelty is that the central bank must now respond to deviations of past inflation expectations from the target. The central bank in this setting monitors surveys of expectations, market prices of assets that provide hedges against inflation, or any other piece of information that summarizes what people expected inflation to be. If these measures differ persistently from target, the central bank responds aggressively via its policy tool to affect slack and through it actual inflation.

By doing so, the central bank can control inflation so that $p_t = p_t^* + \varepsilon_t$ and the error now is:

$$\varepsilon_t = \beta \left(\mathbb{E}_{t-1}(\pi_t) - \hat{\mathbb{E}}_{t-1}(\pi_t)\right) + \kappa(c_t - \hat{c}_t) + \kappa\psi(\hat{r}_t - r_t) + (u_t - \hat{u}_t) + \kappa\psi e_t^l . \tag{2.58}$$

There are two differences relative to the sticky price case. First, insofar as the central bank can measure and respond to past expectations of present inflation, then inflation is no longer affected by forecast errors about future conditions. Second, and the other side of the coin, errors in measuring these expectations are now the main source of control errors.

Turning to the backward-looking Phillips curve, naturally, the policy rule that controls inflation now involves responding to past realizations of inflation:

$$r_t^L = \hat{r}_t + (\kappa\psi)^{-1} (\pi_t^* - \beta\hat{\pi}_{t-1} - \kappa\hat{c}_t - \hat{u}_t) + e_t^l , \tag{2.59}$$

where we allow for the fact that due to measurement lags, past inflation is typically not known precisely and must be estimated by the central bank. This naturally leads to the control errors:

$$\varepsilon_t = \beta (\pi_{t-1} - \hat{\pi}_{t-1}) + \kappa(c_t - \hat{c}_t) + \kappa\psi(\hat{r}_t - r_t) + (u_t - \hat{u}_t) + \kappa\psi e_t^l . \tag{2.60}$$

[14]Within imperfect information models, there are dynamic sticky information models, rational inattention models, models where agents receive imperfect signals on the state of the world, imperfect information arising due to higher-order uncertainty, and many combinations of these.

[15]Alternatives here include models with least-squares learning, imperfect knowledge, extrapolate expectations, diagnostic expectations, and others.

From this view on policy, a major investment by the central bank should go into real-time measurement of inflation, complementing the measurement of real interest rates, marginal costs, and markups.

2.7.3 Alternative central-bank approaches to affect inflation via slack

Likely, lending rates are a small component of marginal costs. That is, c_t is probably very large. Therefore, using lending policies as the main monetary tool in conjunction with the Phillips curve, may not be particularly efficient at controlling inflation. Among the macroeconomic variables that affect marginal costs, the level of output stands out as the main driver as it through its effect on the relative scarcity of credit, the price of intermediate inputs or the level of wages. Therefore, considering $mc_t = \alpha y_t + c_t$ instead may well minimize the variance of the determinants of marginal costs that are orthogonal to output c_t

Once the classical dichotomy does not hold, any of the previously discussed policy tools at the disposal of the central bank can potentially affect real activity and so marginal costs. The setting of interest rates, potentially combined with forward-guidance and going long policies, that were discussed in section 2.3, the printing of banknotes discussed in section 2.4, the quantitative and qualitative easing that vary the size and composition of the central bank balance sheet in section 2.5, and the foreign exchange rate interventions discussed in section 2.6, all have effects on real activity. Through the Phillips curve, they will affect inflation as well, and so could be used to control it.

At a general level, the Phillips curve captures a breakdown of the classical dichotomy in the sense that all shocks now have an effect on both real outcomes and on inflation. That is, focusing on nominal income, $n_t = p_t + y_t$ in logs, the Phillips curve captures the menu of possible combinations of p_t and y_t that may result. This makes nominal income a natural intermediate target for monetary policy in a Phillips-curve approach to controlling inflation. Through its multiple tools, the central bank can steer nominal GDP. The Phillips curve then solves for how efficiently these control inflation.

Taking the view that policy is now choosing n_t, inflation solves:

$$\pi_t = \beta \, \mathbb{E}_t(\pi_{t+1}) + \kappa n_t - \kappa p_t + u_t. \tag{2.61}$$

This gives a stochastic second-order difference equation for p_t:

$$\beta \, \mathbb{E}_t \left(p_{t+1}\right) - (1 + \beta + \kappa) p_t + p_{t-1} = -\kappa n_t - u_t. \tag{2.62}$$

The quadratic equation $x^2 - (1+\beta+\kappa)x + \beta = (1 - x/\theta)(1 - \theta x)$ has two positive roots: θ and $\frac{\beta}{\theta}$ where we choose θ to be the smallest among the two. Therefore, the difference equation has the solution:

$$p_t = \theta p_{t-1} + \beta\theta \sum_{k=0}^{\infty} (\beta\theta)^k \, \mathbb{E}_t \left(n_{t+k} + u_{t+k} \right) \, . \tag{2.63}$$

From this it follows that choosing the policy target for nominal income to follow the process:

$$n_t = \frac{1-\beta\theta}{\beta\theta} \left[p_t^* - \theta p_{t-1} - \beta\theta \sum_{k=1}^{\infty} (\beta\theta)^k \left(\hat{n}_{t+k} + \hat{u}_{t+k} \right) \right] - \hat{u}_t + e_t^n \, , \tag{2.64}$$

will control inflation around the target with control errors:

$$\varepsilon_t = u_t - \hat{u}_t + \beta\theta \sum_{k=1}^{\infty} (\beta\theta)^k \left[\mathbb{E}_t(n_{t+k}) - \hat{n}_{t+k} + \mathbb{E}_t(u_{t+k}) - \hat{u}_{t+k} \right] + \frac{1-\beta\theta}{\beta\theta} e_t^n \, . \tag{2.65}$$

Arguably, the errors in controlling inflation are smaller with this broad nominal income targeting than with the narrow setting of credit policies. the trade-off is likely between the c_t being smaller but the e_t^n being larger. This captures a more general lesson in inflation control. Moving from direct instruments that the policymaker controls to intermediate targets on endogenous variables on the one hand reduces the errors arising from the relation between policy tool and economic variable while on the other hand increases the errors in keeping the intermediate target on target.

2.7.4 The three-equation model

We now turn to the classic three-equation New Keynesian model that incorporates the notion of output gap as a measure of the level of slack in the economy. The output gap, $\hat{y}_t$, is the difference between actual output and its estimated potential, understood as the output level that would arise in the absence of nominal rigidities. To characterize the corresponding New Keynesian Phillips curve, we impose a linear mapping from real marginal costs to output gap: $\hat{mc}_t = \gamma \hat{y}_t$, and sticky prices à la Calvo (1983):

$$\pi_t = \beta \, \mathbb{E}_t(\pi_{t+1}) + \kappa\gamma\hat{y}_t + u_t \, . \tag{2.66}$$

So far we have assumed there is some component of the real marginal cost that is exogenously determined. By fully dropping the classical dichotomy, a two-way relationship between inflation and the real economy emerges. To see this assume log utility and combine the Euler equation (2.1), the definition of the stochastic discount factor and the Fisher equation (2.4):

$$\frac{1}{Y_t} = \beta E_t \left[\frac{1}{Y_{t+1}} \frac{1+I_t}{\Pi_{t+1}} \right] \, . \tag{2.67}$$

Then, log-linearizing and rewriting in terms of the output gap delivers:

$$\hat{y}_t = \mathbb{E}_t(\hat{y}_{t+1}) - (i_t - \mathbb{E}_t(\pi_{t+1}) - r_t^n)\,, \tag{2.68}$$

where r_t^n is the natural rate of interest. This equation is called the IS curve. Higher expected inflation increases current output through an increase in consumption as the incentive to save (the real interest rate) is reduced, other things equal.

