The London School of Economics and Political Science

Essays on Applied Microeconomics

Jack Welcome Fisher

A thesis submitted to the Department of Economics of the London School of Economics and Political Science for the degree of Doctor of Philosophy. London, June 9, 2023.

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. The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent. I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party. I declare that my thesis consists of 52,254 words.

Jack Welcome Fisher

June 9, 2023

Statement of Conjoint Work

I confirm that Chapter 3, "Refinancing Cross-Subsidies in the Mortgage Market", was jointly co-authored with Alessandro Gavazza, Lu Liu, Tarun Ramadorai, and Jagdish Tripathy. A version of this paper has been published as Bank of England Staff Working Paper No. 948. This statement is to confirm that I contributed 20 percent of this work.

Jack Welcome Fisher

June 9, 2023

Acknowledgments

I am deeply grateful to Johannes Spinnewijn for his time, patience, and invaluable advice over the course of my PhD. I have learned so much from his supervision and mentorship, which began prior to my graduate studies when I was a research assistant at STICERD. During and after this period as a research assistant, I also learned a great deal about the research process from Camille Landais to whom I extend my gratitude.

I would like to thank my co-advisors Alan Manning, Alessandro Gavazza, and John Van Reenen. Their enthusiasm, guidance, and differing perspectives have been a source of motivation and greatly benefited my research.

There are too many people to mention who have generously engaged with my work at the LSE and elsewhere. To name a few, I am extremely thankful to: Richard Blundell, Jakob Brounstein, Gaia Dossi, Arnaud Dyèvre, Andrew Ellis, Shadi Farahzadi, Antoine Ferey, François Gerard, Giulia Giupponi, Daniel Gottlieb, Ben Handel, Jonathan Hazell, Xavier Jaravel, Henrik Kleven, Jonathan Kolstad, Lu Liu, Stephen Machin, Isabela Manelici, William Matcham, Guy Michaels, Virginia Minni, Canishk Naik, William Parker, Jörn-Steffen Pischke, Christopher Pissarides, Tarun Ramadorai, Hugo Reichardt, Daniel Reck, Nina Roussille, Emmanuel Saez, Mark Schankerman, Benjamin Schoefer, Kate Smith, Dmitry Taubinsky, Catherine Thomas, Jose Vasquez, Maria Ventura, and Jordi Blanes i Vidal.

Mentors during my undergraduate and master's studies have also had a huge impact on my research. In particular, I owe am indebted to Subir Chattopadhyay and Neil Rankin at the University of York, and Anthony Venables and David Vines at the University of Oxford.

I also want to acknowledge the incredible support from individuals at the LSE: Lubala Chibwe, Jane Dickson, Rhoda Frith, Lakmini Staskus, Emma Taverner, Anna Watmuff, and Mark Wilbor, as well as many others. My graduate studies would also not have been possible without financial assistance from the Economic and Social Research Council and the Royal Economic Society.

Lastly, this thesis would not have been possible without the support of my friends and family; especially Meg, Linda, and Nigel.

Abstract

This thesis is comprised of three chapters that study individual and household behavior in markets where they are subject to choice frictions. Chapters 1 and 2 study workers in the gig economy, while chapter 3 examines households' decision-making in a mortgage market. All the chapters combine theoretical and empirical analysis.

In chapter 1, I use administrative data spanning the UK's food delivery market to estimate worker surplus in this typical gig labor market. Evidence that workers learn about their own value of gig work over time allows for the identification of the joint distribution of gig work valuations and outside options. Estimates imply a median monthly surplus for gig workers equal to one-third of the median employee's monthly income (or £673). In terms of policy, the analysis suggests that attaching fixed benefits to gig work, such as those mandated by California's Proposition 22, is unlikely to raise worker welfare if platforms pass on some of the associated costs through, for example, lower hourly earnings.

Chapter 2 investigates an important dimension of the typical flexibility versus security trade-off that is used to frame gig work, as well as self-employment more generally. Namely, behavioral frictions that prevent workers from fully exploiting flexibility. I study the welfare cost of behavioral biases in intensive margin labor supply decisions for a group of self-employed workers who are free to pick their hours. In the spirit of Chetty-Looney-Kroft (2009), I estimate a salient—large and positive—daily Frisch elasticity to characterize preferences and contrast this with typical daily labor supply, which is subject to behavioral biases. A new sufficient statistics formula translates these deviations into welfare losses ranging between two and six percent of daily income.

Lastly, chapter 3 studies cross-subsidies in the UK mortgage market that are caused by heterogeneity in the timeliness of household refinancing. We build and estimate a model of household mortgage refinancing using rich administrative data on the stock of outstanding mortgages in the UK. The results imply sizeable cross-subsidies from poorer to richer households. This work highlights how the design of household finance markets can entrench financial inequalities.

Contents

List of Figures

List of Tables

Chapter 1

Worker Welfare in the Gig Economy

1.1 Introduction

In both the US and the UK, the last decade has seen rapid growth in the number of solo self-employed workers who connect with customers via digital intermediary platforms (Bertolini et al., 2021; Collins et al., 2019; Katz and Krueger, 2019). The vast take up of this type of work—gig work—suggests that individuals are enjoying a significant surplus, but this casual observation is at odds with qualitative evidence that many workers have negative experiences of gig work (Broughton et al., 2018; Dubal, 2019; Ravenelle, 2019).

In light of a desire to better protect workers in the gig economy with new regulation (Adams et al., 2018; Dubal, 2021; Goldman and Weil, 2021; Harris and Krueger, 2015; Kolsrud, 2018; Prassl, 2018), a comprehensive understanding of the gig work surplus is of first order importance for policymakers. However, little is known about the size of this surplus and how to assess the impact of prospective changes to the gig work environment. For example, what are the welfare effects of rebalancing hourly earnings and fixed benefits, such as sick pay and health insurance? Or of reducing participation costs?

There are three reasons for our limited understanding of the gig work surplus. Firstly, gig work differs from traditional employment in many ways (Boeri et al., 2020), which makes it hard to estimate workers' welfare in the gig economy relative to what they would enjoy in its absence. Gig work involves flexible hours (Chen et al., 2019; Katsnelson and Oberholzer-Gee, 2021; Mas and Pallais, 2017), uncertain pay (Angrist et al., 2021; Cook et al., 2021; Parrott and Reich, 2018), jobs of varying difficulty (Athey et al., 2021), and often takes place at specific times in the worker lifecycle (Jackson, 2019; Koustas, 2018, 2019; Katz and Krueger, 2017). Although researchers have investigated the implications of these amenities individually, there has not yet been a holistic analysis that subsumes these factors.

Secondly, data limitations impose an obstacle; the best data on the gig economy is siloed within different platforms. If a researcher is able to access data from one platform, then observed workers may be selected based on the platform's characteristics (*e.g.*, the structure of financial incentives and the user interface) and they may switch between platforms so that only a partial picture of labor supply is revealed (Caldwell and Oehlsen, 2021). Thirdly, gig workers face frictions and behavioral biases (Camerer et al., 1997a; Chen et al., 2020; Fisher, 2022; Thakral and Tô, 2021), which pose problems to revealed preference approaches. Indeed, the unusual traits of gig work likely make it hard for workers to judge their own surplus and, in turn, to reveal their preferences.

This paper tackles these problems head-on, and estimates the size and structure of the gig work surplus through the joint distribution of gig work valuations and outside options for a typical gig labor market in the UK. I introduce new data sources, evince novel patterns in gig work participation, and bring these contributions together with a structural model of gig work, which can evaluate worker welfare in counterfactual scenarios. The framework uses a new identification strategy, which leverages misperceptions and learning to infer workers' outside options from their endogenous exit decisions.

The results imply a large surplus with the median gig worker enjoying a monthly surplus equal to one-third of the median UK employee's monthly income (or £673). Aggregating to the market level, this equates to a £15bn annual surplus for workers, which mainly accrues to less than full-time participants. The distribution of the surplus causes a trade-off for policymakers between ensuring benefits for full-time gig workers and maintaining gig work's appeal to most participants. The analysis also reflects some workers' negative experiences through misperceptions, which depress the gig work surplus by 17% via an allocative inefficiency.

The ideal experiment to identify an individual's total surplus from gig work would offer workers increasing amounts of money which, if accepted, would prevent them

Table 1.1: Gig Work Surplus Matrix

		Outside Option ν	
		High	Low
Valuation θ	High	$\theta_H - \nu_H \approx 0$	$\theta_H - \nu_L > 0$
	Low	$\theta_L - \nu_H < 0$	$\theta_L - \nu_L \approx 0$

from working in the gig economy. The lowest amount at which a worker accepts is equal to the worker's surplus. Absent this experiment, it is desirable to infer the size and structure of surpluses from gig workers' observable actions.

The crux of the problem is that the difference between workers' valuations of gig work, conditional on working in the gig economy, and their outside options determines their gig work surplus. There is significant unobservable variation in both dimensions. For example, in terms of valuations, some individuals engage in less than 40 hours of gig work a month and others regularly work more than 40 hours a week; these people clearly extract very different value from the gig economy. Further, they likely face contrasting outside options. Part-time workers may engage in leisure activities absent the gig economy, while full-time workers may substitute to another full-time job or a patchwork of part-time jobs.

To make this point clear, table 1.1 considers a gig economy with four types of workers who are defined by a combination of two characteristics: whether their valuation of gig work is high θ_H or low θ_L, and whether their outside option is high ν_H or low ν_L. Individuals with low outside options and high valuations benefit the most from gig work, and those with high outside options and low valuations would be better off outside of the gig economy. For the remaining individuals, the surplus from gig work is ambiguous but small. In this rudimentary setting, identification of workers' valuations $\{\theta_H, \theta_L\}$ and outside options $\{\nu_H, \nu_L\}$, and the proportion of types in the population allows for a description of the gig work surplus $\mathbb{E}[\theta - \nu | \theta > \nu]$.

In a more flexible setting, the task is to identify the joint distribution of gig work valuations and outside options. To do so, I use a new data source and focus on a typical part of the gig economy in the UK: Food delivery by motorcycle. This industry makes up around one fifth of the UK's gig economy (Cornick et al., 2018) and exhibits features that are characteristic of gig work more generally, such as payment on a per job basis and flexible hours (from hereon, "gig economy" will generally refer to this part of the gig economy).

The new data source is administrative data from a vehicle insurer (the "firm"), which provides *mandatory* insurance to a sizeable share of gig workers in this market. The insurance primarily provides cover for damage to a third party that occurs while working, but is best seen from the worker's perspective as a necessary cost. The six largest food delivery platforms send information on each delivery to the firm to facilitate its insurance policy offering. As a result, this data avoids the pitfalls of working with an individual platform.

The firm offers insurance policies with either a variable or fixed premium. The fixed policy is preferable when workers expect to work many hours, while the variable policy minimizes costs for workers who intend to work few hours. Therefore, a cost-minimizing worker's policy choice contains information about their expectation of the hours that they will work in the gig economy. I contrast policy choice with realized hours using the firm's data and show that one in five workers make non-cost-minimizing decisions. Broadly, the lack of cost-minimization could be due to either *ex ante* misperceptions of future hours, or *ex post* shocks that affect hours after the policy choice has been made.

Ex post shocks would suggest that subsequent dynamic behavior is not related to cost-minimizing policy choice. Instead of this—and consistent with workers suffering from misperceptions—optimistic individuals (*i.e.*, those who do not work enough to explain their policy choice) reduce their hours after they enter, and exit faster. Conversely, pessimistic workers (*i.e.*, those who work too much to account for their policy choice) increase their hours initially, and exit at a slower rate. Different factors could drive misperceptions. For example, lower than expected customer demand that leads to low earnings, or unexpectedly high running costs of workers' vehicles—both find support in a survey of gig workers that I present in this paper.

Intuitively, if an individual's valuation of gig work determines the hours that they work, then misperceptions of hours reveal misperceptions of valuations. In this context, the dynamic patterns of participation and engagement in the gig economy also provide evidence of learning. Optimistic individuals enter the gig economy in anticipation of high value but, in learning that this is not the case, reduce their hours and exit faster, while the opposite is true for pessimists. In line with the hypothesis that learning ameliorates misperceptions, half of survey respondents subscribe to learning "a lot" about the "costs versus benefits" of gig work.

Misperceptions and learning lead to a mirror image of the ideal experiment outlined above. Consider an individual who is overly optimistic about their valuation of gig work. They will enter the gig economy and soon discover that their optimism was misplaced. Over time, they will reduce their hours as they learn about their true valuation and, if their perceived valuation drops below their outside option, they will exit the gig economy. In essence, the value of gig work is incrementally reduced, whereas the ideal experiment incrementally increases the outside option. Therefore, outside options can be identified as workers exit at the point where the trajectory of their perceived valuation crosses their outside option.

Equipped with this logic, the task is to estimate money metric valuations of gig work, workers' perceptions of valuations, and the learning process. Identification requires more than data on intensive and extensive margin participation in the gig economy because it does not permit a way to classify individuals based on their misperceptions. Observing workers' insurance policy choices is important to partition them into groups which are informative of their misperceptions; a worker who enters the gig economy and reduces their hours, before eventually exiting, would only indicate optimism if we take learning as given, which is itself something that needs to be evinced. This additional margin of variation is sufficient to identify workers' surpluses from gig work provided some structure, which is necessary to model the complicated environment that workers face. It is also attractive given the ability of the data to speak simultaneously to different features of the economic environment and the usefulness of non-marginal counterfactual evaluations (Mahoney, 2022).

I develop a model of individuals' participation, insurance policy choice, and hours of work in the gig economy to reflect the patterns in the data. Importantly, it allows for heterogeneity in three key areas: both outside options and valuations can vary across workers, and workers differ in their misperceptions of their valuations. The model explicitly connects the intensity of engagement in the gig economy, measured by hours, to valuations of gig work, and describes the learning process.

Workers decide to participate in the gig economy, if they perceive that its value exceeds their outside option. Upon entering, they select either a fixed or variable insurance policy based on their expected hours, which are driven by their perceived valuation. Then, workers learn about their true valuation as they partake in gig work and adjust their hours accordingly since workers' perceptions always drive their engagement. Finally, workers will leave the gig economy, if their perceived value falls below their outside option, or if they receive an exogenous shock. Throughout, individuals' valuations are subject to *ex post* shocks to account for their influence when bringing the model to the data.

Simulated method of moments (SMM) provides estimates of workers' valuations, outside options, and misperceptions. I specify individual-level heterogeneity to follow a joint log-normal distribution, which helps to capture the skewed distribution of hours in the data and allows for economical computation with a convenient but flexible pattern of correlations between worker characteristics. Empirical moments from the administrative data, coupled with external moments on the employment share of the gig economy and labor supply elasticities, identify the model's parameters in conjunction with structural and stationarity assumptions.

The model and estimates of its structural parameters imply a large gig work surplus. The typical worker enjoys a monthly surplus of £1,066, over one third of mean employee's monthly earnings in the UK. But this masks huge dispersion (SD £1,775) with some workers extracting thousands of pounds of surplus from gig work, while others are on the margin of participating. At an hourly rate, the analysis implies workers enjoy 70 to 80% of their wage as a surplus. The bulk of the gig work surplus—55%—is received by workers who work less than 60 hours per month in the gig economy. This result is driven by the fact that the vast majority of individuals only dabble in gig work, and despite larger average surpluses for high-hours individuals.

I highlight several factors that could explain the large gig work surplus. Most gig workers are in the bottom half of the income distribution and often receive negative income shocks prior to entering (Bernhardt et al., 2022; Cornick et al., 2018; Koustas, 2018), which points towards a high marginal utility of income. Workers who especially value gig work amenities are more likely to select into participation. Further, these individuals' outside options may be particularly low. Alternative employment opportunities that can flexibly adapt to changing schedules are rare, and many gig workers are non-nationals who speak English as a second language, which may reduce other earning options.

The concentration of the aggregate gig work surplus amongst workers at the low end of the hours distribution poses a difficulty for policymakers. Broadly, regulators want to compel platforms to offer benefits for regular gig workers.[1] This comes at a cost to platforms, however, some of this will be passed onto workers. Moreover, this cost will likely be borne by all gig participants because platforms cannot *a priori* distinguish who will qualify for benefits, and multi-homing across platforms undermines any targeted incidence. Therefore, if platforms pass on some of the cost through, for example, an hourly wage penalty, the gig work surplus falls sharply because this severely hurts low-hours workers, who generate the majority of the surplus. In other words, policymakers face a steep trade-off between ensuring benefits for full-time gig workers and maintaining gig work's appeal to the majority of participants.

A counterfactual policy evaluation that is calibrated to match aspects of California's Proposition 22 crystallizes this point. In particular, I model the introduction of mandatory benefits for workers who reach certain hours thresholds. I find that such a policy reduces worker welfare if even half of the cost is born by workers through an hourly wage penalty. Yet, this intervention can increase the gig work surplus by up to 11% if there is minimal incidence on workers. These impacts rationalize a complementary minimum wage requirement, which could obstruct firms from passing on the cost of mandated benefits to workers.

In the same vein, I use the rich heterogeneity in the model to consider an innovation that reduces the fixed costs of gig work. As a concrete example, I use the introduction of the variable policy, which affects welfare through four channels: (i) it can help existing gig workers save money, (ii) it allows individuals who wish to work only a few hours in the gig economy to participate, (iii) it can attract pessimistic workers who discover that they value gig work more than expected, and (iv) it can attract optimistic workers who, with hindsight, should stay out of the gig economy. Overall, the introduction of the variable policy increases welfare by 4.7% primarily through increased participation. These numbers correspond to an annual aggregate welfare gain of £709mn for workers, which represents a significant return to innovation and policy in this direction.

Next, I turn to the effect of misperceptions on the gig work surplus. Although the

[1] This may be through new legislation or clarifications around the legal tests that are used to define employment status. See efforts along these lines in, for example, the US, Europe, and India.

population correctly perceives the value of gig work on average, variation in misperceptions gives rise to an allocative inefficiency that stifles the gig work surplus. Namely, pessimistic workers who would receive a positive surplus do not participate, and overly optimistic participants may lose out compared to their outside option. Absent misperceptions the gig work surplus would be 21% higher, which stems equally from eradicating optimistic and pessimistic misperceptions. Notably, even halving the standard deviation of misperceptions attains three quarters of the first-best surplus, which suggests increasing transparency and information around gig work conditions could be a fruitful area for policy.

Relatedly, the results in this paper are congruent with the negative qualitative evidence surrounding gig work. Quantitatively, 45% of each new entering cohort is overly optimistic about gig work and would not enter but for misperceptions. Over the course of these individuals' tenure in the gig economy, they lose on average £1,183 relative to their next best alternative. Fortunately, learning ameliorates these losses such that misperceptions are not severe on average amongst participants.

Literature review. The rise of the digital gig economy has spurred a renewed interest in alternative work arrangements (Boeri et al., 2020; Mas and Pallais, 2020). This paper most closely relates to studies by Chen et al. (2019) and Chen et al. (2020) that use data on Uber rideshare drivers. Both papers aim to estimate the additional surplus that workers receive due to flexible work. To do so, they estimate workers' reservation wages and the market's wage hour-by-hour. The sum of differences between the market and reservation wages over a given schedule corresponds to their measure of a surplus. Then, the value of flexible work is calculated as the difference between the surplus under the flexible schedule and an alternative. They find flexibility almost doubles the gig work surplus on average. Notably, these estimates are much higher than those from discrete choice experiments (Datta, 2019; Mas and Pallais, 2017).

This paper's relative contribution is threefold. Firstly, the welfare gains here correspond to the gig work surplus rather than the additional surplus gig workers receive due to flexible hours. The welfare estimates in this paper plausibly constitute a catch-all surplus that accounts for many of the different features associated with gig work, including but not limited to the value of flexible work. Secondly, the approach in this paper lends itself to the evaluation of prospective policies that aim to improve worker welfare. In particular, the paper takes participation in the gig economy seriously by estimating workers' outside options through a structural model. Thirdly, the data and methodology used in this paper provide a novel way to identify misperceptions about the value of gig work, which leads to the possibility that some workers may be made worse off by entering the gig economy. This mechanism is consistent with qualitative evidence that many individuals are disappointed by what the gig economy offers, but is missing in the economics literature on gig work (Broughton et al., 2018; Dubal, 2019; Ravenelle, 2019).

Naturally, this work builds on many other studies of the gig economy. For example, numerous papers have documented the rise of the gig economy (Bernhardt et al., 2022; Collins et al., 2019; Katz and Krueger, 2019) and carefully assessed its underlying forces (Abraham et al., 2019; Cullen and Farronato, 2021; Ganserer et al., 2022; Garin et al., 2022). There has also been a great deal of descriptive work on the motives and demographics of gig workers (Garin and Koustas, 2021; Hall and Krueger, 2018; Chen et al., 2022), which can explain the wide range of surpluses that gig workers enjoy.

Relative to these studies, I establish novel patterns in gig work hours and survival by leveraging a new cost-structure choice. I argue these results are strongly suggestive of misperceptions and learning. Concurrent work by Pires (2022) also finds evidence of these phenomena in a survey of US gig workers. Moreover, I translate heterogeneity in worker choices over hours, exit, and policy into a distribution of surpluses via a structural model that captures workers' value of gig work and their next best alternative.