The main implication of this new two-way relationship is that indeterminacy now has real effects.[16] Absent appropriate monetary policy, different combinations of inflation rates and output allocations may well satisfy equations (2.66) and (2.68), resulting in multiple equilibria. To prevent this, the three-equation New Keynesian models features a Taylor-type rule that responds to both inflation and the output gap:

$$i_t = r_t^n + \phi_\pi \pi_t + \phi_y \hat{y}_t + v_t\,, \tag{2.69}$$

where v_t is a monetary policy shock while ϕ_π and ϕ_y are inflation and output Taylor coefficients respectively.

Picking the Taylor coefficients to ensure determinacy in this setting involves slightly more algebra. First, rewrite equations (2.66), (2.68) and (2.69) in matrix notation:

$$\begin{bmatrix} \hat{y}_t \\ \pi_t \end{bmatrix} = A \begin{bmatrix} E_t\{\hat{y}_t\} \\ E_t\{\pi_{t+1}\} \end{bmatrix} + \Omega v_t\,, \tag{2.70}$$

$$\text{where } A \equiv \Omega \begin{bmatrix} 1 & 1-\beta\phi_\pi \\ \kappa\gamma & \kappa\gamma + \beta(1+\phi_y) \end{bmatrix} \text{ and } \Omega \equiv \frac{1}{1+\phi_y+\kappa\gamma\phi_\pi}\,.$$

Next, evaluate the stability of a system of linear difference equations using the Blanchard and Kahn (1980) conditions. These state that a unique non-explosive solution will exist whenever the number of eigenvalues of A outside the unit circle is equal to the number of non-predetermined variables. For the above, both eigenvalues of A lie within the unit circle provided: $\phi_y(1-\beta) + \kappa\gamma(\phi_\pi - 1) > 0$.[17] This is an

[16] See Carlstrom and Fuerst (2002) and Nakajima and Polemarchakis (2005) for further discussion on sticky prices and real indeterminacy.

[17] The characteristic polynomial of A is given by:

$$det \begin{vmatrix} \frac{1}{1+\phi_y+\kappa\gamma\phi_\pi} - \lambda & \frac{1-\beta\phi_\pi}{1+\phi_y+\kappa\gamma\phi_\pi} \\ \frac{\kappa\gamma}{1+\phi_y+\kappa\gamma\phi_\pi} & \frac{\kappa\gamma+\beta(1\phi_y)}{1+\phi_y+\kappa\gamma\phi_\pi} - \lambda \end{vmatrix} = 0\,,$$

which can be rewritten as:

$$[1-\lambda(1+\phi_y+\kappa\gamma\phi_\pi)]\,[\kappa\gamma+\beta(1+\phi_y)-\lambda(1+\phi_y+\kappa\gamma\phi_\pi)] - \kappa\gamma(1-\beta\phi_\pi) = 0\,,$$

and simplifies to: $\lambda^2 + a_1\lambda + a_0 = 0$, where $a_1 = -\frac{1+\kappa\gamma+\beta(1+\phi_y)}{1+\phi_y+\kappa\gamma\phi_\pi}$ and $a_0 = \frac{\beta}{1+\phi_y+\kappa\gamma\phi_\pi}$. Both eigenvalues of A lie within the unit circle if and only if both of the following conditions hold:

(i) $|a_0| < 1$

(ii) $|a_1| < 1 + a_0$

Condition *(i)* is trivially satisfied as $0 < \beta < 1$. For condition *(ii)* to be satisfied: $\phi_y(1-\beta) + \kappa\gamma(\phi_\pi - 1) > 0$.

generalized version of the Taylor principle.

The New Keynesian model as presented here is subject to the same caveats highlighted in sections 2.3.4 and 2.3.6: determinacy relies on a terminal condition and policy specifications disregard the zero lower bound constraint on the Taylor rule. The recent monetary experience has made the latter particularly relevant and a number of papers have studied the model's behavior in a liquidity trap. Cochrane (2016) argues that central banks can increase inflation at the zero lower bound by increasing their interest rate targets. Eggertsson (2010) and Eggertsson et al. (2014) emphasize that positive supply shocks, such as increased labor supply or structural reforms, can be contractionary.[18] Christiano et al. (2011), Eggertsson (2011), Woodford (2011b) and Carlstrom et al. (2014) find fiscal multipliers to be larger than normal times and increasing in the duration of the fiscal stimulus. Del Negro et al. (2012) document that the model overestimates the macroeconomic effects of central bank announcements about future interest rates.[19] Cochrane (2017) stresses policy predictions are strengthened as price stickiness is reduced despite effects vanishing at the friction-less limit.

In these papers, policy is often characterized as an exogenous path for the nominal interest rate $\{i_t^*\}_{t=0}^{\infty}$. To overcome the indeterminacy that results from an interest rate peg, Werning (2012) shows that featuring a Taylor rule of the type: $i_t = i_t^* + \phi(\pi_t - \pi_t^*)$, where π_t^* is the equilibrium inflation rate, is observationally equivalent. However, there are many inflation rates that are consistent with i_t^* in equilibrium and, thus, it is the central banker's choice which one she picks to implement. Cochrane (2017) shows that this equilibrium selection process is key in understanding the puzzling predictions of the New Keynesian model at the zero lower bound. Different π_t^* deliver different policy prescriptions.

2.8 Conclusion

In his presidential lecture to the American Economic Association, Christopher A. Sims (2013) concluded:

> The kinds of models that have been the staple of undergraduate macroeconomics teaching, with price level determined by balance between "money supply" and "money demand", and money supply described using the "money multiplier", are obsolete and provide little insight into the policy issues facing fiscal and monetary authorities in the last few years. There are relatively simple models available, though, that could be taught in

[18] Wieland (2019) provides empirical support to this claim.

[19] Many papers have since then attempted progress on addressing the so-called forward guidance puzzle. McKay et al. (2016) focus on incomplete markets while Angeletos and Lian (2018) and Gabaix (2018) abstract from rational expectations and propose a framework with imperfect information and behavioral myopia respectively.

undergraduate and graduate courses and that would allow discussion of current policy issues using clearer analytic foundations.

This article tried to present these simple models, emphasizing their common features, and presenting their differences in terms of different policy choices rather than as opposing schools of thought. In future work we hope to evaluate each of processes empirically by estimating their hypothetical effectiveness.

2.9 Bibliography

Adão, B., I. Correia, and P. Teles (2014): "Short and Long Interest Rate Targets," *Journal of Monetary Economics*, 66, 95–107.

Angeletos, G.-M. and C. Lian (2018): "Forward Guidance without Common Knowledge," *American Economic Review*, 108, 2477–2512.

Atkeson, A., V. V. Chari, and P. J. Kehoe (2010): "Sophisticated Monetary Policies," *The Quarterly Journal of Economics*, 125, 47–89.

Bassetto, M. (2005): "Equilibrium and Government Commitment," *Journal of Economic Theory*, 124, 79–105.

Blanchard, O. J. and C. M. Kahn (1980): "The Solution of Linear Difference Models under Rational Expectations," *Econometrica*, 1305–1311.

Calvo, G. A. (1983): "Staggered Prices in a Utility-Maximizing Framework," *Journal of monetary Economics*, 12, 383–398.

Carlstrom, C. T. and T. S. Fuerst (2002): "Taylor Rules in a Model that Satisfies the Natural-Rate Hypothesis," *American Economic Review*, 92, 79–84.

Carlstrom, C. T., T. S. Fuerst, and M. Paustian (2014): "Fiscal Multipliers under an Interest Rate Peg of Deterministic versus Stochastic Duration," *Journal of Money, credit and Banking*, 46, 1293–1312.

Christiano, L., M. Eichenbaum, and S. Rebelo (2011): "When Is the Government Spending Multiplier Large?" *Journal of Political Economy*, 119, 78–121.

Christiano, L. and Y. Takahashi (2018): "Discouraging Deviant Behavior in Monetary Economics," Tech. rep., National Bureau of Economic Research.

Christiano, L. J. and M. Rostagno (2001): "Money Growth Monitoring and the Taylor Rule," Tech. rep., National Bureau of Economic Research.

Cochrane, J. H. (2011): "Determinacy and Identification with Taylor Rules," *Journal of Political Economy*, 119, 565–615.

——— (2016): "Do Higher Interest Rates Raise or Lower Inflation?" *Manuscript*.

——— (2017): "The New-Keynesian Liquidity Trap," *Journal of Monetary Economics*, 92, 47–63.

Del Negro, M., M. P. Giannoni, and C. Patterson (2012): "The Forward Guidance Puzzle," Tech. rep.

Eggertsson, G., A. Ferrero, and A. Raffo (2014): "Can Structural Reforms Help Europe?" *Journal of Monetary Economics*, 61, 2–22.

Eggertsson, G. B. (2010): "The Paradox of Toil," Tech. rep., FRB of New York Staff Report.

——— (2011): "What Fiscal Policy is Effective at Zero Interest Rates?" *NBER Macroeconomics Annual*, 25, 59–112.

Evans, G. W. and B. McGough (2018): "Equilibrium Selection, Observability and Backward-Stable Solutions," *Journal of Monetary Economics*, 98, 1–10.

Evans, G. W. and G. Ramey (1992): "Expectation Calculation and Macroeconomic Dynamics," *The American Economic Review*, 207–224.

——— (1995): "Expectation Calculation, Hyperinflation and Currency Collapse," Cambridge University Press, 307.

——— (1998): "Calculation, Adaptation and Rational Expectations," *Macroeconomic Dynamics*, 2, 156–182.

Farhi, E. and I. Werning (2018): "Monetary Policy, Bounded Rationality, and Incomplete Markets," .

Gabaix, X. (2018): "A Behavioral New Keynesian Model," *Manuscript*.

García-Schmidt, M. and M. Woodford (2015): "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis," Tech. rep., NBER Working Paper 21614.

——— (2019): "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis," *American Economic Review*, 109, 86–120.

Hall, R. E. and R. Reis (2017): "Achieving Price Stability by Manipulating the Central Bank's Payment on Reserves," .

Hume, D. (1752): *Political Discourses*, A. Kincaid & A. Donaldson.

Magill, M. and M. Quinzii (2014): "Term Structure and Forward Guidance as Instruments of Monetary Policy," *Economic Theory*, 56, 1–32.

McCallum, B. T. (2009): "Inflation Determination with Taylor Rules: Is New-Keynesian Analysis Critically Flawed?" *Journal of Monetary Economics*, 56, 1101–1108.

McKay, A., E. Nakamura, and J. Steinsson (2016): "The Power of Forward Guidance Revisited," *American Economic Review*, 106, 3133–58.

Nakajima, T. and H. Polemarchakis (2005): "Money and Prices under Uncertainty," *The Review of Economic Studies*, 72, 223–246.

Obstfeld, M. and K. Rogoff (1983): "Speculative Hyperinflations in Maximizing Models: Can we Rule them Out?" *Journal of Political Economy*, 91, 675–687.

Reis, R. (2016): "Funding Quantitative Easing to Target Inflation," .

Sargent, T. J. and N. Wallace (1975): ""Rational" Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule," *Journal of Political Economy*, 83, 241–254.

Taylor, J. (1996): "Policy Rules as a Means to a More Effective Monetary Policy," *Monetary and Economic Studies*, 14, 28–39.

Taylor, J. B. (1993): "Discretion versus Policy Rules in Practice," in *Carnegie-Rochester Conference Series on Public Policy*, vol. 39, 195–214.

Werning, I. (2012): "Managing a Liquidity Trap: Monetary and Fiscal Policy," Tech. rep., National Bureau of Economic Research.

Wieland, J. F. (2019): "Are Negative Supply Shocks Expansionary at the Zero Lower Bound?" *Journal of Political Economy*, 127, 973–1007.

Woodford, M. (2011a): *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.

——— (2011b): "Simple Analytics of the Government Expenditure Multiplier," *American Economic Journal: Macroeconomics*, 3, 1–35.

——— (2013): "Macroeconomic Analysis Without the Rational Expectations Hypothesis," *Annual Review of Economics*, 5, 303–346.

Chapter 3

Dual Mandates and Excessive Hawkishness: a Principal-Agent Approach

3.1 Introduction

The aftermath of the recent financial crisis brought under discussion the benefits and reservations of the adoption of a dual mandate by central banks throughout the world. While the initial response of monetary authorities to the credit freeze and the incipient downturn was akin across developed economies, marked disparities in policy design during the late recession and, especially, during the recovery have questioned the suitability of inflation targeting as the unique objective of central banks.

In the first phase of the crisis, 2007-2009, the Federal Reserve (Fed), the European Central Bank (ECB) and the Bank of England (BoE) responded aggressively to the bursting of the US housing bubble by cutting interest rates to effectively zero and implementing an array of unconventional monetary policy measures in an attempt to ease liquidity and fix disrupted financial markets.

During the second stage of the crisis, 2010-2013, the Fed, on the one hand, continued to pursue intense monetary stimulus, conducting large open market operations, in order to propel robust economic growth as well as to enhance the recovery of the labor market. Its European counterpart, on the other hand, remained much more cautious: purchases made by the ECB appear limited in comparison to the large packages of quantitative easing implemented by the US and interest rates were even pushed up twice in 2011, despite the impending sovereign debt crisis in the Euro area.

According to some economists, the difference in mandates is accountable for the divergence in the strategies followed by policymakers on both sides of the Atlantic. Since 1977, the US Federal Reserve Act states that decision making in monetary policy should be guided by the objectives of maximum employment, stable prices and moderate long-term interest rates. Meanwhile, unlike the Fed, the ECB (and, similarly, the BoE) is mandated to deliver price stability, defined as keeping headline inflation around 2 percent in the medium run.

Undoubtedly the performance of the US economy has, by far, exceeded that of the Euro zone during the aftermath of the crisis: the recession was technically over as early as June 2009 and ever since GDP has been growing gradually but steadily with its rate reaching 2.4% in the fourth quarter of 2013. The Euro zone fell into a second and longer recession in the third quarter of 2011, which exited much later; GDP grew 0,3% in the last quarter of 2013 and the unemployment rate is at almost record levels of 12.1%.

While the suitability of a dual mandate is still uncertain, the recent experience has revived the debate. Should strict inflation targeting be complemented by economic growth objectives in any form, for example, targeting output gap or unemployment stabilization?

The ongoing discussion, however, has overlooked an outstanding revelation: it is precisely this 2010-2013 episode the unique example of clear dissension among monetary authorities from the main western economies. Despite long-established differences in objectives, central banks have responded to shocks using similar strategies. Strict inflation targeting has been the predominant approach to monetary policy implementation even in countries with an explicit dual mandate.

The historic trend is actually more surprising than by its recent reversal. As highlighted by The Economist in late 2012, it is common knowledge that policymakers at the US Federal Reserve have always felt more comfortable with the low inflation statutory goal than with full employment and it has not been until the recent crisis that the Fed has publicly emphasized that unemployment and inflation both carry weight in its decisions by setting a threshold for each objective.

Current differences in policy-making are perfectly in line with predictions made by economic theory, while former homogenized behavior is hardly rationalized using existing models. In fact, little or no attention has been devoted to studying the causes behind central banks' marked hawkishness. This paper examines the incentives that shape central bank's decision making in order to address the following question: why do central banks with a dual mandate prioritize inflation stabilization?

I build a model of monetary policy design based on a principal-agent problem framework, where society (or Congress) is forced to delegate the implementation of optimal policy to two independent agents. The monetary authority is entrusted with two tasks, inflation stabilization and economic growth, while the fiscal authority is only responsible for the latter. The agent evaluates the benefits of exerting effort in each of the tasks and the personal costs associated. The key elements are the compensation scheme and the coordination of strategies between the two agents. These can lead to excessive hawkishness by central banks with a dual mandate.

The rest of the paper is organized as follows: Section 3.2 summarizes prevailing monetary macroeconomic models, outlines related micro-foundations and reviews existing literature. Section 3.3 presents the theoretical framework, solving for the optimization problems and introducing main results. Section 3.4 discusses potential extensions and main caveats and, finally, Section 3.5 brings this paper to an end by adding some concluding remarks.