Work mediated by digital platforms inherently involves many different parties and this paper only considers the worker's perspective. There is a body of work that considers the welfare effects for customers (Cohen et al., 2016), the incentives of platforms to manage either side of the market (Akbarpour et al., 2021; Castillo, 2020; Hall et al., 2021; Rochet and Tirole, 2006; Weyl, 2010), and how surpluses are shared between consumers and workers in the context of the knowledge gig economy (Stanton and Thomas, 2021).

There are other parts of the labor economics literature that connect with this research, such as the long history of estimating labor supply elasticities (Blundell and MaCurdy, 1999). These estimates serve to discipline behavioral responses in the structural model. Specifically, I lean on labor supply estimates for solo self-employed taxi drivers from Fisher (2022), which also provides evidence of information frictions affecting labor supply decisions. Further, more recent studies assessing workers' outside options (Caldwell and Harmon, 2019; Caldwell and Danieli, 2020) and their perceptions of these alternatives (Jäger et al., 2022) relate closely and provide useful benchmarks for the results in this paper.

Lastly, work from the industrial organization field on usage-based pricing (or two-part tariffs, or second-degree price discrimination) provides a useful foundation for the policy choice modeling in this setting (DellaVigna and Malmendier, 2006; Economides et al., 2008; Goettler and Clay, 2011; Grubb and Osborne, 2015; Hoffman and Burks, 2020; Lambrecht and Skiera, 2006; Nevo et al., 2016). I contribute to this literature by demonstrating that it is possible to identify heterogeneous outside options with misperceptions, learning, and endogenous exit. The logic is similar to how Bresnahan and Reiss (1990, 1991) exploit variation in the number of entrants under different market conditions to estimate entry costs.

The paper proceeds as follows. Section 1.2 discusses the institutional setting and the data for this study. Section 1.3 presents a series of reduced form facts, which motivate the model presented in section 1.4. Section 1.5 explains how this model is brought to the data, while section 1.6 discusses the estimates' implications for the gig work surplus and counterfactual scenarios. Finally, section 1.7 concludes and presents complementary areas of future research.

1.2 Empirical Setting

This section describes the institutional environment in which gig workers operate, the data available, and the sample that I use for the analysis. To summarize, I will study motorcycle food delivery carried out by solo self-employed workers in the UK. These workers' experiences are emblematic of the broader gig economy in that platforms mediate their work, they are free to enter and exit, they have flexible hours, and they face uncertain wages. The data source is a firm that provides insurance to a sizeable share of this market, and collects administrative data on these individuals from many different platforms. I complement this data with a new survey of workers' experiences in the gig economy.

1.2.1 Institutional Details

Some of the most visible forms of gig work involve moving passengers and goods on the road; Toyota Prii[2] with smart phones fixed on dashboards and motorcycles adorned with insulated food delivery boxes are now quintessential sights for many cities around the world. Indeed, Cornick et al. (2018) find that this makes up over half of gig work in the UK, which is the setting for this study.

In this sense, I will focus on an exemplary part of the gig economy: food delivery by motorcycle. Specifically, solo self-employed workers who carry out this job for intermediary, digital platforms will be the subject of this paper. Like many in the gig economy, these individuals are free to onboard, and pick their own hours and location of work. They are generally paid a set fee per job but, when this is combined with fluctuating demand and supply, as well as other shocks (*e.g.*, traffic and waits at restaurants), they receive an uncertain wage.[3] Further, gig workers are self-employed so they are entitled to very few employment rights beyond health and safety and discrimination protections. For example, they do not receive sick pay and are not guaranteed a minimum wage.

A particular job, if accepted, requires the worker to drive to a restaurant, pickup a meal, and then deliver the meal to the customer. Platforms differ in the ways that they provide information and offer compensation. For example, some platforms tell workers where the customer is located prior to the acceptance of a job, while others only disclose the location of the restaurant. Compensation often adjusts to the distance of a job but this is only done in a coarse fashion.

Importantly, individuals working on UK roads *must* have an enhanced level of vehicle insurance. This additional insurance is called Hire and Reward (H&R) insurance and is a necessity for many gig workers, including motorcycle food delivery workers. This insurance covers damage to third parties while working and further coverage can be purchased to protect one's own vehicle under certain circumstances (*e.g.*, fire and

[2] The plural of Prius by popular choice, see https://pressroom.toyota.com/toyota-announces-the-plural-of-prius/.

[3] Some platforms also provide other financial incentives, such as a bonus for making a set amount of deliveries within a month, but these are uncommon and information on their details is lacking.

theft). This insurance does not cover the food being delivered, which is dealt with by the intermediary platform.

From the perspective of a gig worker, this imposition can be seen as an unavoidable cost to be minimized. The H&R market offers insurance in two forms: variable and fixed. Variable policies are paid approximately by the hour, while the fixed policies insure workers for a 30 day period and are paid for upfront. From a cost-minimization perspective, if one expects to work few hours over the next 30 days, then they should prefer the variable policy since it would not be economical to pay for a full 30 days of coverage. Conversely, the fixed policy is preferable when individuals expect to work many hours. Both policies are easy to use; the fixed policy is paid as a direct debit and the variable policy is paid for via a digital wallet, which can be auto-topped up from workers' bank accounts. When workers choose between either the fixed or variable policy online, both options appear equally prominent, side-by-side with their premiums listed.[4]

1.2.2 Data

The data for this paper comes from a H&R insurer (the "firm") that offers both the variable and fixed policy to prospective gig workers. The firm receives data from many different intermediary platforms in order to facilitate its insurance policies and, therefore, does not suffer from individuals selecting into or switching between work providers. Further, the firm provides insurance to a significant share of gig workers in the food delivery market.

The data contains information on jobs completed by workers and their fixed or variable policy choice, along with the premiums they faced and some worker covariates. In particular, the data contains information about the length of jobs, when they took place, a unique worker identifier, the age and gender of the worker, the type of insurance policy, and the premium. The main omission from the data is worker compensation.

Given the choice environment that workers face *vis-á-vis* a 30 day policy, and the aim of estimating a longer-term and broader surplus from gig work, I aggregate the

[4]Appendix 1.A presents additional evidence which supports the claim that workers are primarily concerned with cost-minimization. In particular, it shows that individuals who change policy tend to switch to more economical policies. This appendix also contains a more detailed discussion of the potential differences between the fixed and variable policies, and ulterior influences that may affect policy choice, such as variance in income.

data from the worker-job-level to the worker-month-level to construct, for example, a monthly hours worked variable.[5] The monthly data also has the advantage that it reduces the influence of high-frequency shocks that affect workers' labor supply decisions.[6] I treat a worker's first appearance in the dataset as their first entrance into the gig economy. While it is possible that workers may have already undertaken gig work, the rarity of policy switchers within the firm's data suggests that switching across firms is not a significant problem. Similarly, I define exit from the gig economy as a worker not reappearing in the data.

The data spans January 2018 to October 2021. I restrict to worker-months observed from the start of 2019 and onward because the insurer was growing rapidly in 2018 and did not offer a consistent menu of policies. This period of time includes the Covid-19 pandemic, which was a period of continuity and even growth for the food delivery market.[7] Some workers have multiple spells in the gig economy; for these workers I keep their first spell, where a spell is defined as working consecutive months with a break of no longer than three months. The fixed policy can offer additional forms of coverage to a worker's own vehicle for a higher premium, while the variable policy provides only third party coverage. In order to adjust for this, I use reports of willingness to pay (WTP) for additional coverage from the survey (discussed below) to correct workers' premiums.[8] At present, I remove switchers from the analysis although the model is being developed to integrate these individuals.[9] Workers can also opt for an annual policy; these policies are taken up by less than a fifth of individuals who seem to be engaged in permanent, full-time work and, as such, they are qualitatively different from the vast majority of workers who are the focus of this study, so I do not include these people in

[5]Thus, the measure of labor supply is the sum of time spent on food deliveries over the course of a month. Workers often spend 20 to 30% of their time idle between jobs, which I account for in my structural estimation.

[6]It is possible that this aggregation also leaves more time for shocks to workers' valuations to realize. Therefore, in order to provide further insight into workers' dynamic behavior, I augment the main monthly analysis in section 1.3 with weekly data. Further, I allow for shocks to workers' valuations in the structural model in section 1.4.

[7]In appendix 1.C, I show that the reduced form evidence is broadly consistent before and during the pandemic.

[8]Precisely, for those fixed policy workers who purchase additional coverage, I deduct the average WTP conditional on the WTP being greater than the extra premium for additional coverage. Further detail on these filters and adjustments are provided in appendix 1.B and I present reduced form evidence and structural estimates, where I focus strictly on third party only policies, in appendix 1.C and 1.F, respectively.

[9]This is unlikely to significantly affect the results because for every 100 gig workers who exit, only seven switch policies. Further, many policy switches take place at the start of a second stint in the gig economy, in which case the first spell is kept in the analysis sample.

the main analysis.[10]

The firm also has quote data, which contains the menu of prices that workers face when they make their participation and policy decision. This is useful for two reasons. Firstly, it reveals the distribution of fixed policy premiums faced by the population without any selection. I leverage the observed selection into policies based on premiums in the estimation. Secondly, it allows for the construction of individual-level "break-even" points. That is, the number of hours at which both the fixed and variable policy entail the same cost. For illustrative purposes, I often calculate an average break-even point as equal to the average observed fixed premium divided by the average hourly premium, which equals 110 hours (=£103/£0.94ph). Sometimes workers receive more than one quote; if these differ, I use the average of a worker's quotes.

Table 1.2 presents summary statistics for the analysis sample broken down by the type of policy, where the observations have been collapsed to the worker-level to ensure representativeness across workers. In total, I observe 86,024 (=16,575$\times$5.19) worker-months. 64% of workers select the variable policy and these workers tend to work less both in terms of hours and the number of jobs that they complete in a month, but they stay in the gig economy longer than their peers on the fixed policy. Hourly premiums for variable policyholders are considerable at £0.94 per hour.

I complement this administrative data with a survey conducted in collaboration with the firm. The survey was sent out in June 2022 to the firm's active customer base who had subscribed to receiving promotional material.[11] The survey contained questions regarding workers' experiences of the gig economy, especially relative to their expectations, and their policy choice. The survey received over 500 responses in total though not all questions were answered by all respondents.

[10] Naturally, it is also more difficult to evaluate whether policy choices are cost-minimizing since it requires a minimum of a year's worth of data. In appendix 1.F, I compare annual and 30 day policyholders behavior, and adjust the empirical moments in the estimation to inform how their exclusion affects the results.

[11] Given selection based on this subscription decision and voluntarily responding, I use the survey to corroborate my interpretation of the reduced form evidence and to benchmark results from the structural model.

Table 1.2: Worker-Level Summary Statistics

Statistic	Variable	Fixed	Both
Number of workers	10,589	5,986	16,575
Mean number of jobs	114.19	274.95	172.25
Mean duration (months)	5.42	4.76	5.19
Mean monthly hours	43.76	87.51	59.56
SD monthly hours	38.24	59.98	51.72
Mean monthly premium (£)	–	93.20	–
SD monthly premium (£)	–	41.85	–
Mean hourly premium (£)	0.94	–	–
SD hourly premium (£)	0.31	–	–

Notes: This table shows summary statistics at the worker level from the analysis sample. The worker-month-level data is collapsed to the worker-level. Then, for example, the mean hours row displays the mean number of hours worked by workers during an average month, and standard deviations are computed across workers. Mean duration is constructed as the average number of months workers spend in the gig economy. Monthly premiums are constructed as total premiums paid in a 30 day period and hourly premiums are constructed as monthly premiums divided by hours worked in the corresponding 30 day period.

1.3 Patterns in Gig Work Participation

This section presents four empirical facts about gig work participation. Firstly, there is dramatic variation in the number of hours worked per month across individuals. Secondly, hours worked do not predict survival in the gig economy. Thirdly, workers do not always make cost-minimizing policy choices. Fourthly, cost-minimization is correlated with trends in hours worked and survival. These patterns are consistent with workers having misperceptions about the value of gig work and learning, and new survey evidence corroborates this interpretation. Namely, workers report that the realities of gig work frequently deviate from their expectations and that they learn about these differences over time.

These facts motivate four features of the model in section 1.4: (i) workers' have different valuations of gig work, which manifest as a distribution of hours worked in the gig economy; (ii) outside options vary across individuals to justify those working few hours—or, equivalently, those with low valuations—remaining in the gig economy; (iii) workers may misperceive their valuations and this can lead to non-cost-minimizing decisions; and (iv) individuals learn about their true valuations over time, which leads to the observed evolution of hours and survival for (non-)cost-minimizing workers.

1.3.1 Hours Worked

There is enormous dispersion in the number of hours worked in a month by different workers, which suggests that workers are extracting very different value from gig work. Figure 1.1 illustrates this dispersion and how it relates to policy choice. Panel 1.1a presents the empirical distribution of hours worked and the share of policies that make up each hours bin, as reflected by the coloring of the bar. There are two key takeaways from this graph. Firstly, a majority of workers (*circa* two thirds) are on the variable policy; the blue area is greater than the red area. Secondly, most individuals do not work many hours in the gig economy but there is a strong right skew. The modal number of hours worked is approximately 20 hours per month while the mean is 60 hours per month with a standard deviation of 52 hours. Panel 1.1b reveals a third fact: the distribution of hours looks very different conditional on policy choice, as one would expect given selection into policies. For the variable policy, most of the mass is compact around its mean of 44 hours per month. Conversely, hours are more dispersed under the fixed policy with a mean and standard deviation of 88 and 60 hours per month, respectively.

The firm's quotes data contains information on individuals' age and gender, which can be linked to work hours data. Over 90% of workers in the sample are male, which limits statistical power to discern differences in gig labor supply between genders. Still, regression analysis suggests that women on the variable policy tend to work approximately 4 hours less per month than their male counterparts (p-value 2.4%), while there is no statistically significant difference in hours worked between genders for the fixed policy. Interacting gender with age yields imprecise estimates. In general, hours tend to increase with age regardless of policy although this is more pronounced on the fixed policy where, on average, individuals under 30 work 12 hours less per month than those over 40.

1.3.2 Policy Choice

Given information on hours worked and quoted premiums it is possible to assess the quality of policy choice from a cost-minimization perspective. To do so, I construct worker-level break-even points, which describe the number of hours above which the

Figure 1.1: Distribution of Hours Worked

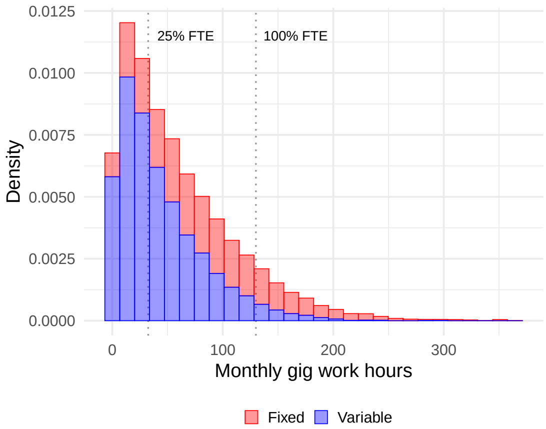

(a) Stacked density by policy choice

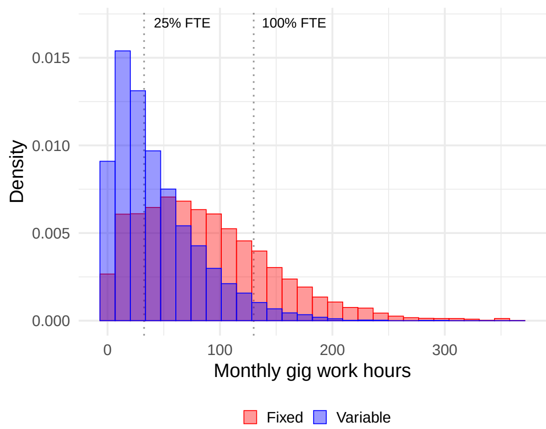

(b) Overlapping density by policy choice

Notes: This figure plots the probability mass of binned hours worked per month for different samples. The sample in panel 1.1a is all workers in the analysis sample and the coloring of the bar is determined by the share of fixed versus variable policy holders. The proportion of the bar that is blue represents the share of workers who are on the variable policy in that hours bin. Two samples are used in panel 1.1b; the distribution of hours worked by fixed and variable policy holders is shown by the red and blue bars, respectively. The dotted grey lines show where different shares of full-time equivalent (FTE) work fall in the distribution of gig work hours. Each observation in a bin is a worker, where repeated worker observations have been averaged over. For this figure, I have removed individuals with less than three monthly observations in order to reduce the impact of noise, which leaves 9,575 workers.

fixed policy is most economical. While the data display patterns firmly consistent with an intention to cost-minimize, many individuals would be better off on the alternative policy and there is some *a priori* evidence of optimism from gig participants.

In section 1.4, this motivates misperceptions over workers' valuations of gig work, which can lead to misperceptions of the number of hours they will work and, in turn, non-cost-minimizing policy choices. Moreover, individuals with higher perceived valuations select into gig work, which can cause the appearance of optimism.

Figure 1.2 provides a convenient lens through which to view policy choice quality (Handel et al., 2020). The graph shows the share of individuals on the fixed policy for different normalized hours bins relative to their break-even point. Normalized hours are constructed as hours of work in the gig economy minus an individual's break-even point. A perfect cost-minimizer would exhibit a step function so that when they work below the break-even point, they are always on the variable policy, and when they are above the break-even point, they are always on the fixed policy. This is illustrated by the dashed green line. Of course, at the break-even point a perfect cost-minimizer would be indifferent between policies and any fixed policy share is compatible with optimization. If one were to take this fictional cost-minimizer and introduce imperfect foresight to their predictions of hours worked, then this would smooth the step function and lead to a monotonically increasing line that crosses the break-even point at 50%.

The data reveals a pattern similar to this, as shown by the blue line in figure 1.2. Workers far from the break-even point (*e.g.*, those on -150 and 150 normalized hours) all but minimize their costs and, moving between these extremes, workers have an increasing tendency to opt for the fixed policy. Therefore, the data is strongly supportive of workers minimizing costs in their policy choice. Yet there are still a significant portion of workers who make non-cost-minimizing choices. This is illustrated by the red shaded regions in figure 1.2, which highlight deviations from the perfect step function. Further, the blue line crosses the break-even point above the 50% level, which could be indicative of optimism from gig workers.

Categorizing workers. For the remainder of this section, I will categorize workers based on a combination of their policy choice and whether their choice was cost-minimizing. Practically, I consider a worker to be cost-minimizing if their policy choice minimized their costs for the majority of months during their tenure in the gig economy.[12] The categories are summarized in table 1.3 alongside their unconditional share of the population. The colors in the matrix correspond to how these groups are depicted in the figures below. Workers who make cost-minimizing choices are grouped together and referred to as "minimizers", and fixed and variable policy holders who make non-cost-minimizing decisions will be called "optimistic" and "pessimistic", respectively.

Broadly, deviations from cost-minimization could be driven by two factors: firstly, *ex post* shocks that affect workers' hours after they have entered the gig economy and, secondly, *ex ante* misperceptions about how much they will work. If *ex post* shocks are responsible for non-cost-minimizing behavior, then these categories should not be predictive of subsequent behavior in the gig economy. Conversely, if *ex ante* misperceptions cause non-cost-minimization, then this partitioning of the data should be correlated with subsequent gig work engagement. Of course, minimizers are also subject to

[12]The results below are robust to other ways of classifying choice quality. Appendix 1.C replicates the analysis below with alternative categories, where non-cost-minimization is calculated by whether either an individual minimized their total insurance premiums over the course of their spell in the gig economy, or whether they minimized their costs for a typical month.

Figure 1.2: Fixed Policy Share by Normalized Hours Worked

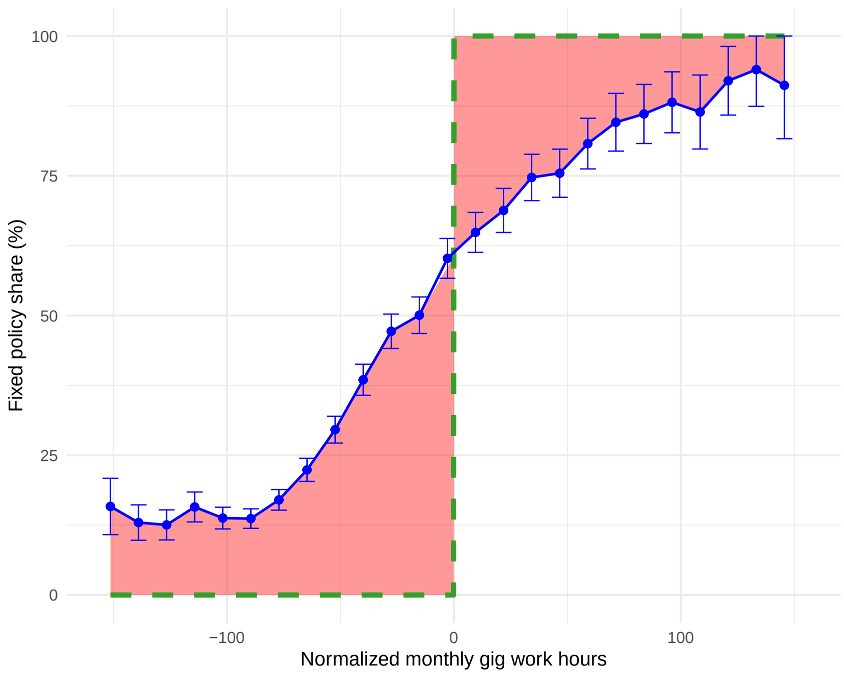

Notes: This figure plots the share of workers who are on the fixed policy by normalized hours bins. Normalized hours are hours minus an individual's break even point, which is constructed from the quote data as a worker's quoted or actual monthly premium divided by the analagous hourly premium. Each observation in a bin is a worker, so the hourly bin that an individual falls into is determined by their average monthly hours. The green dashed line indicates the perfect cost-minimizer's policy choice, which is vertical at the break-even point. Standard errors are constructed by applying the law of large numbers to the average of Bernoulli random variables (*i.e.*, $\sqrt{p \cdot (1-p)/N}$ where p is the share of policies on the fixed policy in a bin and N is the number of observations in that bin).