3.2 Literature Review

This section reviews, on the one hand, prevalent macroeconomics models of monetary policy design, and, on the other hand, the microeconomic literature on contract theory and moral hazard that will serve as a building block for my model.

Models of optimal monetary policy Optimal monetary policy has always been a recurrent topic of interest in Macroeconomics, especially after the hyperinflation episodes that followed World War I. However, the most fruitful period of research started in the mid 70s, and took off during the 80s and 90s (see also Taylor and Williams (2010) for a detailed review).

Economists address the question of how monetary policy should be conducted using either a historical or a model-based approach. The former, which focuses on the empirical analysis of a particular event or case study to provide intuition of what works in practice, has been loosing ground to the latter as sophisticated DSGE models with rational expectations and nominal rigidities have been developed.

Barro and Gordon (1983a) present the first positive theory of monetary and inflation policy in a purely discretionary environment that is relevant for this analysis. Building on the verbal and graphical model of Kydland and Prescott (1977), they align the monetary authority's objective with the preferences of a representative private agent. In other words, the extent of monetary response to a shock is governed, mostly, by society's relative dislike for inflation and unemployment. This theoretical framework lies within the main contributions of the paper and has been used by a large number of succeeding papers in the monetary literature, including the first section of the model I present below.

In addition, the main finding of the paper is that rational expectations lead to an inflation bias: the central bank will keep unemployment rate below its natural level, raising wages, which, ultimately lead to a higher inflation rate. This is explained by the inability of the central bank to commit to a zero inflation rate, the efficient monetary policy, i.e. the inconsistency of optimal plans.

The literature provides different solutions to the inconsistency problem in monetary policy: rules, reputation and delegation. In the first case, the policymaker can credibly commit in advance to a rule that is responsive to changes in observable variables, generally, changes in inflation and the output gap. McCallum (1986), Taylor (1986), Svensson (2001), and many others, discuss the design of these policy rules and estimate their performance relative to discretionary policies. However, as argued by Taylor (2011), despite the popularity of rule-based monetary policy during the 80s and the 90s, in the last decade, there has been a dramatic shift back toward discretionary policies.

A second alternative proposed by Barro and Gordon (1983b) and Backus et al. (1985) involves repeated interaction between the central bank and the private agents. In such a game, they argue, reputational considerations discourage the monetary authorities from pursuing surprise inflation (deviating from optimality) and, therefore, contracting constraints are not required.

Finally, the delegation solution assumes that the government delegates the monetary policy task to the central bank. In these models, the design of institutions plays an important role, given that it might lead to differences in the loss function adopted by the central banker and that desired by the government (or society). Some papers that support this approach include Rogoff (1985), Svensson (1997), and Chortareas and Miller (2003). Among this, Persson and Tabellini (1993) and Walsh (1995) introduce an explicit principal-agent approach, in which a penalization rate is imposed whenever the central bank deviates from the stipulated inflation target.

The main innovation of this set of papers, which I follow throughout the model, is the idea that it is not central banks but policymakers that set interest rates and are, therefore, subject to behavioral motives.[1] Following this notion, my model incorporates a moral hazard rationale to the monetary policy decision-making process. However, it addresses a completely different phenomenon: excessive hawkishness in central banks with dual mandates.

Moral hazard Given the above, the model is closely influenced by the mechanism design literature, more precisely, it formalizes an economic relation between society and the central banker that easily fits into the general principal-agent problem framework. As it is commonly known, the agency problem arises when the principal needs to delegate on a different party, the agent, the fulfillment of a given task. If the agent's actions are unobservable, there is an incentive for the agent to prioritize her own interest over the objective of the principal. In these circumstances, the principal seeks to design a compensation scheme that mitigates the effect of the informational asymmetries and indirectly incentivizes the agent choosing the correct action. Hart and Holmstrom (1987) provide a good introduction to the methodology of moral hazard.

The literature has developed a large number of extensions to this basic analysis. For example, Bernheim and Whinston (1986) consider a model with several principals but a single agent. Another interesting and popular extension by Fudenberg et al.

[1] What is sometimes known as *the who and the how of monetary policy* has drawn attention from researchers recently. Reis (2013) offers an important selection of topics in central bank design with a special mention to the strictness of the central bank's mandate.

(1990) and Allen (1985) among others, investigates the dynamics of the incentive mechanism. They investigate how efficient long-term contracts are relative to short-term ones, in a context in which the agency problem repeats over many periods. Dye (1986) allows for the observability of the agent's effort but after paying a monitoring cost.

The most relevant extension for this paper is Holmstrom and Milgrom (1991). They explore the more realistic case of multidimensional effort and develop a model in which the principal's utility depends on multiple tasks the agent engages in. In this setting, the problem cannot be simplified to a trade-off between risk sharing and incentive design. The principal must now take into account how the incentive provided for one action might affect the incentive to undertake a different task. In fact, the incentive contract varies with the shape of the cost function (degree of substitutability of efforts) and the degree of observability of tasks. The examples considered by the original paper include the remuneration of CEOs, salespeople job description or the debate on incentive pay for teachers, but can easily be extended to any other context.

3.3 Model

This section develops a model of monetary policy design which is two-fold and based on a multitask principal-agent setting. I first focus on the principal optimization problem in which society (or the general public) minimizes deviations in output and inflation from a given target subject to the constraints faced by the policymaker, in order to determine the optimal level of both monetary and fiscal policy.

I then assume that the principal cannot implement policy on its own, which would correspond to the first best in this context, but is forced to delegate the tasks to an agent. More specifically, this is a model with two agents and two tasks: inflation stabilization depends solely on the central bank while output stabilization depends on both the fiscal and the monetary authority. I abstract however from analyzing the fiscal authority and focus mainly on the central bank's optimization problem in the second part of this section.

3.3.1 Principal's Optimization Problem

This is a model of monopolistic competition and nominal price rigidity. Given the former, output produced by monopolists is inefficiently low and leaves room for policymakers to increase the supply of goods through a production subsidy. Note, however, deadweight losses prevent fiscal authorities from fully subsidizing production to its efficient outcome and, consequently generating an output gap. The monetary authority is then responsible, in addition to price stabilization, for closing the output gap through unanticipated monetary policies which have real effects due to staggered prices.

I follow Dixit and Lambertini (2003) in extending the Barro and Gordon (1983a) model to include not only monetary but also fiscal policy in the analysis. The social welfare loss minimization problem is specified as:

$$min \quad \frac{1}{2}\int \left(\pi\left(z\right)-\pi^{F}\right)^{2}+\theta_{F}\left(y\left(z\right)-y^{F}\right)^{2}+2\delta x(z)$$

$$subject\ to$$

$$\pi_{e}=E_{z}\left[\pi(z)\right]\equiv\int\pi(z)\,,$$

$$y(z)=\bar{y}+ax(z)+b\left[\pi(z)-\pi^{e}\right]\,,$$

$$\pi(z)=m(z)+cx(z)\,.$$

The interpretation of parameters is the following: $\bar{y}$ is the natural rate of output growth, $a > 0$ measures the direct effect of fiscal policy, $b > 0$ features the effect of an unexpected increase in the inflation rate on GDP and $c < 0$ captures the effect of fiscal policy on the price level. Monetary and fiscal policy instruments are denoted $m(z)$ and $x(z)$ respectively. The former can represent either the monetary base or the nominal interest rate while the latter is a production subsidy. π^F is the average level of pre-set prices in the economy, while y^F is the average level of GDP growth. Fiscal policy generates a deadweight loss measured by δ. Finally, θ_F is the relative weight given to output fluctuations in the social loss function and represents principal's preferences.