Table 1.3: Worker Categories

		Policy Choice	
		Fixed	Variable
Cost-Minimizing	Yes	Minimizers 13.1%	Minimizers 59.1%
	No	Optimistic 20.1%	Pessimistic 7.7%

Notes: This table shows the constructions of different worker categories, where the color denote how the categories are shown in subsequent figures. The percentages reflect the proportion of worker-months that fall into each category.

both phenomena, but not such that they have revealed it through their policy choice.

1.3.3 Dynamics of Survival and Hours

In this subsection, I leverage repeated observations of gig workers over time to examine dynamic aspects of worker behavior and how this correlates with their categories.

Figure 1.3 depicts individuals' survival probabilities over time. Panel 1.3a shows that hours do not predict survival; grouping workers by the average number of hours they work does not lead to noticeably different survival trajectories and, moreover, any differences are not monotonic in hours. Intuitively, this implies that workers have different outside options since low-hours workers would not remain in the gig economy if they had to sacrifice the same outside option as a full-time gig worker. Further, policy choice is not predictive of survival. Panel 1.3b shows that fixed and variable policyholders have almost indistinguishable survival paths. Thus, hours and policy choice alone are not informative of tenure in the gig economy.

In contrast, categories are a strong predictor of survival, as shown by panel 1.3c. The optimistic group (*i.e.*, those who select the fixed policy but do not work enough to make it worthwhile) initially drop out of the gig economy faster than the other groups; the red line falls below the other two lines in the first period. Conversely, the pessimistic group (*i.e.*, those who select the variable policy but would have saved money on the fixed policy) exit slowest at first, as evinced by the fact the blue line starts above the others. Minimizers' survival probabilities are somewhere in between those of the optimistic and pessimistic groups, which reflects their less severe exposure to forces that could push them off the cost-minimizing policy.

This evidence suggests that *ex ante* misperceptions play an important role in non-cost-minimizing policy choices because, if the categories only reflected *ex post* shocks, then these dynamic patterns should not be evident. Further, the precise path of survival points towards learning. That is, the categories influence survival most at the beginning of a worker's tenure in the gig economy, which implies that workers enter with misperceptions and learn about these over time such that some individuals exit. After sufficient time has passed, misperceptions have all but gone and the categories do not affect survival further.

To evince these patterns, and to confirm that categories are highly predictive of survival relative to hours, I estimate a Cox proportional hazards model with time-varying coefficients.[13] Table 1.4 displays the estimates. Monthly hours, although statistically significant, have little meaningful impact on survival. The top row of estimates suggest

[13] A parametric model also allows for the inclusion of controls although, in appendix 1.C, I show that the visual patterns remain when splitting the sample by age, time periods, and policy coverage. Appendix 1.C also contains a linear probability regression that reveals the same patterns as the Cox proportional hazards model with time varying coefficients.

Figure 1.3: Survival

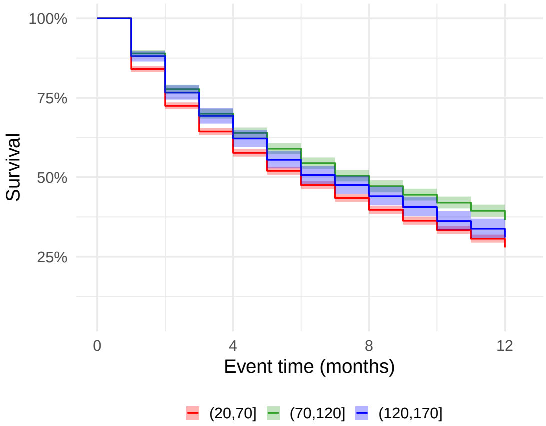

(a) By hours

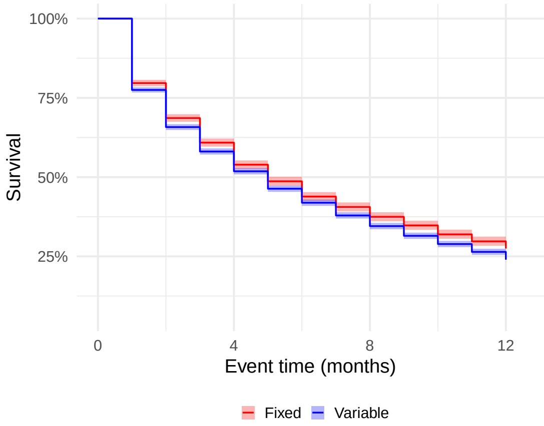

(b) By policy

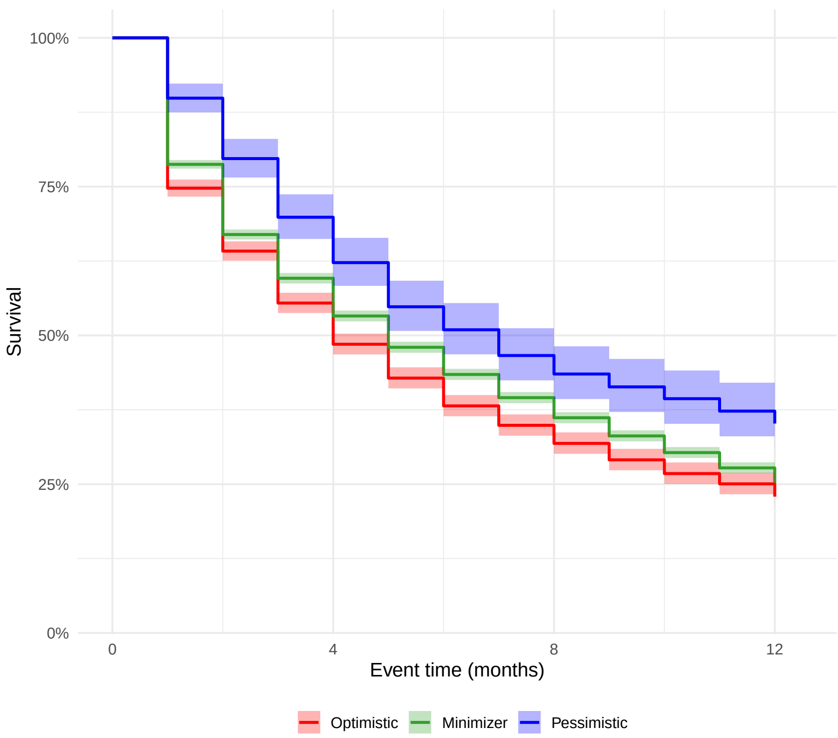

(c) By category

Notes: This figure plots Kaplan-Meier survival curves for different groups in the gig economy. In panel 1.3a, the green, red, and blue lines denote the hours bins [20, 70), [70, 120), and [120, 170), respectively. Panel 1.3b shows fixed and variably policyholders in red and blue, respectively. In panel 1.3c, the green, red, and blue lines denote the minimizers, optimistic, and pessimistic categories, respectively. Event time is tenure month in the gig economy (*i.e.*, $t = 1$ is workers' first month in the gig economy so if an individual does not have a second month in the gig economy, then they exit in the first period).

that increasing a workers average number of hours per month by ten marginally decreases the baseline hazard rate by 1%. Meanwhile, the preferred estimates in column (1) suggest the optimistic and minimizer categories increase the baseline hazard rate by 48% and 22%, respectively, relative to the pessimistic category over the first two months of a worker's spell in the gig economy. Thereafter, their effects wane; the minimizer category's impact is no longer discernible from zero and the optimistic category's effect falls by two thirds—consistent with learning about misperceptions.

Lastly, I show that categories are also correlated with the evolution of hours worked. Figure 1.4 displays how workers' hours evolve over time. The figure shows average hours worked in each week of tenure relative to workers' second week in the gig economy to avoid the fact that workers may not begin working at the start of their first week.[14] Again, the different categories display contrasting behavior. Optimistic workers see their hours initially fall while pessimistic workers see their hours increase at the start of their tenure. Cost-minimizing workers also see their hours fall, though less so than optimistic workers, which is consistent with the survival evidence for minimizers.

This is also supportive of learning and misperceptions. Some workers have severe misperceptions such that, when they become aware of them, they leave the gig economy. For others it is still worthwhile participating but they adjust their intensive margin accordingly.

1.3.4 Survey Evidence

Survey responses from over 300 of the firm's customers provide strong evidence that expectations of gig work often deviate from reality, and that the true value of gig work is learned over time. Figure 1.5 presents responses to four questions contained in the survey. Panel 1.5a shows how individuals found gross earnings (*i.e.*, earnings before costs) relative to their expectations. Earnings expectations appear to be accurate on average, but with significant dispersion such that the majority of workers are left either pleasantly surprised or disappointed in almost equal proportion. Panels 1.5b and 1.5c show that workers report costs and the difficulty of the job, respectively, to be much higher than expected. Approximately three quarters of workers found costs to be more

[14] As with the survival evidence, appendix 1.C shows that the patterns in hours dynamics are robust across different cuts of the data.

Table 1.4: Cox Proportional Hazards Model with Time-Varying Coefficients

	Dependent variable:		
	Tenure in the gig economy (months)		
	All controls	Time controls	No controls
	(1)	(2)	(3)
Mean hours	-0.001^{**}	-0.001^{***}	-0.001^{***}
	(0.0003)	(0.0003)	(0.0003)
Minimizer (<= 2 months)	0.221^{**}	0.181^{*}	0.179^{*}
	(0.100)	(0.095)	(0.095)
Optimistic (<= 2 months)	0.481^{***}	0.378^{***}	0.353^{***}
	(0.104)	(0.098)	(0.098)
Minimizer (> 2 months)	-0.053	-0.067	-0.072
	(0.070)	(0.068)	(0.068)
Optimistic (> 2 months)	0.173^{**}	0.083	0.074
	(0.078)	(0.073)	(0.073)
Low hours	Yes	Yes	Yes
Time controls	Yes	Yes	No
Age	Yes	No	No
Gender	Yes	No	No
Cover	Yes	No	No
Observations	23,969	25,729	25,729
R^2	0.076	0.071	0.066

Notes: *p<0.1; **p<0.05; ***p<0.01. This table shows estimates from a Cox proportional hazards model with time varying coefficients on the categories variable; the effect of this factor variable is allowed to differ between the first two months of a worker's spell and any remaining months. Coefficients reflect percentage changes in the base hazard rate associated with the corresponding variable. The main panel of the table shows estimates for these coefficients and estimates of the coefficient on a worker's average number of hours per month. The table displays three specifications: column (3) includes no controls, column (2) includes only time controls, and column (1) includes time controls and additional covariates from the quotes data. Observations are censored at October 2021. All specifications also include a dummy for low hours because panel 1.3a suggests survival may not be monotonic in hours and to proxy for optimistic misperceptions amongst variable policyholder minimizers. Standard errors are shown in parentheses. The number of observations refers to the number of observations in the stratified data that is used to estimate the time varying coefficients.

Figure 1.4: Hours Worked Over Time by Category

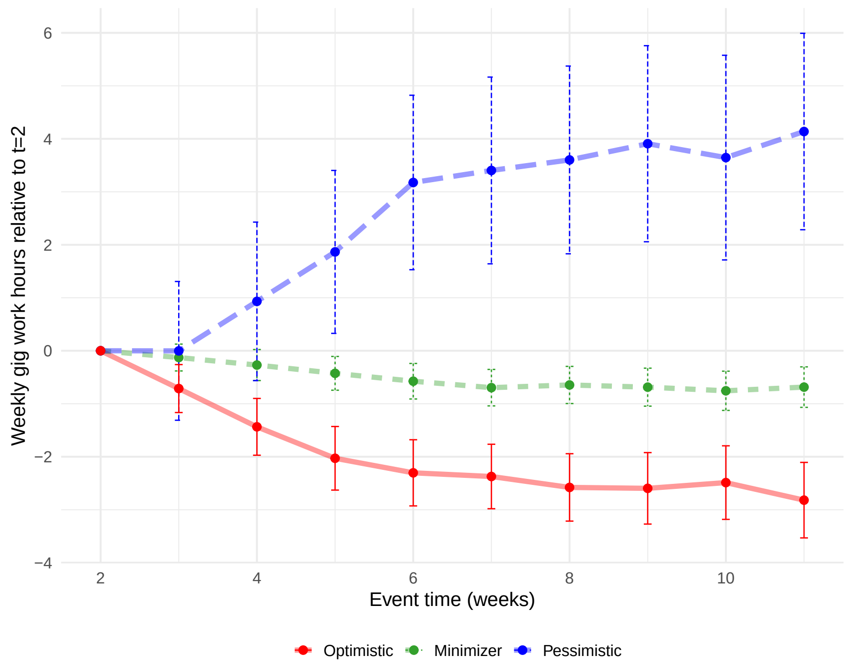

Notes: This figure plots three sets of coefficients from three separate regressions, which are run on a balanced panel of each category of worker. Weekly hours are regressed on fixed effects and event time dummies, where $t = 2$ corresponds to their second month in the gig economy (*i.e.*, event time is tenure in the gig economy), as well as calendar time controls. Intuitively, each coefficient represents the difference in hours at a given point in time from their hours in their first month in the gig economy. SEs are clustered at the worker level with a HC3 weighting scheme.

than expected, while half of them find the work more difficult than anticipated.

Survey responses are generally consistent across policies but, in figure 1.C.6, I show the share of optimists and pessimists' responses relative to minimizers are in line with the hypothesis of misperceptions and learning. Pessimists are more likely to find gig work better than expected relative to optimists and minimizers. Conversely, optimists more frequently report gig work to be worse than expected, while minimizers report their experiences as expected most often. These differences are not statistically significant because of low sample-size, and noise from self-reports of hours and premiums, but are supportive nonetheless.

Inaccurate initial perceptions suggest room for learning. This is confirmed by panel 1.5d, which presents workers' responses to the question of whether they have learned about the "costs versus benefits of this job since [they] started". Around 90% of workers report experiencing learning over the course of their tenure in the gig economy. Inter-

Figure 1.5: Experiences and Learning

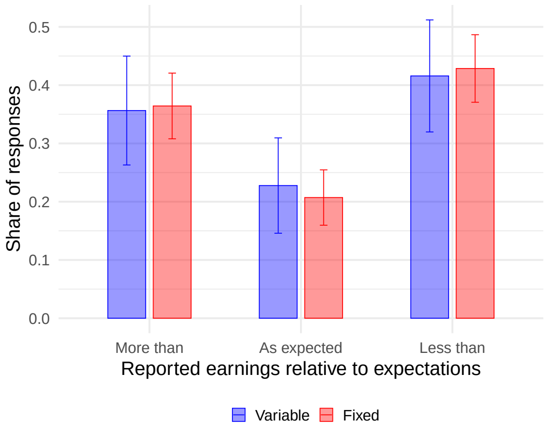

(a) "Are you earning (before costs) more or less than you expected in this job?"

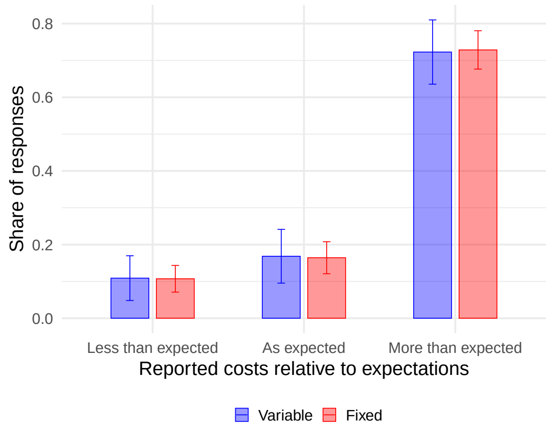

(b) "Are the costs in this job (*e.g.*, fuel, insurance) more or less than you expected?"

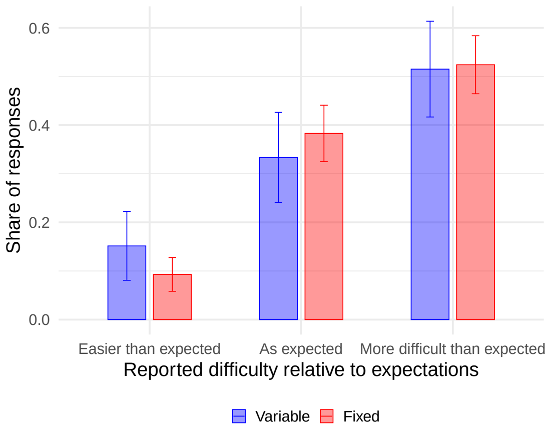

(c) "Is the difficulty of this job more or less than you expected when you started?"

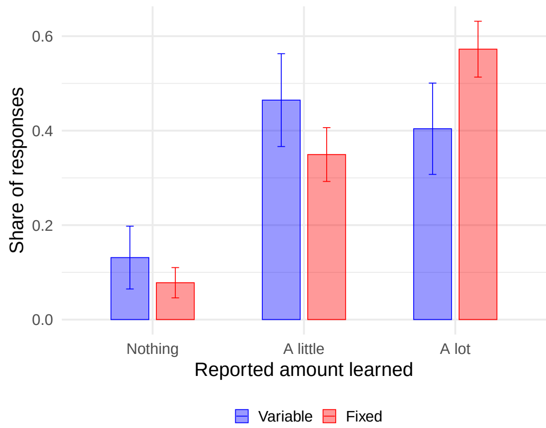

(d) "Have you learned much about the costs vs benefits of this job since you started?"

Notes: This figure plots the share of responses to four questions from the survey. The sample contains 85 variable users and 251 fixed users. Standard errors are calculated by applying the law of large numbers to the average of a random variable that follows the multinomial distribution with one trial (*i.e.*, $\sqrt{p_j \cdot (1-p_j)/N}$ where p_j is the share of responses for a given category j and N is the number of observations).

estingly, some differences between policy holders emerge in this figure. Fixed policy holders are almost 20 percentage points more likely to say that they learned “a lot” relative to variable policy holders.

1.3.5 Discussion

To recap, the distribution of hours worked in the gig economy exhibits are strong right-skew so that, while the majority of individuals work a small number of hours, some workers participate full-time and more. Gig workers suffer from *ex ante* misperceptions about their hours, which cause non-cost-minimizing policy choices. A combination of policy choice and (non-)cost-minimization is predictive of survival and the evolution of hours. In particular, optimists tend to reduce their hours of work in the gig economy and exit faster than other workers.

These pieces of evidence are consistent with a story whereby workers have misperceptions about their valuations of gig work relative to their outside option, but they learn over time about their true valuations. Take an individual who thinks they value gig work more than they in fact do—an optimist—this individual will be more prone to enter and more likely to select the fixed contract. After entering they will learn that their optimism was misplaced and they will reduce their hours, and potentially exit. In an opposite fashion, a worker who initially undervalues gig work—a pessimist—would be more likely to select the variable policy and subsequently increase their hours with a much lower propensity to exit.

Survey evidence lends further support to this narrative. Individuals’ responses confirm that misperceptions of the realities of gig work and learning are common phenomena for gig workers. A model that encapsulates this thesis follows in section 1.4.

1.4 A Theory of Gig Work

In this section, I develop a model of workers’ participation, choice of insurance policy, and hours in the gig economy. The model captures the key features of the economic environment as evinced by the reduced form empirics: Workers have different valuations, which lead to a distribution of hours worked, but they may misperceive these valuations such that many individuals end up on an unnecessarily expensive policy, while some would be better of outside of the gig economy altogether. In addition, workers learn about their true valuation of gig work over time, which manifests in an evolution of hours worked and survival that is correlated with policy choice and (non-)cost-minimization. The model also allows for workers to experience *ex post* shocks to their valuations.