The minimization problem can be rewritten (and simplified) to read as follows:

$$min \quad \frac{1}{2}\int \left(m\left(z\right)+cx(z)-\pi^{F}\right)^{2}+\theta_{F}\left(\bar{y}+ax(z)+b\left[m\left(z\right)+cx(z)-\pi^{e}\right]-y^{F}\right)^{2}+2\delta x(z)$$

$$subject\ to$$

$$\pi_{e}=\int m(z)+cx(z)$$

3.3.2 Solving the Principal's Optimization Problem

As argued by Dixit and Lambertini (2003) plugging π^e into the objective function complicates the algebra substantially. I therefore follow their approach and assume that society has an additional choice variable which is also subject to the constraint, expected inflation π^e. The corresponding first order conditions read:

$$\begin{aligned} w.r.t \quad x(z): &\qquad (\pi(z)-\pi^F)c+\theta_F(a+bc)\left(y(z)-y^F\right)+\delta+\lambda c=0\,,\\ w.r.t \quad m(z): &\qquad (\pi(z)-\pi^F)+\theta_F b\left(y(z)-y^F\right)+\lambda=0\,,\\ w.r.t \quad \pi^e: &\qquad -\lambda-\int\theta_F b\left(y(z)-y^F\right)=0\,. \end{aligned}$$

After some rearrangements, the optimal inflation and output allocations are given by:

$$y^*(z)=y^F-\frac{\delta}{\theta_F a}\,, \tag{3.1}$$

$$\pi^*(z)=\pi_F+\frac{\delta b}{a}-\int\frac{\delta b}{a}\,. \tag{3.2}$$

The optimal inflation rate and output growth can be written as a function of exogenous variables. To solve for optimal monetary and fiscal policy, simply combine the first order conditions with the constraints of the original optimization problem, to get:

$$m(z)=\frac{c}{a}\cdot(\bar{y}-y(z))+(1+\frac{c}{a})\pi(z)+\frac{b\cdot c}{a}\pi^e\,,$$

$$x(z)=\frac{1}{a}\cdot(y(z)-\bar{y})-\frac{1}{c}\cdot\pi(z)-\frac{b}{a}\pi^e\,,$$

and then substitute $\pi(z)$ and $y(z)$ by the optimal values given in equations (3.1) and (3.2).

This is the first best. Nevertheless, the model assumes that society cannot implement policy directly and instead is forced to delegate the job to the central banker and the fiscal authority, the analysis of which follows next.

3.3.3 Agent's Optimization Problem

For the benchmark model, I partially disregard the role of the fiscal authority and focus exclusively on the central banker as the main agent. The monetary policymaker has a dual mandate: to stabilize prices, on the one hand, and to foster economic growth, on the other.[2] Thus, the central banker undertakes both of these tasks by choosing the level of effort e_π and e_y, producing q_π and q_y^{CB} correspondingly.

[2]Traditionally, the literature has viewed this second mandate as the full stabilization of the output gap. However, given the framework I consider below, it is difficult to maintain this interpretation of y because it would necessarily lead to the agent exerting negative effort in certain periods. In order to reconcile the model with the mainstream interpretation of y, one should view e no longer as the effort level but as the input chosen by the monetary authority who cannot choose perfectly its implementation but is subject to some error. The analysis that would follow is, however, similar to the one here presented. Thus, to avoid confusion and for the sake of simplicity, I shall assume that the second objective of the central bank is to pursue sustained GDP growth, rather than closing the output gap.

The agent optimization problem reads as follows:

$$\max_{q_\pi, q_y^{CB}} \quad u\left[w(q_\pi, q_y^{CB})\right] - \psi(e_\pi,\, e_y)$$

$$subject\ to$$

$$q_\pi = e_\pi + \nu_\pi\ ,$$

$$q_y^{CB} = e_y + \nu_y\ ,$$

where the agent features a utility function, which is separable in income and effort as it is standard in the moral hazard literature. Note further that $w(q_\pi,\, q_y)$ is the reward scheme, in order words, compensation given by society for accomplishing the job, and $\psi(e_\pi,\, e_y)$ is the personal cost the central banker incurs when exerting effort.[3]

Cost function specification I follow Holmstrom and Milgrom (1991)'s basic setup in their well-known multitask principal-agent analysis. They assume the agent's cost function adopts a flexible functional form that allows for different degrees of substitutability among the tasks undertaken by the agent:

$$\psi(e_\pi,\, e_y) = \frac{1}{2}\left(c_\pi e_\pi^2 + c_y e_y^2\right) + \rho e_\pi e_y \qquad \text{with} \qquad 0 \leq \rho \leq \sqrt{c_\pi c_y}\,.$$

There are three possible cases. At one extreme, suppose $\rho = 0$, implying that the two efforts are technologically independent. At the other extreme, $\rho = \sqrt{c_\pi c_y}$, meaning that the two efforts are perfect substitutes. Anything in between, $\rho \geq 0$, is referred to as the effort substitution problem, whereby increasing the effort on one of the tasks raises the marginal cost of effort on the other task. In order to formalize the existing trade-off faced by the agent when deciding whether to prioritize inflation targeting or output stabilization and to what extent, I consider the latter case.

Reward scheme specification The design of the reward scheme is the core assumption in this paper given that it shapes the incentives of the central banker to exert the effort required to achieve the first-best solution. In this sense, the reward is taken to be inversely proportional to the dispersion between society's desired value and its perception of agents' actions. In the case of inflation, society observes q_π directly while in the case of output society observes a combined measure of what both the central banker and the fiscal authority are implementing in order to pursue stable output growth.[4]

[3] To keep the analysis as simple as possible, I here consider an economic compensation i.e. wage, however, it is possible to build a dynamic model in which the central banker does not (only) care about the wage but is concerned about her chances of remaining in her position during the following period. This would involve incorporating career concerns incentives.

[4] The interpretation of the outcome of effort here deviates slightly from the standard principal agent problem. I shall not regard q_π as the resulting inflation rate in the economy but rather as

Thus, consider the following specification:

$$w(q_\pi, q_y) = \lambda \left[-\frac{1}{2}(q_\pi - \pi^*(z))^2 \right] + (1-\lambda) \left[-\frac{1}{2}(q_y - y^*(z))^2 \right] \qquad ,$$
$$\text{where} \qquad q_y = Q(q_y^{CB}, q_y^{FA})\,.$$

Given the above expression, instead of considering a central banker who receives a monetary incentive for accomplishing tasks, I present an agent who is punished for not performing her tasks correctly. Therefore, when optimizing, the central banker shall now consider the disutility of the punishment, instead of the utility of the compensation.

Next, consider $Q(q_y^{CB}, q_y^{FA})$, which indirectly relates the agents' efforts regarding output stabilization and the resulting, observable, measure of output which society uses to determine compensation. I assume that $q_y^{CB}$and q_y^{FA} are complements to each other to some extent. This eliminates incentives to completely free ride on the other agent and the possibility of assigning one task to one agent to achieve the first best. This paper explores different degrees of complementarity and how it affects model outcomes. In particular, I consider the general case of imperfect complements such that $Q(q_y^{CB}, q_y^{FA})$ is a Cobb-Douglas function.

Fiscal authority's behavior A crucial assumption is the extent to which the central banker knows policy choices made by the fiscal authority. For the benchmark set of results, I assume that q_y^{FA} is known by the central banker at the time of choosing monetary policy. This means that it has full information regarding the fiscal's authority behavior. I acknowledge that this assumption may be unrealistic because it requires the fiscal authority being the first mover or the central bank having full access to the fiscal authority's decision making process while the opposite does not hold. However, in such a context, the central banker does not need to form expectations, simplifying the analysis considerably.

Income utility specification Traditionally, the literature specifies a risk neutral agent which maximizes a constant absolute risk aversion (CARA) utility function. The agent only has control over the mean, and not the riskiness, of its compensation. In fact, the coefficient of risk aversion only affects how the principal designs the incentive contract. Given that the reward scheme specification considered above is not quadratic, for the sake of simplicity, I consider a risk neutral agent with a linear utility function in income.

The agent maximization problem looks as follows:

an informative measure of what the central banker is doing in relation to price stabilization (the central bank's target for inflation).

$$
\begin{aligned}
\min_{q_\pi, q_y^{CB}} \quad & u\left(\lambda\left[\frac{1}{2}(q_\pi - \pi^*)^2\right] + (1-\lambda)\left[\frac{1}{2}(Q(q_y^{CB}, q_y^{FA}) - y^*)^2\right]\right) \\
& + \frac{1}{2}\left(c_\pi e_\pi^2 + c_y e_y^2\right) + \rho e_\pi e_y \\
& \textit{subject to} \\
& q_\pi = e_\pi + \nu_\pi \, , \\
& q_y^{CB} = e_y + \nu_y \, .
\end{aligned}
$$

3.3.4 Solving the Agent's Optimization Problem

Recall from above that q_y^{FA} is known by the central banker when she solves its optimization problem, the utility function $u(\cdot)$ is linear and further simplify $q_i = e_i \; \forall i \in \{\pi, y\}$. Suppose that the target that society observes is: $q_y = (q_y^{CB})^\alpha (q_y^{FA})^{1-\alpha}$ i.e. imperfect complements case.

The central banker's maximization problem becomes:

$$
\begin{aligned}
\min_{q_\pi,\, q_y^{CB}} \quad & \lambda\left[\frac{1}{2}(q_\pi - \pi^*)^2\right] + (1-\lambda)\left[\frac{1}{2}((q_y^{CB})^\alpha (q_y^{FA})^{1-\alpha} - y^*)^2\right] \\
& + \frac{1}{2}\left(c_\pi q_\pi^2 + c_y (q_y^{CB})^2\right) + \rho q_\pi q_y^{CB}
\end{aligned}
$$

such that the first order conditions read:

$$
w.r.t \quad q_\pi : \qquad \lambda[q_\pi - \pi^*] + c_\pi q_\pi + \rho q_y^{CB} = 0 \, , \tag{3.3}
$$

$$
w.r.t \quad q_y^{CB} : \qquad \alpha \cdot (1-\lambda)[q_y - y^*]\left(\frac{q_y^{FA}}{q_y^{CB}}\right)^{1-\alpha} + c_y q_y^{CB} + \rho q_\pi = 0 \, . \tag{3.4}
$$

This is a system of non-linear equations with no closed-form solution, which complicates the interpretation of results. In order to build intuition, instead of solving for q_y^{CB} and q_π directly, I first consider and discuss the most extreme cases, by evaluating the first order conditions at particular parameter values of λ and q_y^{FA}. This will familiarize the reader with how the incentive mechanism works in the model.

Based on reward Suppose that society rewards the central banker based only on the result of inflation. This implies that $\lambda = 1$. The corresponding results are $q_\pi = \frac{\pi^*}{(1+c_\pi) - \frac{\rho^2}{c_y}} > 0$ and $q_y^{CB} = -\frac{\rho \pi^*}{(1+c_\pi)c_y - \rho^2} < 0$. It is difficult to rationalize the resulting negative target for output growth. So it might be convenient to restrict choice of q_y^{CB} to a value of at least zero, thus, obtaining $q_\pi = \frac{\pi^*}{(1+c_\pi)}$ and $q_y^{CB} = 0$. In any case, the interpretation is straightforward. Whenever the monetary policymaker is only rewarded for the task of inflation stabilization, she will completely prioritize the control of prices and will subordinate the choice of an output target i.e. the value q_y^{CB} is found as a residual using the constraints the central banker faces.

An important characteristic of this result is that it will hold for any specification of $Q(q_y^{CB}, q_y^{FA})$. This is intuitive; if the central banker is not compensated for the stabilization of output, from her point of view, the aggregation of individual output targets is completely irrelevant for the optimization problem.

Based on fiscal authority's behavior Let's now consider the fiscal authority's set of choices and how the central bank reacts to this information. Without formalizing the government's optimization problem, I consider analytically the two polar scenarios that are of interest from the central bank's point of view and represent graphically results for the intermediate case, where the fiscal authority chooses any value lying between the two extremes.

First, say the fiscal authority is well behaved. In other words, suppose the government or Treasury sets its target value exactly as the general public would i.e. $q_y^{FA} = y^*$. In this case, the first order conditions of the problem become:

$$\begin{aligned}
w.r.t \quad q_\pi: &\qquad \lambda[q_\pi - \pi^*] + c_\pi q_\pi + \rho q_y^{CB} = 0\,, \\
w.r.t \quad q_y^{CB}: &\qquad \alpha \cdot (1-\lambda)[(q_y^{CB})^\alpha (y^*)^{1-\alpha} - y^*]\left(\frac{y^*}{q_y^{CB}}\right)^{1-\alpha} + c_y q_y^{CB} + \rho q_\pi = 0\,.
\end{aligned}$$

Even though the fiscal authority is exerting the maximum level of effort, the central banker cannot fully free-ride on the government and free herself from its growth responsibility given the complementarity assumption. Thus, although there is some incentive to prioritize the first task over the second because a decent reward is already guaranteed at the time of decision making, such incentive is not a strong as one would initially expect.

Next, suppose the fiscal authority sits idle and completely fails to fulfill the task for which she is responsible. This implies that the outcome is $q_y^{FA} = 0$. Given this behavior, the central banker best responds by setting:

$$\begin{aligned}
q_\pi &= \frac{\lambda \pi^*}{\lambda + c_\pi}\,, \\
q_y^{CB} &= 0\,.
\end{aligned}$$

Clearly, if the fiscal authority does not comply with its job and the reward depends partly on what she does, there is no incentive for the central banker to devote time, effort or energy, which is costly, in fostering economic activity given that, in any case, the reward will ultimately depend on what is done relative to inflation. What is key to this result, however, is that the targets of both agents are aggregated using a Cobb Douglas function (imperfect complements), so that one agent choosing a zero target will lead to a zero aggregate target no matter what the second agent does.

Note that this outcome closely resembles that of a model with a reward scheme uniquely determined by the inflation outcome as discussed above. The intuition is that in both cases there is effectively no compensation for stabilizing output. The only difference lies in the role of parameter λ, which was set equal to one in the former case. The larger the value assigned to λ, the closer the inflation outcome will be to the optimal value.

Finally, let's perform full comparative statics. The previous cases are extreme and, therefore, difficult to reconcile with the real world. The next step to better understand how the choices of the fiscal authority influence the decisions made by the central banker is to evaluate the intermediate cases. I do this numerically by making use of the Implicit Function Theorem together with equations (3.3) and (3.4) to derive:

$$\frac{\partial q_y^{CB}}{\partial q_y^{FA}} = -\frac{\alpha \cdot (1-\lambda) \cdot (1-\alpha) \cdot \left(\frac{q_y^{CB}}{q_y^{FA}}\right)^{\alpha} \left[\left(\frac{q_y^{FA}}{q_y^{CB}}\right)^{1-\alpha} + [q_y - y^*](q_y^{CB})^{-1}\right]}{\alpha \cdot (1-\lambda) \cdot \left(\frac{q_y^{CB}}{q_y^{FA}}\right)^{\alpha-1} \left[\alpha \left(\frac{q_y^{FA}}{q_y^{CB}}\right)^{1-\alpha} - (1-\alpha)[q_y - y^*](q_y^{CB})^{-1}\right] - \frac{\rho^2}{\lambda + c_\pi}}. \tag{3.5}$$

While the above relationship is difficult to interpret, it can be shown that there is a threshold at which the derivative above changes sign representing a clear change in the incentive system. To show this graphically, Figures 3.1 to 3.4 plot the relationship between the central bank's optimal choice for q_y^{CB} and what the fiscal authority has previously done q_y^{FA} for the range $0 \leq q_y^{FA} \leq y^*$.

Table 3.1: Calibration

	Parameter	Value
λ	Inflation reward weight	0.6
α	Cobb-Douglas parameter	0.4
y^*	Optimal output level	3.5%
π^*	Optimal inflation rate	2%
c_y	Output effort cost	0.001
c_π	Inflation effort cost	0.001
ρ	Effort substitution parameter	0.4

To do this, Table 3.1 first presents the calibration of model parameters. Given that the objective of this numerical exercise is not quantitative, parameters values are not to be taken strictly.