1.4.1 The Model

A worker i is endowed with an individual-specific quadruple $\{\theta_i, \nu_i, \phi_i, P_i\} \in \mathbb{R}_+^4$ that contains their true valuation of gig work θ_i, their outside option ν_i, their initial misperception of their valuation ϕ_i, and their fixed policy premium P_i. If the worker enters the gig economy, upon entering they decide between the fixed and variable policy $\omega \in \Omega = \{\omega_F, \omega_V\}$ and then, each period, they pick how many hours to work in the gig economy $h \in \mathbb{R}_+$. These choices entail a normative flow utility for worker i of

$$u(h, \omega; \theta_i) = \theta_i \cdot \frac{h^{1-\frac{1}{\varepsilon}}}{1-\frac{1}{\varepsilon}} - \big(p(\omega) + \kappa\big) \cdot h - P_i(\omega),^{15} \tag{1.1}$$

where $\varepsilon \in (1, +\infty)$ governs the response of hours to the value of gig work and the variable cost of hours. The variable cost is made up of an exogenous linear cost to working $\kappa \in \mathbb{R}_+$ and the hourly insurance premium $p(\omega)$, which equals $p \in \mathbb{R}_+$ if the worker opts for the variable policy $\omega = \omega_V$ and is zero otherwise. The worker also faces a fixed premium $P_i(\omega)$, which equals P_i if the worker chooses the fixed policy $\omega = \omega_F$ and is zero otherwise.[16] If the worker decides not to enter the gig economy, then they receive their outside option ν_i every period. Utility and surplus are always measured with the normative utility function described by equation (1.1).

When the worker is making their decision about gig work participation, they misperceive the value of gig work. That is, before they enter the gig economy, they perceive their value of gig work to be $\hat{\theta}_{i,0} = \phi_i \cdot \theta_i$. If the worker decides to participate in the gig economy, they will learn about their misperception over time. Concretely, their misper-

[15] The following analysis goes through if this utility function also had an individual-specific intercept, however, in practice this would not be separately identified from the outside option.

[16] Heterogeneity in the fixed premium P_i is motivated by the fact that the firm personalizes prices for the fixed policy but not the variable policy.

ception will erode such that after t periods it is equal to

$$\Phi(t, \phi_i) = \frac{t}{t+\lambda} + \frac{\lambda}{t+\lambda} \cdot \phi_i, \tag{1.2}$$

where $\lambda \in \mathbb{R}_+$ determines the speed of learning (an increase in λ implies slower learning).[17] This functional form is microfounded by a model of individual level Bayesian learning (see appendix 1.D.2). Equation (1.2) implies that $\lim_{t\to\infty} \Phi(t, \phi_i) = 1$ so the worker will all but perceive their true valuation after sufficient time has passed.

At any point in time, the worker will behave in accordance with their perceived value of gig work $\hat{\theta}_{i,t} = \Phi(t, \phi_i) \cdot \theta_i$, which generates a perceived flow utility for worker i at time t given by

$$u_i(h, \omega; \hat{\theta}_{i,t}) = \hat{\theta}_{i,t} \cdot \frac{h^{1-\frac{1}{\varepsilon}}}{1-\frac{1}{\varepsilon}} - \big(p(\omega) + \kappa\big) \cdot h - P_i(\omega).^{18} \tag{1.3}$$

Misperceptions mean that individuals can find themselves on non-cost-minimizing policies, and that the opportunity to work in the gig economy can lead to welfare losses when individuals enter because of an inflated valuation.

The worker may exit the gig economy at any period. This could be due to either endogenous exit, if a worker's perceived value falls below the threshold that makes gig work worthwhile relative to their outside option, or because of an exogenous shock that removes them from the gig economy with probability η. This latter feature captures persistent shocks to valuations or outside options that render the surplus from gig work negative.

Given this series of decisions—participation, policy choice, hours, and exit—the model is solved via backward induction. Exit is either exogenous or equivalent to participation, so I start with hours, then policy choice, and then participation and exit.

Hours worked. Hours are chosen each period by workers conditional on their participation and policy in order to maximize perceived flow utility. Therefore, worker i in

[17] As well as learning about misperceptions, workers may also learn to improve their productivity on the job, which is not explicitly modelled here. Intuitively, this would be an alternative force that raises the true and perceived value of gig work over time. Since systematic dynamics in the model are captured by learning about misperceptions, I expect this mechanism to influence estimated misperceptions and the rate of learning. The direction of the effect is ambiguous *ex ante*.

[18] Note that the subscript i for the utility function $u_i(\bullet)$ captures the role of individual specific fixed policy premiums P_i, which are suppressed as an argument of the function (this is true also for the value functions $V_i(\bullet)$ and $\tilde{V}_i(\bullet)$ that are defined later).

period t will pick hours to maximize equation (1.3)

$$h^*_{i,t}(\omega) = \left(\frac{\hat{\theta}_{i,t}}{p(\omega)+\kappa}\right)^{\varepsilon}.^{19} \tag{1.4}$$

Policy choice. Workers make their policy decision before they have entered into the gig economy, and so suffer from their initial misperception. They believe that their flow utility during their tenure in the gig economy will remain constant (*i.e.*, they are naïve about their learning) so they pick whichever contract yields a higher flow utility. Denote the value function

$$V_i\big(\omega;\hat{\theta}_{i,0}\big) = u\big(h^*_i(\omega),\omega;\hat{\theta}_{i,0}\big). \tag{1.5}$$

Therefore, individuals pick the policy which maximizes their perceived flow utility

$$\omega^*_i = \arg\max_{\omega}\big\{V_i\big(\omega_F;\hat{\theta}_{i,0}\big), V_i\big(\omega_V;\hat{\theta}_{i,0}\big)\big\}. \tag{1.6}$$

Participation and exit. Since workers believe they will remain in the gig economy until they exogenously exit, at which point they receive their outside option, they will decide to enter the gig economy if and only if

$$V_i\big(\omega^*_i;\hat{\theta}_{i,0}\big) > \nu_i. \tag{1.7}$$

Similarly, worker i will exit at time t if their perceptions evolve such that

$$V_i\big(\omega^*_i;\hat{\theta}_{i,t}\big) \leq \nu_i. \tag{1.8}$$

Panel 1.6a in figure 1.6 shows the mechanics of the model when there are no misperceptions and outside options are constant across individuals. If a worker's valuation exceeds an initial threshold, which is shown by the dashed red line, then they will decide to enter the gig economy and select the variable policy. Under this policy, their hours increase with their valuation, however, if their valuation exceeds a further threshold shown by the dashed green line, then they will opt for the fixed policy. Workers' hourly wage rate jumps up when they cross this threshold because they no longer face the

[19] In the empirical implementation of the model, a time period is set to one month and hours are truncated from above at 480 (=16×30) to reflect time constraints, and workers fully account for this, although it does not bite much in practice.

Figure 1.6: Model Illustration

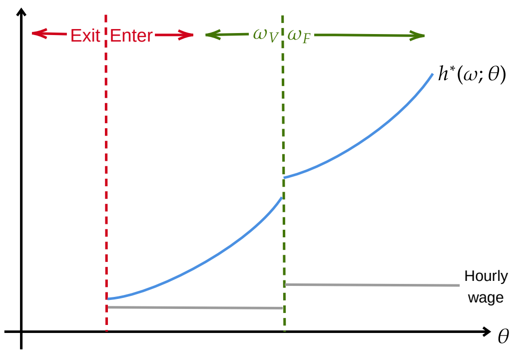

(a) The Worker's Problem

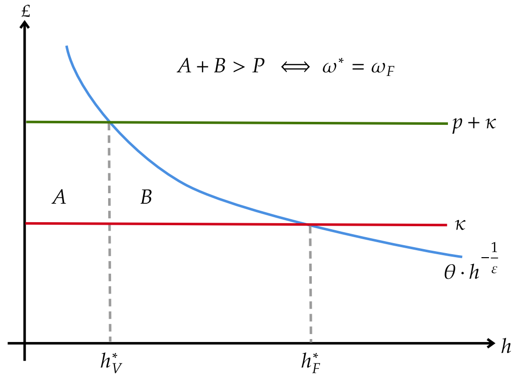

(b) Money-metric Identification

Notes: These figures illustrate different aspects of the model. Panel 1.6a shows the workers problem for a fixed outside option ν with no misperceptions $\phi = 1$ and a range of valuations θ. The dashed red line denotes the threshold valuation required for a worker to enter the gig economy, and the dashed green line is the valuation threshold at which a worker would prefer to be on the fixed policy. The blue curve denotes the hours worked per month, while the grey line denotes the hourly wage rate. Panel 1.6b illustrates how the model identifies money-metric valuations. The blue curve denotes the marginal benefit of an additional hour in the gig economy, and the red and green lines show the marginal cost under the variable and fixed policy, respectively.

hourly premium.

1.4.2 Introducing Shocks to Valuations

To help bring the model to the data, I introduce transitory shocks to workers' valuations. This is important because, although the reduced form evidence supports the hypothesis of misperceptions and learning, *ex post* shocks still affect the cost-minimizing nature of workers' policy choices. Consequently, any patterns in the data are the result of a confluence of these factors that the model should reflect. It is also evident that hours vary within workers across months, which further motivates shocks to the valuations that drive hours.

I consider that the workers' valuations may also be subject to independently and identically distributed shocks $\rho_{i,t} \in \mathbb{R}_+$ each period. The distribution of shocks is known to workers and they leave the worker's valuation unchanged in expectation $\mathbb{E}[\rho_{i,t}] = 1$. Thus, on any given period, the worker's true valuation is $\theta^{\rho}_{i,t} = \rho_{i,t} \cdot \theta_i$, although it will be perceived to be $\hat{\theta}^{\rho}_{i,t} = \rho_{i,t} \cdot \Phi(t, \phi_i) \cdot \theta_i$. That is, the worker's normative flow utility is given by

$$u_i(h,\omega;\theta_{i,t}) = \theta_{i,t}^{\rho} \cdot \frac{h^{1-\frac{1}{\varepsilon}}}{1-\frac{1}{\varepsilon}} - \big(p(\omega)+\kappa\big)\cdot h - P_i(\omega), \tag{1.9}$$

although it is perceived to be

$$u_i(h,\omega;\hat{\theta}_{\rho,i,t}) = \hat{\theta}_{i,t}^{\rho} \cdot \frac{h^{1-\frac{1}{\varepsilon}}}{1-\frac{1}{\varepsilon}} - \big(p(\omega)+\kappa\big)\cdot h - P_i(\omega). \tag{1.10}$$

If a shock is sufficiently low, then the worker may not want to work in the gig economy. In this case, they can access their outside option at a discount $\nu_i \cdot (1-\psi)$, where $\psi \in (0,1)$, and work zero hours in the gig economy. I refer to this as temporary exit. Temporary exit reflects two features of reality. Firstly, if an individual participates regularly in the gig economy, they have less time to invest in and raise the value of their outside option. Secondly, it captures any refundable fixed costs to enter into the gig economy. Moreover, temporary exit reflects intermittent short breaks in gig work that are observed in the data.

Conditional on participation and policy choice, hours are chosen according to equation (1.4) with $\hat{\theta}_{i,t}^{\rho}$ replacing $\hat{\theta}_{i,t}$, if the worker does not temporarily exit. Workers will temporarily exit if their shock $\rho_{i,t}$ is sufficiently low. In this case, they will receive their outside option at a discount $\nu_i \cdot (1-\psi)$ so, formally, they will make a temporary exit and work zero hours if and only if

$$\begin{aligned} u(h_{i,t}^*(\omega),\omega;\hat{\theta}_{\rho,i,t}) &= \frac{1}{\varepsilon-1}\cdot\frac{\big(\hat{\theta}_{i,t}^{\rho}\big)^{\varepsilon}}{\big(p(\omega)+\kappa\big)^{\varepsilon-1}} - P_i(\omega) \le \nu_i\cdot(1-\psi) \\ \iff \quad \rho_{i,t} &\le \left(\big(\nu_i\cdot(1-\psi)+P_i(\omega)\big)\Big/\frac{1}{\varepsilon-1}\cdot\frac{\big(\hat{\theta}_{i,t}\big)^{\varepsilon}}{\big(p(\omega)+\kappa\big)^{\varepsilon-1}}\right)^{\frac{1}{\varepsilon}} = \Gamma(\nu_i,\hat{\theta}_{i,t},P_i). \end{aligned} \tag{1.11}$$

In summary, hours worked are determined by

$$h_{i,t}^*(\omega) = \begin{cases} \left(\frac{\hat{\theta}_{i,t}^{\rho}}{p(\omega)+\kappa}\right)^{\varepsilon} & \text{if } \rho_{i,t} > \Gamma(\nu_i,\hat{\theta}_{i,t},P_i), \\ 0 & \text{otherwise.} \end{cases} \tag{1.12}$$

Workers' participation decisions must incorporate the possibility of temporary exit and engage more fully with the dynamic nature of the problem they face. They discount the future with discount factor $\beta \in (0,1)$ such that they perceive their discounted sum of utility under policy ω for worker i at time t to be

$$\begin{aligned}
\tilde{V}_i(\omega;\hat{\theta}_{i,t}) =& \mathbb{E}\left[\sum_{t=0}^{\infty}(\eta\cdot\beta)^t\cdot\max\left\{\frac{1}{\varepsilon-1}\cdot\frac{\left(\hat{\theta}_{i,t}^{\rho}\right)^{\varepsilon}}{\left(p(\omega)+\kappa\right)^{\varepsilon-1}}-P(\omega),\nu_i\cdot(1-\psi)\right\}\right]+\sum_{t=0}^{\infty}(1-\eta^t)\cdot\beta^t\cdot\nu_i \\
=&\frac{1}{1-\eta\cdot\beta}\cdot\Bigg(\mathbb{P}\left(\rho_{i,t}\leq\Gamma(\nu_i,\hat{\theta}_{i,t},P_i)\right)\cdot\nu_i\cdot(1-\psi)\dots \qquad (1.13)\\
&\dots+\mathbb{P}\left(\rho_{i,t}>\Gamma(\nu_i,\hat{\theta}_{i,t},P_i)\right)\cdot\left(\frac{1}{\varepsilon-1}\cdot\frac{\left(\hat{\theta}_{i,t}\right)^{\varepsilon}}{\left(p(\omega)+\kappa\right)^{\varepsilon-1}}\cdot\mathbb{E}\left[\rho_{i,t}^{\varepsilon}|\rho_{i,t}>\Gamma(\nu_i,\hat{\theta}_{i,t},P_i)\right]-P_i(\omega)\right)\Bigg)\dots \\
&\dots+\frac{\nu_i\cdot\beta\cdot(1-\eta)}{(1-\beta)\cdot(1-\eta\cdot\beta)},
\end{aligned}$$

where the expression makes use of the fact that exogenous exit is an absorbing state. Equation (1.13) breaks down the value function into three lines. The first line contains the utility derived from temporary exit multiplied by the probability of temporary exit on any given period. This line is also multiplied by $1/(1-\eta\cdot\beta)$ to reflect the fact that these pay-offs form a geometric sequence with common ratio $\eta\cdot\beta$. The second line contains the expected flow utility conditional on working in the gig economy scaled by the probability of receiving a sufficiently large shock to work in the gig economy. The last line reflects the expected utility derived from the full value of a worker's outside option, which is received upon exit.

Therefore, worker i selects policy

$$\omega_i^* = \arg\max_{\omega}\left\{\tilde{V}_i(\omega_F;\hat{\theta}_{i,0}),\tilde{V}_i(\omega_V;\hat{\theta}_{i,0})\right\}, \qquad (1.14)$$

and enters the gig economy if and only if

$$\tilde{V}_i(\omega_i^*;\hat{\theta}_{i,0}) > \frac{\nu_i}{1-\beta}. \qquad (1.15)$$

Exit at a given time t occurs endogenously if and only if

$$\tilde{V}_i\left(\omega_i^*;\hat{\theta}_{i,t}\right) \leq \frac{\nu_i}{1-\beta}. \qquad (1.16)$$

1.4.3 Discussion

The model sets out a sequence of decisions which are designed to capture the way workers engage with the gig economy. Here I discuss some aspects of the model.

Heterogeneity. The model allows for heterogeneity both across and within workers. Individuals differ in three key dimensions: gig work valuations, outside options, and misperceptions. A desire to capture the reduced form evidence in the most simple way motivates introducing heterogeneity in these specific areas. Firstly, there is significant variation in the average number of hours worked between workers, which motivates the flexibility of individual specific valuations of gig work. Secondly, workers make different policy choices; some make cost-minimizing decisions and, for those who do not, there is significant dispersion in how far their hours are from their break-even point. Differences in workers perceptions of their valuations are necessary to capture this feature of the data. Finally, it is intuitive that workers face different outside options and that this is critical to their participation in the gig economy. For example, one worker's participation may be precarious because they compare gig work with an alternative job, which is only marginally worse. For this individual, slight changes in their valuation could push them out of gig work. Conversely, an individual who has recently been made unemployed and has few employment prospects is much more attached to the gig economy.

The model also incorporates heterogeneity within a worker's spell in the gig economy. Workers are subject to transitory shocks that affect their valuation of gig work, as well as permanent shocks that remove them from gig work. These permanent shocks capture persistent negative shocks to the surplus that workers derive from gig work. Learning also generates dynamics in within-worker labor supply to the gig economy.

Gig work valuations. Workers' valuations of gig work encompass a broad range of factors that determine how gig work affects their utility. Perhaps most intuitively, one can think of the wage as being an important determinant of valuations, but even this is mediated by other factors such as workers' marginal utility of income. Amenities also likely play an important role, for example, flexible hours, absence of a boss, and varying demands of the job. All these characteristics of gig work, and more, make up valuations. The model allows workers to value these aspects differently but it does impose that, conditional on working the same number of hours, workers valuations are equivalent. In other words, workers who work the same number of hours value gig work equally, although the value may be derived from different factors. This does not mean workers with the same number of hours receive the same surplus since they will have different misperceptions and outside options.

Outside options. Outside options constitute workers' welfare in a world where they do not work in the gig economy. Broadly speaking, workers have two margins of adjustment. Firstly, they can replace the hours worked in the gig economy with another activity (*e.g.*, leisure or some alternative work). Secondly, they can adjust the bundle of activities that they undertake in a day. That is, workers can reorganize all the activities in their day and not just the time they would spend in the gig economy. For example, a worker who complements a full-time job with gig work may instead take up two part-time jobs, if gig work is not available. The model matches these features of reality in two ways: it has a linear cost to hours, which reflects the opportunity cost of time, and a fixed outside option that captures the re-bundling effect.[20]

Misperceptions. Workers can have disappointing experiences of gig work if their misperceptions cause them to be overly optimistic about the value of gig work. This can manifest visibly as workers being on the non-cost-minimizing policy, and as workers promptly reducing their hours and leaving the gig economy. In the model, workers are assumed to be ignorant of the possibility that they misperceive the value gig work. In other words, they act as if they have perfect knowledge of their valuation. This assumption allows the model to abstract from several potential influences. In particular, workers neither shade their valuations (Capen et al., 1971; Smith and Winkler, 2006; Thaler, 1988), nor stay in the gig economy to learn about their valuations, nor pick their hours to maximize an uncertain return.

Another important assumption is that workers misperceive the value of gig work rather than another parameter in the model. In particular, misperceptions could be about outside options, where optimism about gig work maps to pessimism about outside options in a similar way to that found by Jäger et al. (2022). However, the dynamics of hours and survival in the data, as well as the survey evidence, motivates the choice to embed misperceptions in the value of gig work. In particular, it is hard to conceive how misperceptions in outside options could generate dynamics within the gig economy.

[20]Individual-level outside options ν_i may also pick up some opportunity cost of time because this is assumed to be homogeneous across workers through κ, but in reality it may vary.

The gig work surplus. The gig work surplus differs across and within workers because of the model's rich heterogeneity. Differences in surpluses across workers are driven primarily by variation in gig work valuations and outside options, but they can also differ because of suboptimal labor supply decisions driven by misperceptions. Within workers, shocks to valuations cause the surplus to vary over time and to eventually go to zero because of exogenous exit shocks.

The model's gig work surplus measure does not capture the possibility that outside options could change with, for example, a mass exodus from the gig economy. Plausibly, in the short-run, outside options may fall in such a situation as the supply of labor to other markets increases while demand remains unchanged leading to a fall in wages. Thus, the gig work surplus identified here serves as a lower bound. In the long run, the prospect for outside options is unknown in such a scenario.

1.5 Estimation

This section describes how I estimate the structural parameters in the model and provides a guide to how variation in the data helps identification. Further, I present the estimates and model fit before section 1.6 analyses their implications for welfare and counterfactual scenarios in the gig economy.

1.5.1 Simulated Method of Moments Estimator

The variable premium p is set equal to its average of £0.94 in the data and I feed the model the empirical distribution of quoted fixed premiums P_i, which can be seen in figure 1.B.1. These are the premiums offered by the firm—not necessarily taken by customers—and so do not suffer from selection. Moreover, they are all but uncorrelated with hours worked, which implies they are unrelated to worker valuations θ_i from the perspective of the model. Through conversations with the firm it is also apparent that customers' policies are not priced based on any proxy for misperceptions ϕ_i. Therefore, I assume individuals' fixed premiums are independent of their other characteristics.

To reduce the burden of estimation, I also fix the exogenous exit rate η equal to the exit rate of pessimistic workers. In the model, these workers have no reason to leave aside from exogenous shocks. Further, the elasticity parameter ε is mechanically adjusted in order to ensure an intensive margin labor supply elasticity compatible with empirical evidence. Lastly, β is set to a standard value.