Figure 3.1: Benchmark case

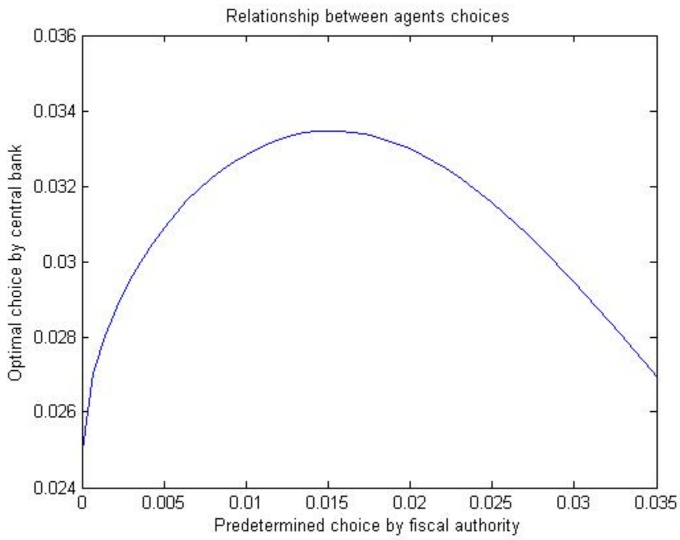

Notes: This figure plots the relationship between the central bank's optimal choice for q_y^{CB} and what the fiscal authority has previously done q_y^{FA} for the range $0 \leq q_y^{FA} \leq y^*$ using the benchmark calibration.

The values of y^* and π^* are implicitly given by the principal's maximization problem above, see equations (3.1) and (3.2). This involves assigning values to a larger number of parameters and acknowledging any realized shocks. For simplicity, I assume π^* is equal to 2%, consistent with the explicit inflation target of a large number of central banks around the world. For y^*, I take the US average GDP growth over the last sixty years. The costs are assumed to be small to minimize distortions and the degree of effort substitution taken to be in line with the mechanism design literature.

The two remaining parameters, λ and α, are key for the analysis. For the general case, I assume that the monetary policymaker's reward depends on the completion of both the tasks to which she is entitled, although inflation targeting has a slightly higher weight on the compensation scheme. This is relaxed in Figure 3.4. For the Cobb-Douglas coefficient, α, I consider three different cases, the benchmark case being that the target of the fiscal authority has a slightly higher weight in determining the overall output target observed by the principal.

Figure 3.1 depicts an inverted U-shape relation between the choice of the fiscal authority and the subsequent choice by the monetary authority.[5] The intuition is straightforward. The central banker must balance two conflicting forces, while

[5] Note that the graph does not apply for the extreme case $q_y^{FA} = 0$ which I have already discussed. The reason is that the graph here presented is derived from the agent's first order condition. However, if $q_y^{FA} = 0$, then the minimization problem the monetary policymaker faces is different and the results derived earlier apply. i.e. $q_y^{CB} = 0$

Figure 3.2: Alternative scenario I

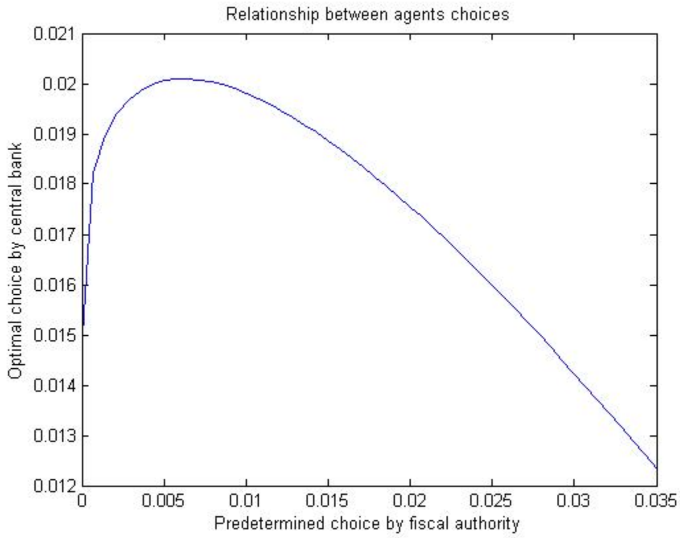

Notes: This figure plots the relationship between the central bank's optimal choice for q_y^{CB} and what the fiscal authority has previously done q_y^{FA} for the range $0 \leq q_y^{FA} \leq y^*$ under the assumption that $\alpha = 0.6$. The remaining parameters are calibrated as summarized by Table 3.1.

putting effort in one task increases the reward on that task, it simultaneously raises the cost of the other task. At the same time, she must coordinate with the government. If the choice of the fiscal authority is very far or very close to the optimal value, the central banker has less incentives to pursue economic growth, as the marginal benefit from exerting more effort in the second task is smaller than focusing on the first. In the former case, this occurs because the central banker's effort has to make up, in some way, for what the fiscal authority is not doing; the key being that benefits of this extra effort are shared while the costs are completely burdened by the monetary authority. In the latter case, the central banker is guaranteed a high compensation irrespective of her actions, such that the extra compensation that might be obtained by exerting extra effort is smaller than what would be obtained if the extra effort is devoted to price stabilization.

Figures 3.2 and 3.3 show that a similar pattern holds despite the change in parameter values. The inverted U-shape is preserved, although the distribution is shifted to the left, in the first case and to the right, in the second case. Finally Figure 3.4 shows the results derived above in the "based on reward" scenario. If there is no reward to output stabilization, then the central banker will completely disregard this task, no matter what the fiscal authority does.

In sum, the main result derived from the theoretical model is that excessive hawkishness of central banks with a dual mandate can be driven by two different factors: an unbalanced compensation scheme that prioritizes inflation over output

Figure 3.3: Alternative scenario II

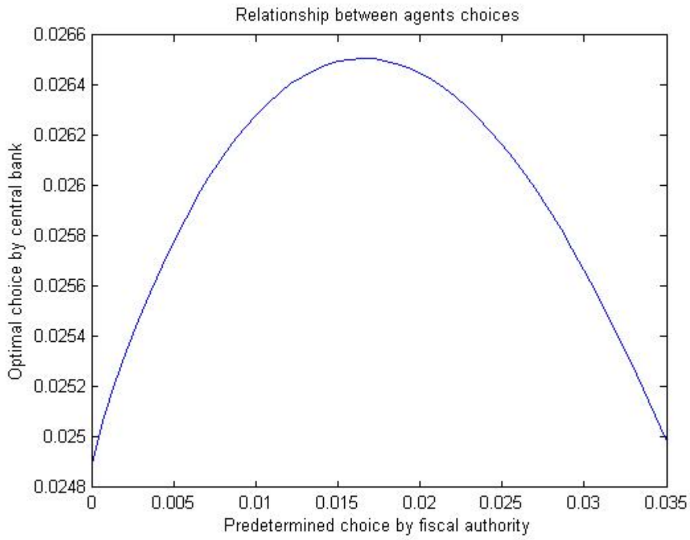

Notes: This figure plots the relationship between the central bank's optimal choice for q_y^{CB} and what the fiscal authority has previously done q_y^{FA} for the range $0 \leq q_y^{FA} \leq y^*$ under the assumption that $\alpha = 0.1$. The remaining parameters are calibrated as summarized by Table 3.1.

stabilization when rewarding central bankers for completing their job and a strategic response to the target decided by the fiscal authority with whom the central banker shares the responsibility of encouraging economic growth.

3.4 Discussion

This section relaxes some model assumptions and briefly discusses potential extensions. Formal analysis of such extensions, however, is left for future research. In parallel, I also comment on the main caveats of the current set-up.

3.4.1 Model extensions

While the model is successful in capturing the underlying structure of incentive interaction in policy making, it is a stripped-down representation of the decision making process. The most straightforward extensions involve incorporating time and uncertainty to the baseline framework.