The remaining parameters are estimated using SMM. This requires an assumption about the distribution of heterogeneity in the population. I assume that individuals' valuations θ_i and outside options ν_i follow a joint log-normal distribution. This has three main advantages. Firstly, it helps to capture the skewed distribution of hours that is evident in the data. Secondly, it allows for an easily specified but rich pattern of correlations between these individual characteristics. Thirdly, the log-normal distribution permits convenient closed form expressions that are helpful with computation. For similar reasons, I specify *ex ante* misperceptions to follow an iid log-normal distribution with a mean equal to one.[21] Thus, in summary, individual heterogeneity is distributed according to

$$\begin{pmatrix} \theta_i \\ \phi_i \\ \nu_i \end{pmatrix} \sim \log\mathcal{N}\left(\begin{bmatrix} \mu_\theta \\ -\sigma_\phi^2/2 \\ \mu_\nu \end{bmatrix}, \underbrace{\begin{bmatrix} \sigma_\theta^2 & 0 & \sigma_{\theta,\nu} \\ 0 & \sigma_\phi^2 & 0 \\ \sigma_{\nu,\theta} & 0 & \sigma_\nu^2 \end{bmatrix}}_{=\Sigma} \right).$$

In practice, I estimate elements of the Cholesky decomposition L of the covariance matrix $\Sigma = LL^T$ in order to ensure the latter is positive semidefinite, where I fix two elements of the lower triangular matrix L to reflect the constraint imposed on the covariance matrix. Further, outside options are truncated at £10,000 since the model's parameters are identified from observed participants, thus, the data cannot speak to individuals who are far from entering the gig economy. The idiosyncratic shocks to worker valuations are assumed to be log-normal iid distributed $\rho_{i,t} \sim \log\mathcal{N}(\mu_\rho, \sigma_\rho^2)$ with $\mu_\rho = -\sigma_\rho^2/2$ so that $\mathbb{E}[\rho_{i,t}] = 1$.

This leaves ten parameters to estimate in the model; six parameters describing the joint log-normal distribution, the linear cost to hours κ, the rate at which misperceptions correct λ, the variance of shocks that affect workers valuations σ_ρ^2, and the sunk

[21] Previous iterations of the model, which allow for misperceptions to be systematically different from one and for a correlation with valuations, do not improve the model fit.

portion of gig workers' outside options ψ. Let $\boldsymbol{\zeta}$ denote this vector of parameters

$$\boldsymbol{\zeta} = \{\mu_\theta, \sigma_\theta^2, \sigma_\phi^2, \mu_\nu, \sigma_\nu^2, \sigma_{\theta,\nu}, \sigma_\rho^2, \psi, \kappa, \lambda\}.$$

I construct the difference between the jth model moment $\hat{m}_j(\bullet)$ and the jth data moment $m_j(\bullet)$ to be $e_j(\bullet)$ so that

$$e_j(\tilde{X}, X|\boldsymbol{\zeta}) = \hat{m}_j(\tilde{X}|\boldsymbol{\zeta}) - m_j(X),$$

where X denotes the observed data and $\tilde{X}$ denotes the simulated data. The estimated parameters are those that minimize the weighted sum of errors

$$\hat{\boldsymbol{\zeta}} = \arg\min_{\boldsymbol{\zeta}} \; \boldsymbol{e}(\tilde{X}, X|\boldsymbol{\zeta})^T W \boldsymbol{e}(\tilde{X}, X|\boldsymbol{\zeta}),$$

where $\boldsymbol{e}(\bullet)$ is the stacked deviations of the moments $(J \times 1)$ and W is a weight matrix $(J \times J)$.

Standard errors for the parameters are computed according to

$$\hat{\mathbb{V}}(\hat{\boldsymbol{\zeta}}) = \frac{N+1}{N} \cdot \left(\hat{\mathcal{J}}^T W \hat{\mathcal{J}}\right)^{-1} \left(\hat{\mathcal{J}}^T \hat{\Sigma} W \hat{\Sigma}^T \hat{\mathcal{J}}\right) \left(\hat{\mathcal{J}}^T W \hat{\mathcal{J}}\right)^{-1},$$

where N is the number of simulations, $\mathcal{J}$ is the Jacobian matrix of the moment conditions $(J \times K)$, and Σ is the variance-covariance matrix of the moment conditions $(J \times J)$ (Hansen, 1982). The hat notation $\hat{\bullet}$ reflects the fact that these components are functions of the data. In particular, the variance-covariance matrix of the moment conditions is estimated by block bootstrapping at the worker level and recomputing the moments. For moments that do not come from the main dataset, I construct their variances and assume zero correlation with the other moments.

The ten parameters of the model are identified with 23 empirical moments:

1. The labor market share of this section of the gig economy.
2. The proportion of workers on the variable policy.
3. The mean number of hours worked in a month by workers on the variable policy.
4. The standard deviation of the number of hours worked in a month by workers on the variable policy.

5. The mean number of hours worked in a month by workers on the fixed policy.
6. The standard deviation of the number of hours worked in a month by workers on the fixed policy.
7. The share of variable policy workers who are not on the cost-minimizing policy.
8. The share of fixed policy workers who are not on the cost-minimizing policy.
9. The mean distance between the break-even point and the number of hours worked in a month for non-cost-minimizing variably policy workers.
10. The standard deviation of the distance between the break-even point and the number of hours worked in a month for non-cost-minimizing variably policy workers.
11. The mean distance between the break-even point and the number of hours worked in a month for non-cost-minimizing fixed policy workers.
12. The standard deviation of the distance between the break-even point and the number of hours worked in a month for non-cost-minimizing variably policy workers.
13. The initial hazard rate of cost-minimizing variable policy holders.
14. The initial hazard rate of cost-minimizing fixed policy holders.
15. The initial hazard rate of non-cost-minimizing fixed policy holders.
16. The initial decline in hours per month of non-cost-minimizing variable policy holders.
17. The initial decline in hours per month of non-cost-minimizing fixed policy holders.
18. The mean quoted fixed premium for fixed policy holders.
19. The mean quoted fixed premium for variable policy holders.
20. The standard deviation of quoted fixed premiums for fixed policy holders.
21. The standard deviation of quoted fixed premiums for variable policy holders.
22. The frequency of zero hour months.
23. The average within worker standard deviation in hours.

The labor market share is derived from two sources separate to the firm's administrative data. Bertolini et al. (2021) report that 17.1% of the UK workforce work for digital platforms on at least a monthly frequency from a survey of 2,201 workers, and Cornick et al. (2018) state that 21% of the 95 gig workers in their sample work in food delivery. Hence, a labor market share of 3.59% (=17.1%×21%).[22] I reconcile the labor market share with the model through a stationarity assumption (*i.e.*, that the employment share of the gig economy is in steady state).

Lastly, I lean on the labor supply elasticity literature to pin down ε. Given the short term nature of gig work for most workers (the standard duration is less than six months), a Frisch elasticity seems the most appropriate. Moreover, given these are self-employed workers with total control over their hours one would expect this group of individuals to exhibit a high Frisch elasticity. I take the Frisch elasticity estimate of 0.80 (SE 0.10) from Fisher (2022), which is estimated on a sample of self-employed taxi drivers who are subject to exogenous variation in their wage rates due to London tube strikes. I combine this with online earnings data for delivery riders and scale this to account for idle time.[23]

In practice, I find the minimum of the objective function using a multi-start simplex search method, the weight matrix is set to normalize the error function to percentage deviations, and I simulate six million workers.

1.5.2 Sources of Identification

While the parameters are identified from a combination of structural assumptions and moments of the data, it is helpful to consider the particularly close linkages between some parameters and moments. Workers' hours are determined by an unobserved distribution of valuations of gig work $\{\mu_\theta, \sigma^2_\theta\}$. Valuations can be made up of many factors (*e.g.*, productivity, the wage rate, an individual's marginal utility of income, and their preference for the type of work), and the higher an individual's valuation the more they want to work in the gig economy. The model implicitly assumes that, aside from misperceptions, participants who work the same hours have the same valuations—and,

[22] I assume independence between these two statistics and use that fact that for two independent random variables X and Y with finite first and second moments $\mathbb{V}(X \cdot Y) = \mathbb{V}(X) \cdot \mathbb{V}(Y) + \mathbb{V}(X) \cdot \mathbb{E}(Y)^2 + \mathbb{E}(X)^2 \cdot \mathbb{V}(Y)$ in order to construct the variance-covariance matrix of the empirical moments.

[23] For example, see Glassdoor or Indeed.

thus, the distribution of valuations is implied by the hours distribution.

The money-metric nature of this valuation comes primarily from the policy choice, which can be framed as the opportunity to buy a higher wage rate at some upfront cost. Panel 1.6b in figure 1.6 shows the logic behind this. The additional benefit from switching to the fixed policy is the hourly earnings that are saved rather than paid to the firm (*i.e.*, area A), plus the sum of marginal benefits net of marginal costs for the additional hours that are worked (*i.e.*, area B). Therefore, the linear cost to work κ must imply a solution to this problem that matches the data. To the extent that policy choice pins down a money-metric, variation in fixed premiums P_i also helps identification.

Naturally, the prevalence of non-cost-minimizing choices and a measure of how far these choices were from justifying their policy choice implies the distribution of misperceptions σ^2_ϕ. For example, an individual on the fixed policy who would save money on the variable policy must have overestimated their valuation (*i.e.*, $\phi_i > 1$) but this alone does not provide information about the extent of the individual's over optimism. However, if the individual was only one hour [100 hours] away from the break-even point, then their misperception must have been small [large]. Selection into participation plays an influential role in determining the distribution of misperceptions amongst the pool of gig workers.

The speed of learning λ is primarily identified from the change in workers' hours upon entering the gig economy. A greater initial adjustment in hours, for example, indicates a quick rate of learning.

Outside options $\{\mu_\nu, \sigma^2_\nu\}$ are most connected to the employment share of the gig economy and survival rates. Intuitively, the employment share can identify a constant outside option; with a fixed distribution of valuations, a homogeneous outside option can be adjusted to ensure the correct employment share. The estimation leverages endogenous exit in the model to identify heterogeneity in outside options. With knowledge of the learning process and valuations, it is possible to infer outside options as the perceived valuations at which optimistic workers decide to leave the gig economy.

The covariance between valuations and outside options $\sigma_{\theta,\nu}$ comes from conditioning the moments on a combination of policy choice and whether the policy choice was cost-minimizing. For example, given knowledge of misperceptions and the speed of learning, differential exit rates amongst variable minimizers and optimistic workers is

Table 1.5: Parameter Estimates

$\zeta_{1:5}$	$\hat{\zeta}_{1:5}$	$\zeta_{6:10}$	$\hat{\zeta}_{6:10}$
μ_θ	3.70	μ_ν	12.66
	(0.07)		(0.03)
σ^2_θ	0.30	σ^2_ν	4.54
	(0.03)		(0.03)
$\sigma_{\theta,\nu}$	-1.08	σ^2_ϕ	0.14
	(0.05)		(0.01)
λ	1.00	κ	39.41
	(0.13)		(1.13)
σ^2_ρ	0.03	ψ	0.27
	(0.00)		(0.00)
μ_P	101.82	σ_P	24.08
p	0.94	β	$0.95^{1/12}$
η	0.07	ε	$\Delta \widehat{\log}(h) \cdot \log(1 + p/\hat{\kappa})$

Notes: The top panel of this table presents estimates of the structural parameters from the model. Standard errors are contained in the parentheses, which are estimated as described in subsection 1.5.1. The second panel shows the variable policy premium and the mean of the quoted fixed premium distribution used in the estimation, which correspond to the empirical averages of these parameters, as well as the set discount factor.

indicative of outside options. Again, selection into participation plays a key role in mediating the covariance amongst gig workers compared to the population as a whole.

The variance of shocks σ^2_ρ is identified from the mean within worker standard deviation of hours worked, and the fraction of the outside option that is sunk due to regularly participating in the gig economy ψ is inferred from the prevalence of interruptions to workers' spells in the gig economy (*i.e.*, the fraction of months where workers work zero hours but reappear in the data subsequently).

1.5.3 Parameters and Model Fit

Table 1.5 presents the structural estimates and their associated standard errors in the top panel. The lower panel of the table shows the fixed model parameters. Namely, the variable premium p, the mean and standard deviation of the fixed premium distribution $\{\mu_P, \sigma_P\}$, the discount factor β, exogenous exit rate η, and the elasticity parameter ε. The estimates map to a joint distribution of characteristics in the simulated population, which are shown in table 1.6.

Across the whole population, the mean and standard deviation of valuations θ_i is equal to 47 and 28, respectively. The distribution of valuations exhibits a right skew so

Table 1.6: Simulated Population Characteristics

Statistic	Population	Participants
Mean valuation θ	47.24	211.17
SD valuation θ	28.07	48.81
Mean misperception ϕ	1.00	1.04
SD misperception ϕ	0.39	0.11
Mean outside option ν	9,744	674
SD outside option ν	1,258	491
Correlation $\rho_{\theta,\nu}$	-0.62	-0.32
Correlation $\rho_{\theta,\phi}$	0.00	-0.27
Correlation $\rho_{\nu,\phi}$	0.00	0.53

Notes: This table presents statistics that describe the simulated population. The first column shows these statistics for the entire population, while the second column conditions on participating in the gig economy.

that the median valuation equals 41. Workers with higher valuations are more likely to participate in the gig economy, all else equal, so participants exhibit higher valuations equal to 211 on average. This valuation would translate to a variable policy flow utility (*i.e.*, $u(h_i^*(\omega_V), \omega_V; \theta_i)$) of £1,398. I discuss the magnitude of these flows net of outside options in the next section.

The mean outside option for participants equals £674 with a standard deviation and median of £491 and £608, respectively. Since the average participant works approximately 50% FTE in the gig economy, it is useful to compare this with the median earnings of a part-time worker in the UK.[24] The average gig workers' outside option is just one quarter of this amount, which implies low quality and/or low pay alternative opportunities for these individuals. The mean outside option for the whole population is large—£9,744—because the vast majority of workers will not be drawn into the gig economy given any plausible variation in their economic environment.

Ex ante, the population level correlation between gig work valuations and outside options is ambiguous. Wealthy and high income individuals, who conceivably have high outside options, also likely have low valuations of gig work since they have low marginal utilities of income. Conversely, those with high valuations would work more in the gig economy which entails a greater opportunity cost of time that could be captured by a greater outside option.[25] The estimation suggests the former mechanism

[24]The Office for National Statistics classifies this as somebody who works less than 30 hours per week.

[25]The linear cost to working κ can also capture the opportunity cost of time but this parameter is restricted to be homogeneous so outside options can still pick up the opportunity cost of time associated with high valuations.

Figure 1.7: Model Fit

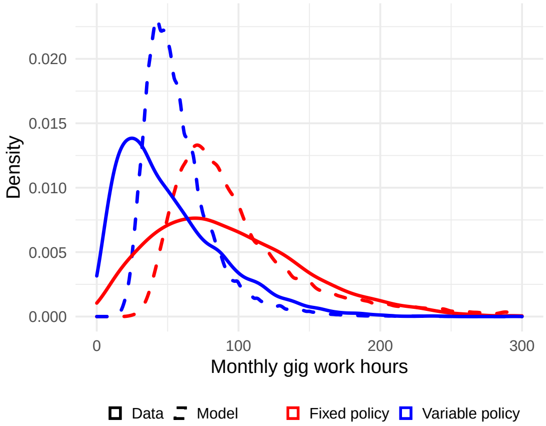

(a) Hours Distribution

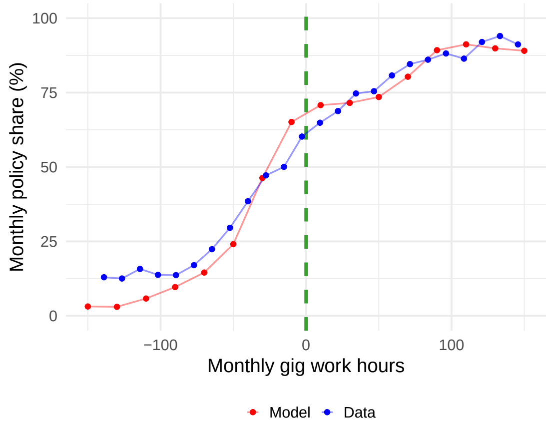

(b) Policy Choice

Notes: These figures plot comparisons of the model output with the data. Panel 1.7a shows the density of hours. The dashed lines reflect the model while the solid lines illustrate the data. The fixed policy holders are shown in red and the variable policy holders are shown in blue. Panel 1.7b is analogous to figure 1.2 but with the models predictions overlaid in red.

dominates and finds a moderate negative correlation between gig work valuations and outside options of -0.62. For gig participants, the negative correlation is largely undone by selection into the gig economy—workers with high outside options must have high valuations to justify their engagement—such that for those in the gig economy the correlation is much smaller at -0.32.

Misperceptions amongst gig workers are not severe because of learning. Nonetheless, the average gig worker is slightly optimistic because those with high misperceptions are more prone to entering. The average participant overestimates the value of gig work ($\phi_i - 1$) by 4%. Further, gig workers' misperceptions ϕ_i exhibit a significant standard deviation of 11%.

The participation decision drives a negative correlation of -0.27 between misperceptions and valuations for gig workers. Intuitively, this reflects two forces. Firstly, regression to the mean; individuals with a high valuation are likely to have a less extreme misperception given the variables are uncorrelated. Secondly, individuals with low valuations are more likely to have optimistic misperceptions, if they are participating in the gig economy.

Aside from individual specific characteristics, the model implies a reasonable degree of concavity of utility with respect to hours worked in the gig economy (*i.e.*, $1-1/\varepsilon$)

Table 1.7: Model Fit

Moment	Data	Model
Labor market share (%)	3.6	3.8
Variable policy share (%)	66.8	68.9
Mean hours per month, variable policy	53.4	58.1
SD hours per month, variable policy	48.4	37.5
Mean hours per month, fixed policy	95.3	104.5
SD hours per month, fixed policy	68.6	74.8
Share non-cost-minimizing (%), variable policy	11.6	9.4
Share non-cost-minimizing (%), fixed policy	60.6	62.3
Mean hours per month from cost-minimizing, variable policy	40.1	40.9
SD hours per month from cost-minimizing, variable policy	36.6	47.3
Mean hours per month from cost-minimizing, fixed policy	55.6	40.8
SD hours per month from cost-minimizing, fixed policy	41.7	31.0
Hazard rate for cost-minimizers (%), variable policy	28.9	23.8
Hazard rate for cost-minimizers (%), fixed policy	20.7	17.6
Hazard rate for non-cost-minimizers (%), fixed policy	29.2	38.3
Decline in hours per month for non-cost-minimizers, variable policy	-10.2	-9.7
Decline in hours per month for non-cost-minimizers, fixed policy	2.6	2.9
Mean quoted fixed premium (£), variable policy	97.0	95.6
Mean quoted fixed premium (£), fixed policy	106.5	104.4
SD quoted fixed premium (£), variable policy	25.1	19.9
SD quoted fixed premium (£), fixed policy	24.3	26.5
Share of zero hours months (%)	5.5	5.9
Mean within worker SD hours per month	31.0	37.7

Notes: This table presents the targeted empirical moments alongside their model implied counterparts. The first column contains the empirical moments and the second column contains the model analogue.

equal to 0.59. The speed of learning parameter implies that misperceptions erode (*i.e.*, $1/(1+\lambda)$) by 50% at the end of the first period of work in the gig economy.

Concerning the stochastic element of the model, the standard deviation of valuations shocks σ_ρ is estimated to be 0.18. This implies that constant individual level valuations are responsible for 79% of the estimated variation in valuations (*i.e.* $\sqrt{\mathbb{V}(\theta_i)/\mathbb{V}(\theta_{i,t}^{\rho})}$). These shocks can cause individuals to temporarily exit the gig economy and receive their outside option at a discount ψ estimated to be 27%.

Table 1.7 compares the empirical moments with those from the model when it is evaluated at the parameter values from table 1.5. Overall, the model fits the 23 empirical moments well; the model's predictions are close to the data. Figure 1.7 provides visual confirmation by contrasting the data with model predictions of policy and hours choices. The quality of the fit supports the view that the estimation captures structural elements of workers engagement with the gig economy.

1.6 Welfare and Counterfactuals

This section describes the gig work surplus implied by the model and estimates of its structural parameters, and considers worker welfare in counterfactual scenarios. Specifically, I analyse the impact of mandatory benefits that workers qualify for by working a sufficient number of hours in the gig economy, and how fixed costs influence the gig work surplus. For the former, I construct a counterfactual experiment to reflect elements of California's Proposition 22, which involves hours thresholds to qualify for a health insurance stipend and, for the latter, I study the introduction of the variable policy. Lastly, I examine how misperceptions stymie the gig work surplus.

Broadly, the analysis reveals a large gig work surplus, which is concentrated amongst low-hours workers. The distribution of the gig work surplus across hours worked suggests that policymakers face a tricky trade-off between guaranteeing costly worker protections and maintaining the appeal of gig work to the majority of participants. Counterfactual analysis of mandatory benefits for workers who reach an hours threshold confirms this; worker welfare falls if they bear even half of the economic incidence associated with the costs of this policy. Fixed costs to gig work, which make dabbling in the gig economy unattractive, impose significant losses. The allocative inefficiency that arises from misperceptions is considerable. Eradicating misperceptions increases the gig work surplus by 21%, which stems almost equally from correcting optimistic and pessimistic perceptions.

1.6.1 The Gig Work Surplus

In this subsection, I measure the size and distribution of the gig work surplus by subtracting workers' outside options from their utility flows to construct a monthly measure of the surplus from gig work. This means, for example, a worker i who temporarily exits receives a negative surplus equal to $-\psi \cdot \nu_i$ and that, naturally, non-participants receive no surplus.

Figure 1.8 presents the estimated distribution of the monthly gig work surplus. The mean monthly surplus for a gig worker equals £1,066, but this masks significant heterogeneity with a standard deviation of £1,775. Moreover, the ratio of the 30^{th} to 70^{th} percentile is equal to 6.5. This median monthly surplus equals £673 or, equivalently,

Figure 1.8: The Gig Work Surplus Distribution

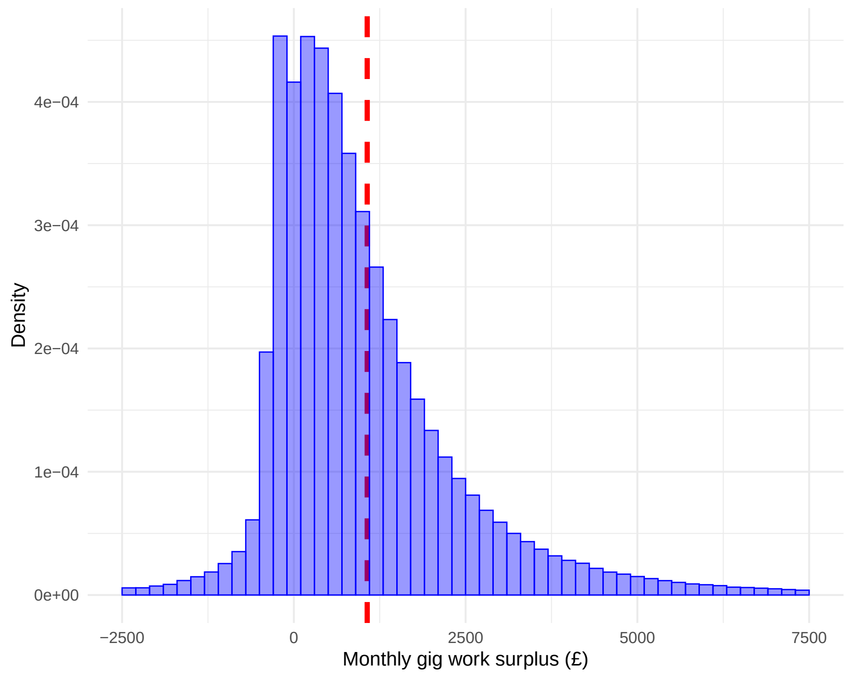

Notes: This figure shows the histogram of the monthly gig work surplus from the estimated model. The red dashed line shows the mean of this distribution.

34% of the median employee's monthly income in the UK.

A small number of workers suffer a negative surplus because of their misperceptions and temporary exit. In an average month, only 16% of participants receive a negative surplus due to misperceptions.[26] However, 45% of each new entering cohort would be better off with their outside option and they lose £1,183 on average. Further, in a typical month, a worker who temporarily exits loses £291 relative to the full value of their outside option. These results reconcile the tension between the huge take up of gig work, which implies a surplus for workers, and the prevalence of negative stories surrounding gig work.

It is also informative to consider where the gig work surplus falls along the hours distribution. Figure 1.9 presents the average monthly gig work surplus by hours bin and the share of the total gig surplus that each bin accounts for. Two clear patterns emerge. Firstly, the surplus increases on average as hours increase. Further, the relationship is roughly linear so that the hourly surplus is more or less constant across the hours distribution. Secondly, the total worker surplus generated by the gig economy is

[26] Jäger et al. (2022) offers a useful comparison to this result. This paper finds that "if workers had correct beliefs, at least 10% of jobs... would not be viable", which is close to the 16% number found here.

Figure 1.9: The Gig Work Surplus by Hours Worked

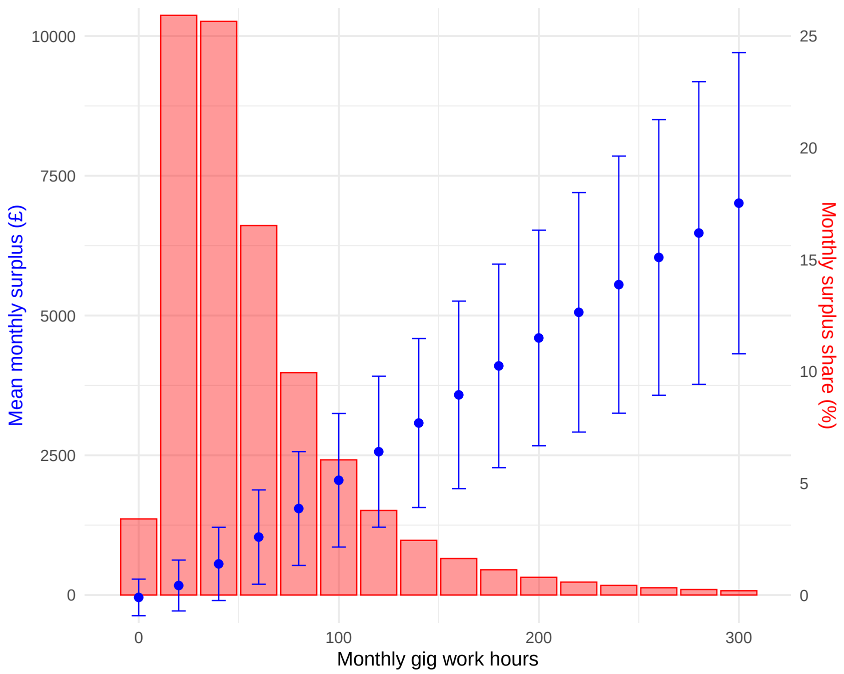

Notes: This figure plots the mean monthly gig work surplus for participants by different hours bins. The error bars show a one standard deviation in the monthly gig work surplus. The red bars show the share of the total gig work surplus accounted for by each hours bin.

concentrated amongst workers who work less the 50% full-time, which has important consequences for the counterfactual experiments that follow. In particular, it makes it difficult for regulators to enshrine protections for those engaged full-time in the gig economy without damaging the surplus of low hours workers, who comprise the majority of the gig work surplus.

I provide three benchmarks to compare these results against. Firstly, the mean hourly surplus is £12.04, which represents between 70 to 80% of the likely wage in the market after accounting for idle time.[27] Secondly, I asked survey respondents "How much would your earnings have to drop per month for you to stop doing this type of work?". With significant caveats, the response to this question should reflect the difference between an individual's value of gig work and their outside option. Comfortingly, the median response is £462, which is roughly in line with the estimates presented

[27] Although the mean hourly surplus is estimated to be less than the likely hourly wage, there is nothing in the model—or in reality—that says the gig work surplus should be constrained by the going wage rate. For example, consider a worker whose next best alternative to gig work is a traditional 9 to 5 job. Further, assume this worker's reservation wage during those hours is greater than their hourly earnings for gig work. The value of gig work for this individual is determined equally by their reservation wages during hours that they do not participate in the gig economy, as it is by wages during their gig work hours.

here.[28] Thirdly, Chen et al. (2019) estimate a median "base" surplus of £571 using an exchange rate of dollars to pounds of 0.8. Thus, the results presented here are consistent with the scant evidence on the gig work surplus.

Discussion. The gig work surplus is unambiguously large, both in terms of its level (*i.e.*, the monthly surplus) and the rate at which it accrues (*i.e.*, the average hourly surplus). Taking the numbers above seriously, at an aggregate level, this part of the gig economy alone generates £15bn (= 32.8 million $\times$ 3.59% $\times$ £1,066 $\times$ 12) in worker surplus annually.

Broadly, workers can derive a large surplus through either a high valuation of gig work or a low outside option, or both. In terms of valuations, workers can derive value from the gig economy through many channels. Income earned from gig work provides utility via consumption. These additional earnings may be particularly valuable if workers have a high marginal utility of income. There are reasons to believe this is the case; most of these individuals are in the bottom half of the income distribution and are likely to receive negative income shocks prior to entry into the gig economy (Cornick et al., 2018; Koustas, 2018). Moreover, a desire from workers to top-up their income with gig work implies a higher marginal utility of income than if this were not the case, *ceteris paribus*. Given institutional arrangements around self-employment, there is also greater scope for tax avoidance, as well as evasion, which can further inflate the value of earnings in the gig economy.

In addition to income, gig work entails a unique combination of amenities. To name a few, workers can pick their hours, do not have a boss, are paid weekly, and receive some income insurance because they can earn more or less depending on their circumstances. Individuals who especially value these amenities are more likely to select into gig work, which elevates the gig work surplus.

A large gig work surplus could also come from low outside options. This could be leisure, which may deliver little utility due to, for example, underemployment, complementarity with income, or other poor employment opportunities. Since gig work hours often fit around traditional work hours, the alternative employment offering may be particularly bad. In the UK, many gig workers are low-skilled migrants with English as

[28] Survey responses were censored from above at £1,000 so it is not possible to compare means.

a second language, which could make traditional employment difficult to find. At an extreme, although food delivery platforms implement right to work checks, there are anecdotal stories that these checks are possible to circumvent (*e.g.*, by working under a different identity) such that gig participants may not have the opportunity to work elsewhere.

1.6.2 Hours-Based Benefits: Proposition 22's Health Insurance Stipend

In 2020, California passed a ballot initiative—Proposition 22—that provided gig workers with a range of new protections, while denying these workers the broader benefits received by employees.[29] One of the benefits under the Proposition includes a health insurance stipend for workers who meet an hours threshold. Precisely, workers who average more than 25 hours per week over a quarter are entitled to 100% of the average premium for a specified health insurance policy. Workers who average between 15 and 25 hours per week over the corresponding period are entitled to half that amount.[30]

While the setting for this paper is the UK, the Proposition provides a useful benchmark to think about the scale and structure of protections that may become available to gig workers in the UK and elsewhere. For this counterfactual experiment, I consider that workers who exceed 100 hours per month receive a fixed pecuniary benefit of £400, and that workers who work between 60 and 100 hours per month receive £200. These numbers are approximately in line with the aforementioned hours threshold and health insurance stipend.

The legislation places the statutory incidence on platforms to pay workers who qualify for benefits but the economic incidence will fall on a combination of the platforms, customers, and workers themselves. Importantly, it is the economic incidence, which depends on a number of factors outside the scope of this paper, that will determine the welfare effects of this policy (Besley and Case, 2000; Gruber and Krueger, 1991; Gruber, 1994).

To this end, I assume that the incidence on workers will manifest as a reduction in all their hourly earnings. An hourly earnings penalty is the most transparent way to model how platforms may pass on any costs. The important aspect of this assumption

[29] Proposition 22 exempted digital rideshare and delivery platforms from Assembly Bill 5.

[30] The text of the law can be found here https://vig.cdn.sos.ca.gov/2020/general/pdf/topl-prop22.pdf, and the section relating to the health insurance stipend is at the bottom of page 32.

is that it imposes a variable cost on workers. Moreover, it is reasonable to assume that all workers will face this penalty because *a priori* platforms cannot determine which workers will qualify for benefits, and multi-homing across platforms will undermine any targeted incidence.

I consider a range of incidence s from zero to full in order to study the welfare impacts of the counterfactual policy. The penalty on hourly earnings $c(\bullet)$ associated with incidence s is found by numerically solving

$$c(s) \cdot H\big(c(s)\big) = s \cdot \Big(200 \cdot N_{200}\big(c(s)\big) + 400 \cdot N_{400}\big(c(s)\big)\Big),$$

where $H(\bullet)$ is the aggregate number of hours worked, and $N_{200}(\bullet)$ and $N_{400}(\bullet)$ denote the number of workers who qualify for the respective benefits.

This policy is non-marginal; it introduces sizeable non-convexities into workers' economic environment and, thus, lends itself to evaluation in a structural model. The presence of hours-based benefits induces several labor supply responses from gig workers that depend on the extent to which the economic incidence falls on workers. Intuitively, the greater the incidence on gig workers, the higher the hourly wage penalty that individuals face. Workers can exhibit four potential labor supply responses: (i) non-participants may join the gig economy as they are encouraged by the benefits, (ii) some gig workers may exit as the wage penalty reduces their surplus below zero, (iii) individuals close enough to the hour thresholds may discretely increase their hours to qualify for the benefits, and (iv) the remainder of participants will reduce their hours as their hourly earnings fall to cover some proportion of the cost of the benefits. Specifically, workers' labor supply in the gig economy will follow

$$h^*_{i,t}(\omega) = \begin{cases} \left(\frac{\hat{\theta}^{\rho}_{i,t}}{p(\omega)+\kappa+c}\right)^{\varepsilon} & \text{if } \rho_{i,t} > \rho_{25}(\nu_i, \hat{\theta}_{i,t}), \\ 25 \times 4 & \text{if } \bar{\bar{\rho}}(\nu_i, \hat{\theta}_{i,t}) < \rho_{i,t} < \rho_{25}(\nu_i, \hat{\theta}_{i,t}), \\ \left(\frac{\hat{\theta}^{\rho}_{i,t}}{p(\omega)+\kappa+c}\right)^{\varepsilon} & \text{if } \rho_{15}(\nu_i, \hat{\theta}_{i,t}) < \rho_{i,t} < \bar{\bar{\rho}}(\nu_i, \hat{\theta}_{i,t}), \\ 15 \times 4 & \text{if } \bar{\rho}(\nu_i, \hat{\theta}_{i,t}) < \rho_{i,t} < \rho_{15}(\nu_i, \hat{\theta}_{i,t}), \\ \left(\frac{\hat{\theta}^{\rho}_{i,t}}{p(\omega)+\kappa+c}\right)^{\varepsilon} & \text{if } \underline{\rho}(\nu_i, \hat{\theta}_{i,t}) < \rho_{i,t} < \bar{\rho}(\nu_i, \hat{\theta}_{i,t}), \\ 0 & \text{if } \rho < \underline{\rho}(\nu_i, \hat{\theta}_{i,t}), \end{cases} \tag{1.17}$$

Figure 1.10: Proposition 22 Counterfactual

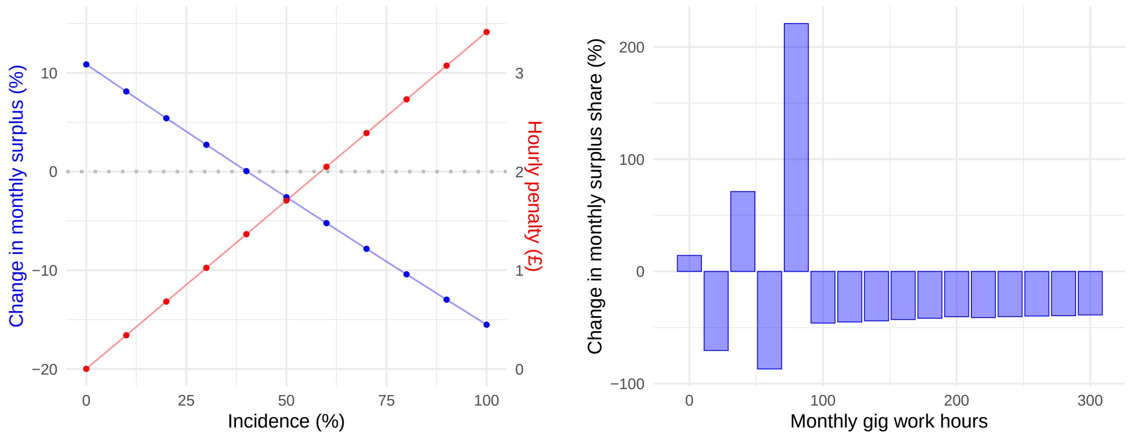

(a) The Gig Work Surplus by Incidence

(b) % Change in the Gig Work Surplus Share by Hours

Notes: This figure plots outcomes under the Proposition 22 counterfactual policy. Panel 1.10a plots the percentage change in the total gig work surplus as a function of incidence in blue. The red line plots the hourly wage penalty associated with the different degrees of incidence. The dotted grey line shows the total surplus in the *status quo*. Panel 1.10b plots the percentage change in the share of the gig work surplus associated with each hours bin at an incidence level of 40%.

where $\underline{\rho}(\bullet)$, $\bar{\rho}(\bullet)$, $\rho_{15}(\bullet)$, $\bar{\bar{\rho}}(\bullet)$, and $\rho_{25}(\bullet)$ are defined in appendix 1.E. Workers' policy choices and participation from the gig economy follow as before via backward induction.

I quantitatively evaluate the welfare effects of this policy. From the policymaker's perspective there is a clear trade-off; higher mandated benefits help those workers who qualify to receive the transfer, but all workers bear the cost through lower hourly wages. This latter factor is exacerbated by the fact that a higher incidence on wages entails a higher wage penalty which discourages work further and necessitates a higher penalty again. Consequently, the efficacy of this policy in terms of raising gig worker welfare will depend critically on the degree of incidence and a complicated set of behavioral responses.

The welfare impacts of this policy are illustrated in figure 1.10. Panel 1.10a shows the gig work surplus under a range of different incidence levels. The analysis suggests that a moderate degree of incidence on workers—anything greater than 40%—would cause this hypothetical policy to decrease worker welfare. This level of incidence leads to a £1.37 drop in hourly earnings. That said, if there is minimal pass-through to workers, the policy could increase worker welfare by as much as 11%.

Panel 1.10b shows how the distribution of the gig work surplus changes across hours worked when incidence is evaluated at 40%. Note that this keeps the size of the gig work surplus the same. The share of the surplus that falls into the bins (40,60] and (80,100] grows substantially because workers bunch at 60 and 100 hours in order to received the mandated benefits. Everywhere else the share of the surplus falls since other hours bins see a decrease in the number of workers and the remaining workers face an hourly wage penalty, which makes them worse off.

The results of this analysis suggest that an appropriately structured minimum wage could complement mandatory benefits by legally limiting the degree of incidence on workers. Indeed, Proposition 22 included a minimum wage requirement that applies while workers are actively on jobs, but it is unclear whether the level set has any bite. It is also possible that platforms could find another margin through which to make workers pay but these are at least less obvious and could be tackled additionally.

1.6.3 Higher Fixed Costs: Introduction of the Variable Policy

The introduction of the variable H&R policy constitutes a real life example of a reduction in the fixed costs associated with gig work. Therefore, in this counterfactual experiment I compare the gig work surplus in the *status quo* with a world without the variable policy. I also consider the removal of the fixed policy for completeness.

Without the variable policy it is possible that fixed policy premiums may adjust. Although there is little evidence of this in this market, I consider welfare absent the variable policy with and without adjustment in fixed premiums. To capture potential adjustment I make two assumptions. First, insurers' costs are linear in hours worked; discussions with the firm indicate that exposure to risk (*i.e.*, time on the road) is the greatest single driver of claims and that the relationship is linear up to a reasonable approximation. Second, the fixed policy market is competitive such that profits are zero (Einav and Finkelstein, 2011).

Under these assumptions I can back out the expected claims from an hour of driving as observed total fixed policy premiums collected divided by the sum of hours worked under the fixed policy. This implies an hourly expected claims cost C of £0.95 per hour, which is very close to the variable policy premium. I assume that the firm will adjust all fixed policy premiums by a proportional amount α, so that a given monthly premium

Table 1.8: % Change in the Gig Work Surplus under Different Scenarios

Scenario	Agg. Surplus	Employment Share	Mean Surplus
Status quo	£40	3.8%	£1,066
No variable policy (NVP)	-4.7	-4.8	0.1
NVP with endogenous P_i	0.8	-2.4	3.2
No fixed policy	0.6	0.3	0.4
No misperceptions	20.6	23.0	-2.0
Half SD of misperceptions	15.2	9.3	5.4
No optimists	8.7	-8.3	18.6
No pessimists	11.7	31.5	-15.0

Notes: This table shows the welfare affects of removing the variable policy. Column 1 shows the scenario considered and columns 2, 3, and 4 show the *per capita* surplus, the employment share of the gig economy, and the gig work surplus conditional on working, respectively. The top panel shows the base levels of these variables, and the bottom panel shows the percentage changes under the different scenarios.

P_i will become $\alpha \cdot P_i$. I solve for α by requiring zero profits

$$H(\alpha) \cdot C = Q(\alpha),$$

where $H(\bullet)$ is aggregate hours worked and $Q(\bullet)$ denotes total premiums collected.

Again, this is a non-marginal counterfactual experiment that is well suited to a structural model. The introduction of the variable policy could spur two labor supply responses: (i) workers may switch policies and reduce the number of hours that they work, and (ii) some workers may enter the gig economy for the first time on the variable policy. These new entrants could benefit or be made worse off, with hindsight, depending on their perceptions.