First, the framework here considered is a one-period static model. A dynamic version of this model would, among other things, allow the reward scheme to include other motives in addition to the wage paid to the central banker. In other words, the benefit the central banker derives from performing her tasks is not only the economic compensation she receives, but also a higher probability of remaining in this position in future periods. This would involve incorporating career-concerns and reputation

Figure 3.4: Alternative scenario III

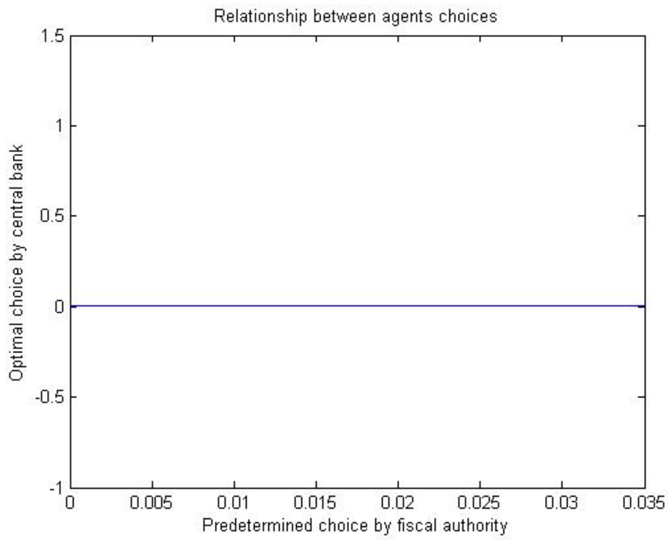

Notes: This figure plots the relationship between the central bank's optimal choice for q_y^{CB} and what the fiscal authority has previously done q_y^{FA} for the range $0 \leq q_y^{FA} \leq y^*$ under the assumption that $\lambda = 1$. The remaining parameters are calibrated as summarized by Table 3.1.

elements à la Holmström (1999) in which agents care about the public perception of their talent because current performance affects future compensation. Note, however, that results would remain to be qualitatively equal; if the probability of staying in office is proportional to an evaluation of performance which gives a greater weight to inflation stabilization, then the central banker will again focus predominantly on this task.

The second and, possibly, most interesting extension incorporates uncertainty to the model by assuming that both agents face the corresponding optimization problem simultaneously. By dropping the full information assumption, the central banker is forced to form expectations of the fiscal authority's choice. In this context, the degree of risk aversion of the monetary authority will play a key role.

Intuitively, the more unpredictable the government's behavior is, the less the central banker will want to rely on its support and, thus, the more she will focus on the task that depends uniquely on her: inflation stabilization. The main conclusion would then be modified: greater uncertainty regarding government's decisions leads to greater hawkishness given central bankers' risk aversion. This extension is particularly relevant for the case of developing economies and its formalization would allow to empirically measure the magnitude of its effect.

3.4.2 Challenges

Before concluding, there are two unresolved challenges that are worth acknowledging.

First of all, the model disregards the incentives behind the fiscal authority's decision. The paper takes the output target as given without formalizing the corresponding agent's optimization problem, which would possibly have to take into account the target (or expectations of such) adopted by the central bank. This would yield some sort of feedback loop that would complicate the analysis further. The current approach is based on the assumption that among all the incentives that shape the decisions made by a government, the behavior of the central bank does not play a prominent role, and, thus, can be ignored without loss of generality.

Secondly, this is a purely descriptive model given that I withdraw from explicitly designing the punishment that would induce the choice of the efficient outcome of society. The theoretical framework is used to describe and formalize the motivations, incentives and interactions that influence the implementation of monetary policy in an intuitive way. A quantitative version of the model that would allow to take it to the data is currently missing.

3.5 Conclusion

This paper looks at the incentives that determine how much weight policymakers attach to the different targets in an economy with a dual mandate. Based on a multitask principal-agent problem framework as in Holmstrom and Milgrom (1991), I construct a model of monetary policy where the central banker is viewed as an agent that is entrusted with two conflicting tasks: inflation stabilization and economic growth, the responsibility of which is shared with the government. The key to the model is the design of the compensation scheme. More specifically, whether it induces the monetary authority to divide its effort optimally among the two tasks. Excessive hawkishness of central banks with a dual mandate is rationalized by this framework. The central banker will tend to prioritize inflation stabilization in any of two scenarios: on the one hand, whenever the compensation scheme is unbalanced, i.e performance measures overstate inflation over output outcomes; on the other hand, if the choices of the fiscal authority are either too close or too far from the optimal target.

Extending the model to include uncertainty and risk aversion is an interesting avenue for future research. This would involve policymakers forming expectations over the government's actions, bringing the model closer to reality, especially for the case of developing economies. Empirical analysis relating the behavior of central banks and the degree of transparency of governments should then follow.

3.6 Bibliography

Allen, F. (1985): "Repeated Principal-Agent Relationships with Lending and Borrowing," *Economics Letters*, 17, 27–31.

Backus, D., J. Driffill, et al. (1985): "Credibility and Commitment in Economic Policy," *Working Paper*.

Barro, R. J. and D. B. Gordon (1983a): "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, 91, 589–610.

——— (1983b): "Rules, Discretion and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics*, 12, 101–121.

Bernheim, B. D. and M. D. Whinston (1986): "Common Agency," *Econometrica*, 923–942.

Chortareas, G. E. and S. M. Miller (2003): "Central Banker Contracts, Incomplete Information, and Monetary Policy Surprises: In Search of a Selfish Central Banker?" *Public Choice*, 116, 271–295.

Dixit, A. and L. Lambertini (2003): "Interactions of Commitment and Discretion in Monetary and Fiscal Policies," *American Economic Review*, 93, 1522–1542.

Dye, R. A. (1986): "Optimal Monitoring Policies in Agencies," *The Rand Journal of Economics*, 339–350.

Fudenberg, D., B. Holmstrom, and P. Milgrom (1990): "Short-Term Contracts and Long-Term Agency Relationships," *Journal of Economic Theory*, 51, 1–31.

Hart, O. D. and B. Holmstrom (1987): "The Theory of Contracts," in *Advances in Economic Theory: Fifth World Congress*, Cambridge University Press.

Holmström, B. (1999): "Managerial Incentive Problems: A Dynamic Perspective," *The Review of Economic Studies*, 66, 169–182.

Holmstrom, B. and P. Milgrom (1991): "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics & Organization*, 7, 24.

Kydland, F. E. and E. C. Prescott (1977): "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85, 473–491.

McCallum, B. T. (1986): "The Optimal Inflation Rate in an Overlapping-Generations Economy with Land," in *New Approaches to Monetary Economics*, Cambridge University Press.

Persson, T. and G. Tabellini (1993): "Designing Institutions for Monetary Stability," in *Carnegie-Rochester Conference Series on Public Policy*, Elsevier.

Reis, R. (2013): "Central Bank Design," *Journal of Economic Perspectives*, 27, 17–44.

Rogoff, K. (1985): "The Optimal Degree of Commitment to an Intermediate Monetary Target," *The Quarterly Journal of Economics*, 100, 1169–1189.

Svensson, L. E. (1997): "Inflation Targeting in an Open Economy: Strict or Flexible Inflation Targeting?" *Working Paper*.

——— (2001): "Price Stability as a Target for Monetary Policy: Defining and Maintaining Price Stability," in *The Monetary Transmission Process*, Springer.

Taylor, J. B. (1986): "Improvements in Macroeconomic Stability: the Role of Wages and Prices," in *The American Business Cycle: Continuity and Change*, University of Chicago Press.

——— (2011): "Legislating a Rule for Monetary Policy," *The Cato Journal*, 31, 407.

Taylor, J. B. and J. C. Williams (2010): "Simple and Robust Rules for Monetary Policy," in *Handbook of Monetary Economics*, Elsevier.

Walsh, C. E. (1995): "Is New Zealand's Reserve Bank Act of 1989 an Optimal Central Bank Contract?" *Journal of Money, Credit and Banking*, 27, 1179–1191.