Table 1.8 compares the welfare outcomes under the *status quo* and the counterfactuals. Without the variable policy and no adjustment in fixed policy premiums the gig work surplus is reduced by 4.7%, which stems mainly from a reduction in gig work participation. This means that, in aggregate, the introduction of the variable policy constituted a £709mn boon for workers. In contrast, welfare would be marginally higher without the fixed policy because, in the *status quo*, overly optimistic individuals select into the fixed policy but would be better off on the variable policy. The removal of the fixed policy benefits these workers more than it harms those who are correctly on the fixed policy.

Interestingly, if fixed policy premiums adjust to the removal of the variable policy, then worker welfare is slightly higher in a world without the variable policy. Participants in the gig economy exhibit a negative correlation between valuations and misperceptions, so that many low hours workers who select the variable policy are overly optimistic. Consequently, a small amount of fixed costs can prevent these individuals from making the mistake of entering. However, this scenario entails that fixed premiums should have risen by 30% since 2017—something that has not been observed. This counterfactual highlights that competition amongst firms that serve gig workers is crucial to the gig work surplus.

1.6.4 Reduced Misperceptions: Alleviating an Allocative Inefficiency

Misperceptions cause an allocative inefficiency in the gig economy: some optimistic workers participate in the gig economy when, with hindsight, they should not, meanwhile some pessimistic individuals would be better off inside the gig economy but are not. This inefficiency is material. The "No misperceptions" counterfactual in table 1.8 reveals that the gig work surplus would be 20.6% higher absent misperceptions.

Interestingly, both sides of this allocative inefficiency are roughly equally responsible for the welfare loss albeit in very different ways. Correcting all optimists' misperceptions causes a 8.3% reduction in participation but a 18.6% increase in the mean surplus. In contrast, if all pessimists are disabused of their misperceptions, there is a 31.5% increase in the employment share and a 15% fall in the typical surplus.

The model suggests that policies aimed at reducing misperceptions could be a fruitful pursuit. Halving the standard deviation of misperceptions obtains three quarters of the welfare gains from eradicating misperceptions. Think-tanks have touted policies that require platforms to increase transparency by, for example, providing predictions of hourly earnings to workers. It remains to be seen how these information treatments would translate to a reduction in misperceptions.

1.7 Conclusion

This paper evinces new empirical facts about workers' engagement with the gig economy: the majority of individuals work less than 50% FTE but a minority of workers work in excess of full-time; individuals select fixed and variable cost structures in line with their hours but many of them still make non-cost-minimizing choices; these decisions are correlated with survival and hours dynamics. Survey evidence supports the view that this results from misperceptions about the value of gig work combined with learning.

I develop a structural model of participation in the gig economy, which captures these core features and is amenable to the evaluation of prospective gig work policies and other counterfactuals. The model fits the data well and, together, they produce precise estimates of the model's structural parameters.

The estimates imply that participants enjoy a surplus of £1,066 per month from gig work. When aggregated, this implies an annual worker surplus of £15bn from a labor market that was nascent a decade ago. Despite this, the model can explain many workers' negative experiences of gig work. Misperceptions cause some individuals to enter the gig economy when they would be better off with their outside option. Learning and endogenous exit ameliorate these losses, although misperceptions still materially reduce the size of the gig work surplus.

Counterfactual experiments reveal that policymakers must be conscious of the incidence of any mandatory benefits that they impose. If the associated cost is passed onto workers in terms of lower hourly earnings, for example, then the gig work surplus will fall sharply as gig work becomes less attractive for the majority of participants. Further analysis suggests that convexifying workers' budget constraints through the gig economy can yield significant welfare gains by allowing individuals to fine-tune their bundle of leisure and consumption. For example, the introduction of the variable policy, which reduced fixed costs for gig work, is estimated to have increased the gig work surplus by 4.7%.

In my view, this study has two main limitations that can serve as a foundation for future research. Firstly, valuations of gig work are inferred from the hours individuals work; the model imposes that individuals who work more in the gig economy value each hour of work more. This does not rank individuals' surpluses since outside options are heterogeneous, but one can imagine other useful information (*e.g.*, wages and preferences over amenities) would enrich the model.

Secondly, individuals' outside options are fixed. In reality, workers choose a bundle of activities and the amount of time to devote to each activity subject to time constraints. Viewing gig work through this perspective and engaging with the IO literature on bundles and discrete-continuous choice problems could offer new insights and, in practice, make outside options endogenous to the opportunities available to workers.

Appendices

1.A Policy Choice

I argue that cost-minimization is the main motivation for workers when they chose their H&R insurance policy. In this appendix, I present supportive evidence and I discuss other potential influences.

1.A.1 Cost-Minimization Motive

The primary difference between the fixed and variable policy is the tariff structure. The variable policy is paid for by the hour while the fixed policy comes with a fixed monthly cost. Therefore, gig workers will select the policy that minimizes their costs *ceteris paribus*. In practice, this means that workers who work less than 110 will generally be better off on the variable policy.

This logic is supported by figure 1.2, which shows the share of workers on the fixed policy across different hours bins. The diagram reveals that workers are increasingly likely to take the fixed policy as they work more hours. This pattern is consistent with a rational cost-minimizer who suffers from imperfect foresight. As described in the main text, a cost-minimizer with perfect foresight would behave in accordance with the dashed green line. Introducing noise into this individual's problem would distort the step function into an upward sloping line that crosses the break-even point at 50%. The exact shape of the line would depend on the distribution of the noise. Thus, the data is consistent with workers that try to cost minimize but who are subject to *ex ante* misperceptions and *ex post* shocks.

The few workers who switch, and who are excluded from the main analysis, can offer further support to the cost minimization hypothesis. Figure 1.A.1 shows that workers who switch policy initially make very poor decisions from a cost-minimization perspec-

Figure 1.A.1: Fixed Policy Share by Hours Worked, Before and After Switching

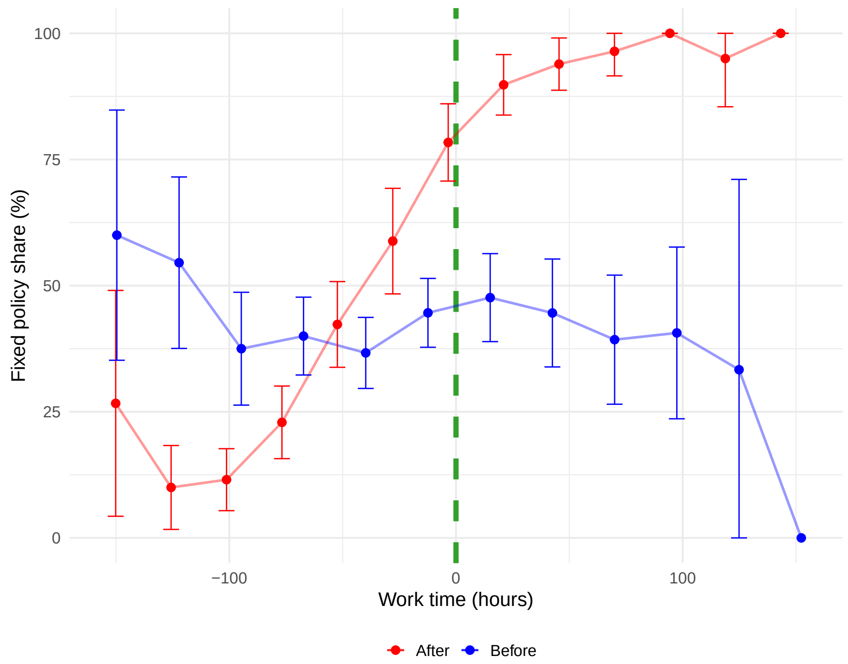

Notes: This figure plots the share of workers who are on the fixed policy by hours bins for switchers before and after they switch. Each observation in a bin is a worker, so the hourly bin than an individual falls into is determined by their average monthly hours. The green dashed line indicates the perfect cost-minimizer's policy choice, which is vertical at the break-even point. Standard errors are constructed by applying the law of large numbers to a Bernoulli random variable (*i.e.*, $\sqrt{p \cdot (1-p)/N}$ where p is the share of policies on the fixed policy in a bin and N is the number of observations in that bin). In total, this figure contains 588 workers.

tive. The blue line describes choices before switching; there is no discernible increase in the probability of opting for the fixed policy as hours worked increase. However, after switching, policy choices reflect a strong cost-minimizing tendency. The increasing red line indicates that the fixed policy share increases with hours. This evidence supports the view that switchers seek to correct non-cost-minimizing choices, which in turn supports the hypothesis that individuals want to minimize costs in their policy choice when they first enter the gig economy.

1.A.2 Alternative Motives

Other factors besides cost minimization could influence the choice between the fixed and variable policies. In this subsection, I provide a taxonomy of these factors.

Coverage. The hourly policy only offers third party coverage, therefore, drivers who desire more comprehensive coverage may select a fixed policy despite higher costs. I deal with this by adjusting premiums for reported WTP for additional coverage from the survey, when individuals opt for greater coverage. Further, I check the robustness of reduced form and structural results to restricting attention to third party only policies in appendix 1.C and 1.F.

Insurance value. The different policy choices imply different variances in costs and overall income. In this case, the policy that minimizes variance in income would have some value aside from its implications for the expected level of costs. Although the fixed policy minimizes variance in costs, the variable policy minimizes variance in overall incomes since increases in costs from more hours worked are offset with higher earnings from work. The overall result is that income is more steady. As such, any insurance value would push in the direction of selecting the variable policy.

Quantitatively, reasonable degrees of risk aversion and the likely magnitude of fluctuations in income imply this story is not a significant driver of any patterns in the data. Consider an individual with CARA preferences and a degree of risk aversion equal to 0.0016 (Handel and Kolstad, 2015). In this case, reducing the standard deviation of monthly income by £35 would increase money-metric utility by less than £1. Moreover, workers exhibit a slightly excessive tendency to opt for the fixed policy, which goes against this mechanism.

Engine size. Only scooters with an engine size of up to 125cc can opt for the variable policy. Fortunately, I can observe engine size with the quote data and exclude them from the analysis. Less than 1% of engine sizes exceed the threshold.

Liquidity. If agents are illiquid, then they may not be able to afford the up-front cost of the fixed insurance policy. This would push such individuals towards the variable policy regardless of their expected hours. Therefore, some of these individuals may find it more economical to select the fixed policy but are not able to do so. Given that few variable policyholders would reduce costs by being on the fixed policy, this does not seem to be a significant friction. Moreover, if illiquid workers are able to access credit, then liquidity issues should not affect their policy choice.

Taxi meter effect. The taxi meter effect would occur in this context if workers receive higher utility gross of insurance costs from the same number of hours on the fixed policy than on the variable policy (Lambrecht and Skiera, 2006). The motivating example, and origin for the name of the effect, is that taxi rides are less enjoyable simply because the customer can see the fare tick up on the taxi meter. Such an effect would push workers to choose the fixed policy over the variable policy *ceteris paribus*.

Present bias. Sophisticated present bias pushes workers towards selecting the fixed policy, while the effect of naïve present bias is ambiguous. An agent who is aware of their present bias (*i.e.*, sophisticated) may prefer to opt for the fixed policy to correct their inefficiently low level of labor supply (Lockwood, 2020). In the same vein, naïve present bias may push workers to choose the fixed policy since they overestimate how much they will work in the future (Augenblick and Rabin, 2019). On the other hand, naïve present bias may cause workers to favour the variable policy to avoid the upfront cost of the fixed policy, if workers suffer from present-bias over money or the timely consumption opportunities that it brings.

1.B Data Cleaning and Filters

In this appendix, I describe the data cleaning procedures and the subsequent filters I apply to the data in order to arrive at my analysis sample. I also discuss the survey and quotes data in more detail.

1.B.1 Data Cleaning and Filters

The raw data is in calendar month tranches, which were received from the firm, with observations at the job level (*i.e.*, each observation is a food delivery). I collapse each tranche down to the user-policy level, or just to the user level in the case of variable policy holders. Then, a variable policy holder's monthly observation is associated with a calendar month, while I merge fixed policy holders' policies that were divided over two calendar months such that their month corresponds to a 30 day policy duration. I then drop any user-policy-month-year duplicates, and removed any fixed policies that exceed 30 days. Some fixed policy holders' policies are further fragmented because of changes to their policy over the course of its duration (*e.g.*, a customer might have switched their coverage). I combine these policies by checking start and finish dates of the fragmented policies to see combinations of policies that consist of precisely 30 days. After this step, I drop any policies that are shorter than 14 days. Lastly, I trim observations according to monthly work time and premiums at the 0.1% percentile in order to remove outliers.

Hours of work may be understated for variable policy holders in the their first month because an observation is associated with a calendar month such that they may begin work halfway through a month. To deal with this, I *pro rata* work hours of these individuals according to what they did work while they were active in that month and, if there is less than two weeks left in the month, then I drop these partial observations. I have tested robustness to the two week threshold (*e.g.* using less than one week or less than three weeks) and the effect on the data is minimal. A similar problem arises for both fixed and variable policyholders in their final month before exiting; if workers leave after one week into their final month then their hours are not reflective of their engagement in the gig economy. I resolve this issue analogously.

I apply four filters after data cleaning. Firstly, I remove annual policy holders, hence, their omission from the data cleaning discussion. Secondly, I keep only individuals who start after January 1 2019 since before this point the firm did not offer a consistent menu of policies. Thirdly, I identify workers who have more than one stint in the gig economy by flagging breaks of four months or longer. For these workers, I keep only their first stint in the gig economy. Finally, I remove individuals from the main analysis who switch policy during their initial spell.

1.B.2 Survey Data

I complement the administrative data with survey data. The survey was conducted through the firm by emailing customers, if they had subscribed to receive promotional material. Fixed policyholders were over-sampled because they have a greater tendency to subscribe affirmatively. 500 workers started the survey, of which 336 completed it. Of these, 251 are on the fixed policy.

The survey was sent out twice (the second time as a reminder) in June 2022. Therefore, the workers surveyed are not necessarily in my administrative data and, despite

Figure 1.B.1: Empirical Distribution of Fixed Premiums

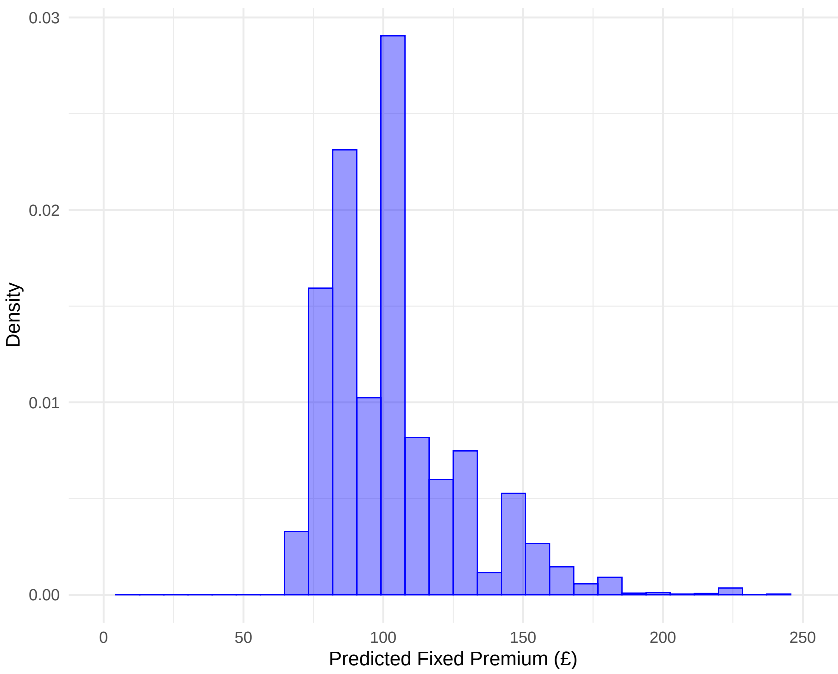

Notes: This figure plots the empirical distribution of offered fixed policy premiums from the firms quote data.

efforts, they cannot be merged. Therefore, to construct categories (*e.g.*, minimizers) I rely on self-reports of hours and the premiums that they face.

1.B.3 Quotes Data

The firm sent data on the quotes that they have offered to all enquirers over the analysis period. This allows for the observation of an unselected distribution of quotes and, in particular, I can view the premiums offered to variable policyholders.

From the fixed policyholders, it is clear that quoted premiums typically fall shy of realized premiums. To resolve this issues, I non-parametrically predict realized premiums with quoted premiums for fixed policy holders by calculating the average realized premium within £5 quoted premium bins. Then, I use this model to construct counterfactual realized premiums for all other customers. Figure 1.B.1 shows the resulting distribution of premiums.

1.C Robustness Checks: Reduced Form Evidence

In the main body of the paper, I present a number of reduced form facts. In this appendix, I present robustness checks which show that the patterns in the data persist across different subgroups of the population, with different definitions of categories, and over different time periods in the data. I also present additional empirical evidence.

Figure 1.C.1 shows the analogue of figure 1.2 for various cuts of the data. Panel 1.C.1a shows the share of individuals on the fixed policy for those under and over 30, separately. Broadly, the two age groups show the same pattern, although over 30s are slightly more effective at minimizing costs by selecting the variable policy more frequently, when hours are below the break-even point. Panel 1.C.1b shows the equivalent data when including all types of coverage and when keeping soley third party only policies. Since only fixed policies can include levels of coverage other than third party only, the blue line is mechanically lower than the red line, where the latter reflects the fixed policy share for all types of coverage. Panel 1.C.1c shows the share of individuals on the fixed policy for the main analysis sample and for the sample where individuals whose predicted premiums exceed £175 are excluded. Panel 1.C.1d shows the fixed policy share for the period of time before the Covid pandemic, and for the course of the Covid pandemic during which there were several lock-downs in UK, where food delivery riders could still work. For transparency, figure 1.C.2 shows the fixed policy share by non-normalized hours where the dashed green line represents a perfect cost minimizer who faces the average fixed premium and the variable premium.

Figure 1.C.4 shows that the patterns of survival persist for the categories across various robustness checks. Panels 1.C.3a and 1.C.3b use alternative definitions to construct categories. The former categorizes workers based on whether they minimize their bill of the course of their tenure in the gig economy; and the latter compares average hours with break-even points to categorize workers. Panels 1.C.3c and 1.C.3d show the patterns of survival for different categories, using the baseline definition, for under and over 30s, respectively. Panels 1.C.3e and 1.C.3f illustrate the same data for the period before the Covid pandemic and for the course of the Covid pandemic, respectively. Panel 1.C.3g exhibits the pattern of survival for different categories, restricting to those on third party only policies. Panel 1.C.3h shows the survival trajectories for non-third party

Figure 1.C.1: Fixed Policy Share by Normalized Hours for Different Samples

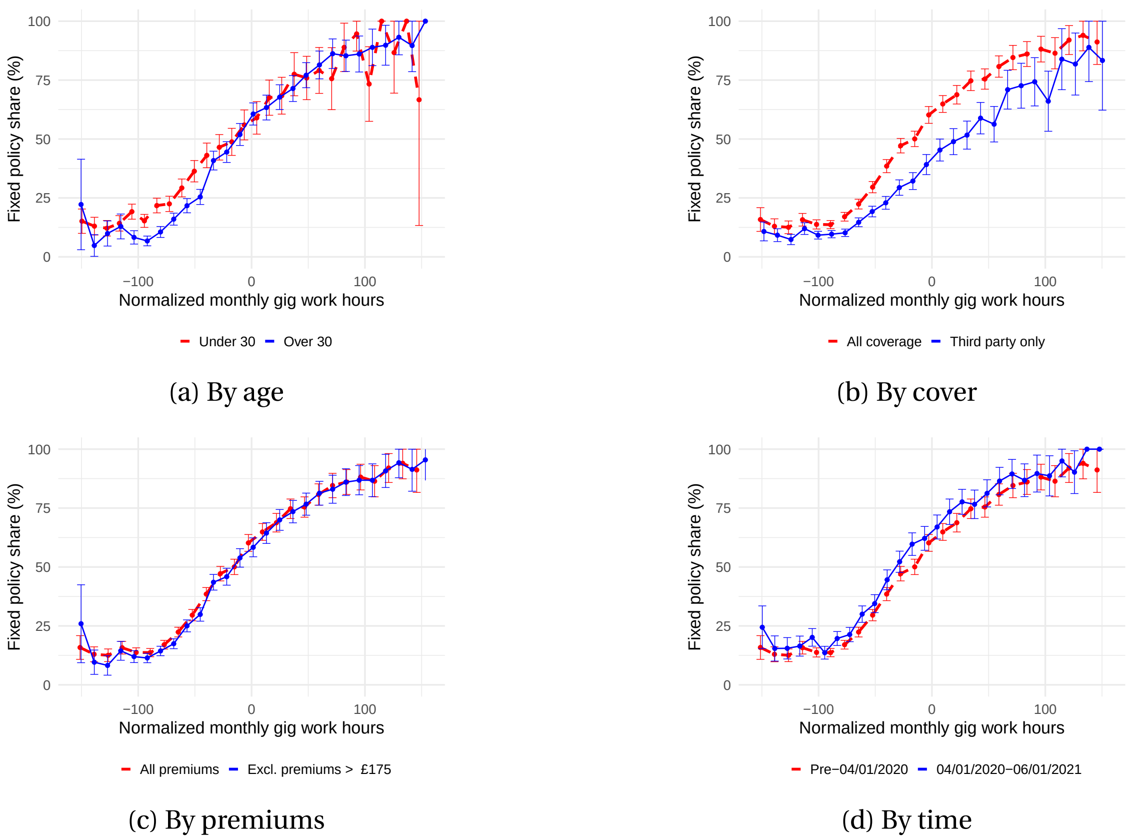

Notes: These figures plot the share of individuals on the fixed policy by different groups. Standard errors are calculated by applying the law of large numbers to the average of a random variable that follows the binomial distribution with one trial (*i.e.*, $\sqrt{p_j \cdot (1-p_j)/N}$ where p_j is the share of responses for a given category j and N is the number of observations).

Figure 1.C.2: Raw Fixed Policy Share

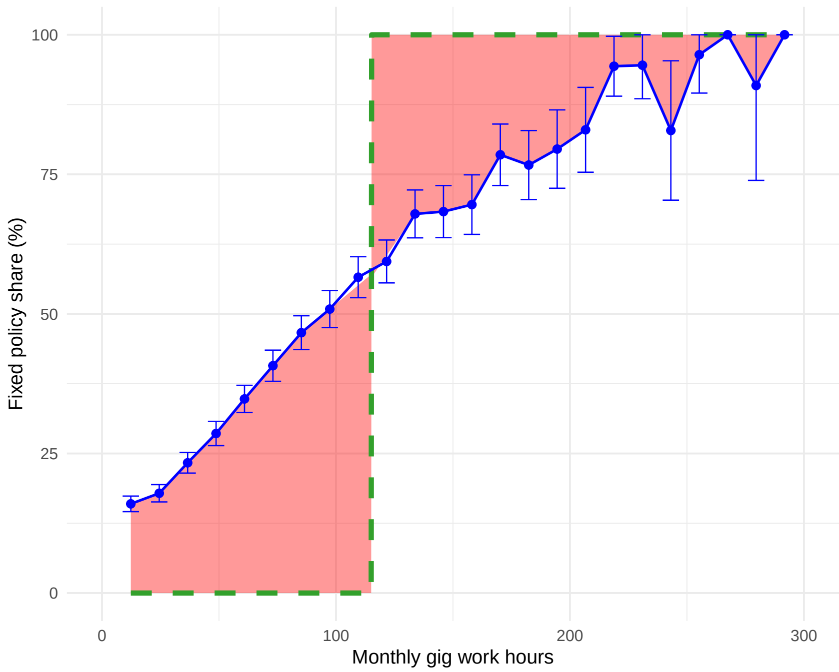

Notes: This figure plots the share of workers who are on the fixed policy by monthly hours bins. Each observation in a bin is a worker, so the hourly bin that an individual falls into is determined by their average monthly hours. The green dashed line indicates the perfect cost-minimizer's policy choice, which is vertical at the break-even point. Standard errors are constructed by applying the law of large numbers to the average of Bernoulli random variables (*i.e.*, $\sqrt{p \cdot (1-p)/N}$ where p is the share of policies on the fixed policy in a bin and N is the number of observations in that bin).

only policies by categories. These are necessarily fixed policies, so the diagram excludes pessimistic workers. Panel 1.C.3i shows the survival curves of minimizers broken down by fixed and variable policy holders. Panel 1.C.3j shows the survival patterns of workers who select different types of coverage. Panel 1.C.3k shows that survival function for workers preceding and during the Covid pandemic.

Table 1.C.1 reproduces table 1.4 with alternative controls. Columns (1) to (4) vary the controls included in the Cox proportional hazards model. In order, the columns exclude cover, gender, age, and low hours controls. Table 1.C.2 and 1.C.3 show the results from linear probability models, which show the same patterns as table 1.4 using the same controls. Table 1.C.2 shows an OLS regression where the outcome variable is the exiting the gig economy within the specified periods of time, while 1.C.3 shows the analogue conditional on reaching that tenure.

Figure 1.C.4 shows the trajectory of hours, like figure 1.4, for subgroups of the data.

Figure 1.C.3: Survival by Categories for Different Definitions & Samples

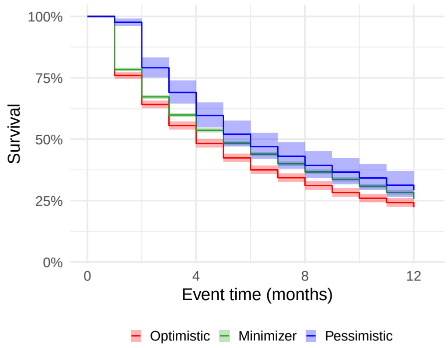

(a) Alternative categories: bill minimization

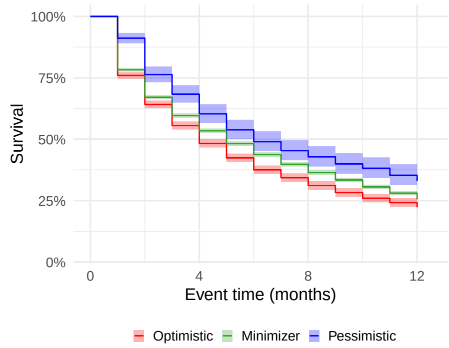

(b) Alternative categories: average hours

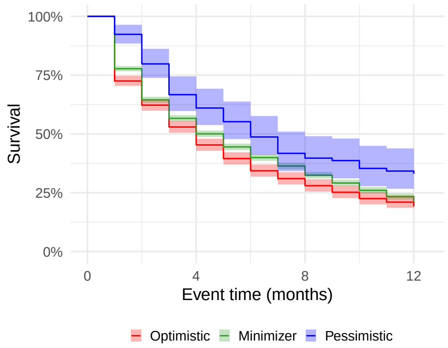

(c) Under 30s

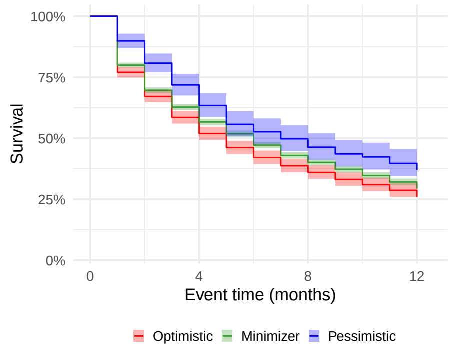

(d) Over 30s

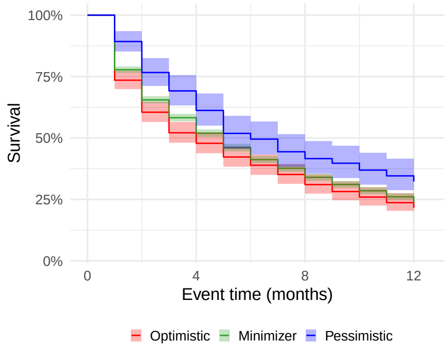

(e) Pre-04/01/202

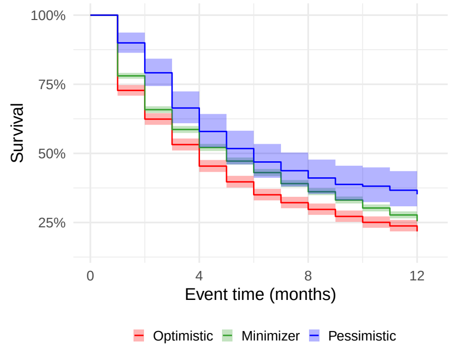

(f) 04/01/2020-06/01/2021

Notes: This figure plots Kaplan-Meier survival curves for different groups in the gig economy. The green, red, and blue lines denote the minimizers, optimistic, and pessimistic categories, respectively. Event time is tenure month in the gig economy (*i.e.*, $t = 1$ is workers' first month in the gig economy so if an individual does not have a second month in the gig economy, then they exit in the first period).

Figure 1.C.3: Survival by Categories for Different Definitions & Samples

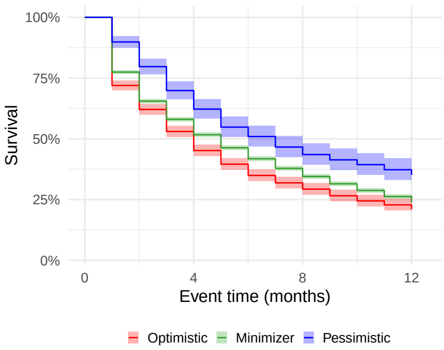

(g) Third party only

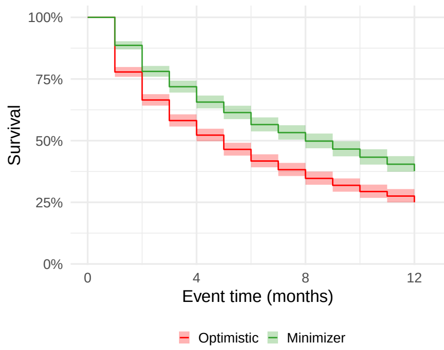

(h) Non-third party only

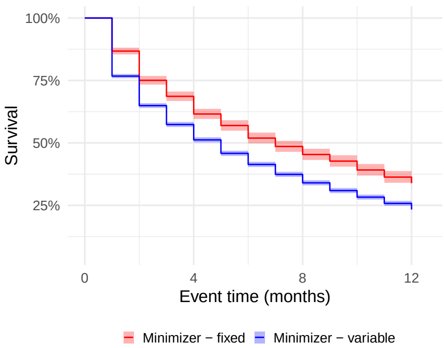

(i) By type of minimizer

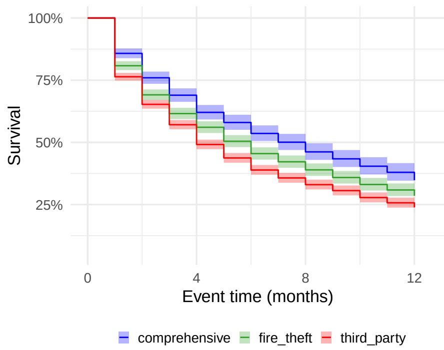

(j) By cover

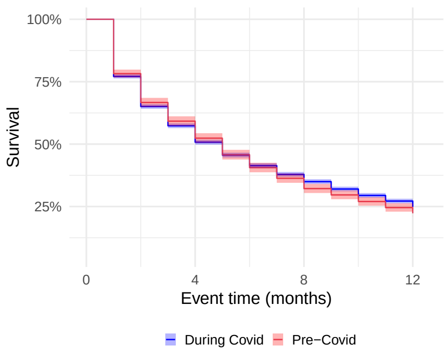

(k) By time

Notes: This figure plots Kaplan-Meier survival curves for different groups in the gig economy. The green, red, and blue lines denote the minimizers, optimistic, and pessimistic categories, respectively. Event time is tenure month in the gig economy (*i.e.*, $t = 1$ is workers' first month in the gig economy so if an individual does not have a second month in the gig economy, then they exit in the first period).

Table 1.C.1: Cox Proportional Hazards Model with Time-Varying Coefficients

	Dependent variable:			
	Tenure in the gig economy (months)			
	No cover ctrl	No gender ctrl	No age ctrl	No low hours ctrl
	(1)	(2)	(3)	(4)
Mean hours	-0.001^{***}	-0.001^{**}	-0.001^{**}	-0.006^{***}
	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Minimizer (<= 2 months)	0.179^{*}	0.222^{**}	0.262^{***}	0.181^{*}
	(0.099)	(0.100)	(0.099)	(0.099)
Optimistic (<= 2 months)	0.388^{***}	0.480^{***}	0.535^{***}	0.322^{***}
	(0.103)	(0.104)	(0.104)	(0.104)
Minimizer (> 2 months)	−0.095	−0.053	−0.007	-0.289^{***}
	(0.070)	(0.070)	(0.070)	(0.070)
Optimistic (> 2 months)	0.082	0.172^{**}	0.230^{***}	−0.122
	(0.076)	(0.078)	(0.078)	(0.078)
Low hours	Yes	Yes	Yes	No
Time controls	Yes	Yes	Yes	Yes
Age	Yes	Yes	No	Yes
Gender	Yes	No	Yes	Yes
Cover	No	Yes	Yes	Yes
Observations	23,969	23,969	24,013	23,969
R^2	0.075	0.075	0.073	0.038

Notes: *p<0.1; **p<0.05; ***p<0.01. This table shows estimates from a Cox proportional hazards model with time varying coefficients on the categories variable; the effect of this factor variable is allowed to differ between the first two months of a workers spell and any remaining months. Coefficients reflect percentage changes in the base hazard rate associated with the corresponding variable. The main panel of the table shows estimates for these coefficients and estimates of the coefficient on a workers average number of hours per month. Standard errors are shown in parentheses.

Table 1.C.2: Linear Probability Model

	Dependent variable:			
	Tenure <= 2 months		2 < Tenure <= 6 months	
	Controls	No Controls	Controls	No Controls
	(1)	(2)	(3)	(4)
Mean hours	0.0002**	0.0002*	0.0001	0.0002**
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Optimistic	0.083***	0.091***	0.010	0.025***
	(0.010)	(0.009)	(0.010)	(0.009)
Pessimistic	−0.076***	−0.066***	0.076***	0.067***
	(0.020)	(0.021)	(0.020)	(0.020)
Low hours	Yes	Yes	Yes	Yes
Time controls	Yes	No	Yes	No
Age	Yes	No	Yes	No
Gender	Yes	No	Yes	No
Cover	Yes	No	Yes	No
Observations	14,795	15,924	14,795	15,924
R^2	0.163	0.077	0.135	0.009

Notes: *p<0.1; **p<0.05; ***p<0.01. This figure plots the coefficients from an OLS regression, where the outcome is a binary variable that takes value one, if that worker survived for either less than two months—columns (1) and (2)—or between 2 and 6 months—columns (3) and (4). Columns (2) and (4) have no controls, similar to column (3) in figure 1.4. Columns (1) and (3) control for all available covariates, as in column (1) in figure 1.4

Table 1.C.3: Linear Conditional Probability Model

	Dependent variable:			
	Tenure <= 2 months		2 < Tenure <= 6 months	
	Controls	No Controls	Controls	No Controls
	(1)	(2)	(3)	(4)
Mean hours	0.0002**	0.0002*	0.0003**	0.0005***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Optimistic	0.083***	0.091***	0.105***	0.125***
	(0.010)	(0.009)	(0.013)	(0.013)
Pessimistic	−0.076***	−0.066***	0.038*	0.049*
	(0.020)	(0.021)	(0.023)	(0.025)
Low hours	Yes	Yes	Yes	Yes
Time controls	Yes	No	Yes	No
Age	Yes	No	Yes	No
Gender	Yes	No	Yes	No
Cover	Yes	No	Yes	No
Observations	14,795	15,924	9,174	9,805
R^2	0.163	0.077	0.263	0.021

Notes: *p<0.1; **p<0.05; ***p<0.01. This figure plots the coefficients from an OLS regression, where the outcome is a binary variable that takes value one, if that worker survived for either less than two months—columns (1) and (2)—or between 2 and 6 months conditional on surviving beyond 2 months—columns (3) and (4). Columns (2) and (4) have no controls, similar to column (3) in figure 1.4. Columns (1) and (3) control for all available covariates, as in column (1) in figure 1.4.

Panels 1.C.4a and 1.C.4b shows the trajectory of hours for categories using alternative definitions, which are analogous to those in figure 1.C.4. Panels 1.C.4c and 1.C.4d show the trajectory of hours over time for under and over 30s respectively. Panels 1.C.4e and 1.C.4f show the dynamics of hours before and during the Covid pandemic. Panel 1.C.4g shows the baseline figure but without enforcing a balanced panel. Lastly, panel 1.C.4h shows the trajectory of hours for third party only policyholders.

Figure 1.4 shows the dynamics of hours at the weekly level. They can also be displayed at the monthly level—this is done in figure 1.C.5. I do this for all three definitions of the categories that I use.

Figure 1.C.6 shows responses to the survey of gig workers' experiences for self-reported optimistic and pessimistic workers, respectively. To construct this figure, I subtract the share of responses by minimizers from those of the other categories in order to illustrate the relative prevalence of responses. Moreover, although the differences are not statistically significant, in order to get some reasonable precision I aggregated all the questions asked in the survey about workers experiences. These are questions about earnings, costs, and the difficulty of work. The figure shows that pessimistic workers are most likely to report aspects of gig work are better than they expected. Meanwhile, minimizers are most likely to report experiences as expected. Lastly, optimistic workers most frequently report gig work to be worse than expected.

Finally, figure 1.C.7 shows the aggregate responses to all experience questions at the finest level of response. There is a tendency to find gig work worse than expected, which is consistent with some evidence of optimism in fixed policy choice and results from the model. However, the majority of the mass falls in the "As expected" and "A little worse" bins, which indicates misperceptions are not too severe—again consistent with the model's results.

Figure 1.C.8 shows the distribution of individual break-even points.

Table 1.C.4 presents summary statistics for the worker-level covariates from the quotes data.

Table 1.C.5 shows the average level of worker-level covariates within each category.

Figure 1.C.4: Hours Worked Over Time by Categories for Different Definitions & Samples

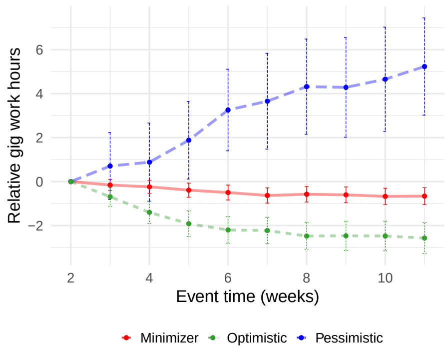

(a) Alternative categories: bill minimization

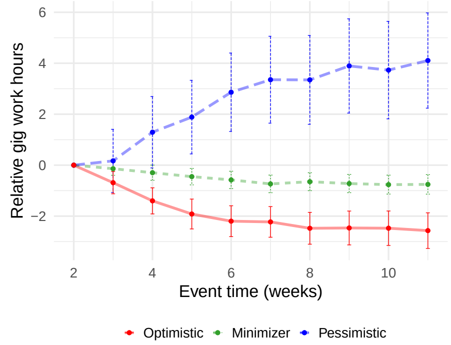

(b) Alternative categories: average hours

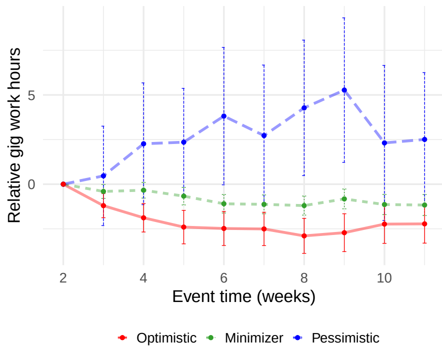

(c) Under 30s

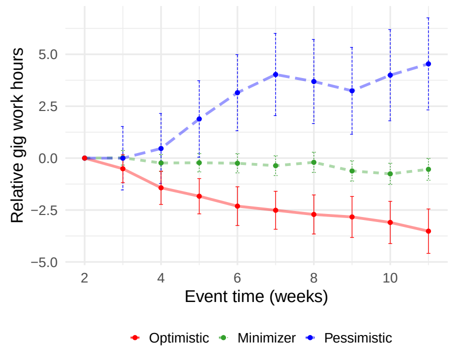

(d) Over 30s

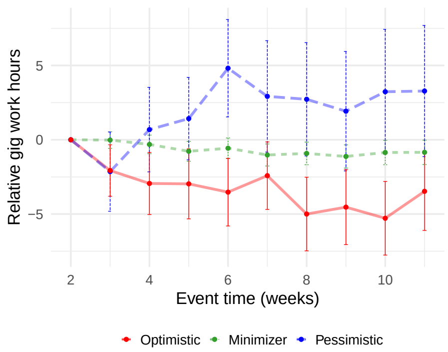

(e) Pre-04/01/202

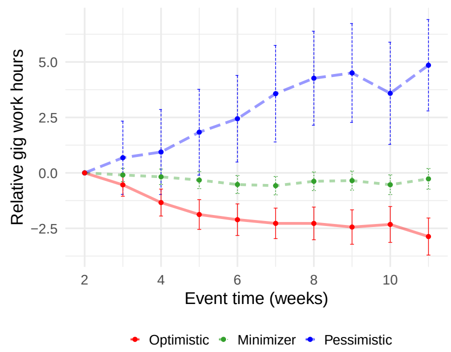

(f) 04/01/2020-06/01/2021

Notes: These figures plot three sets of coefficients from three separate regressions, which are run on a balanced panel (unless otherwise specified) of each category of worker. Weekly hours are regressed on fixed effects and event time dummies, where $t = 2$ corresponds to their second month in the gig economy (*i.e.*, event time is tenure in the gig economy), as well as calendar time controls. Intuitively, each coefficient represents the difference in hours at a given point in time from their hours in their first month in the gig economy. SEs are clustered at the worker level with a HC3 weighting scheme.

Figure 1.C.4: Hours Worked Over Time by Categories for Different Definitions & Samples

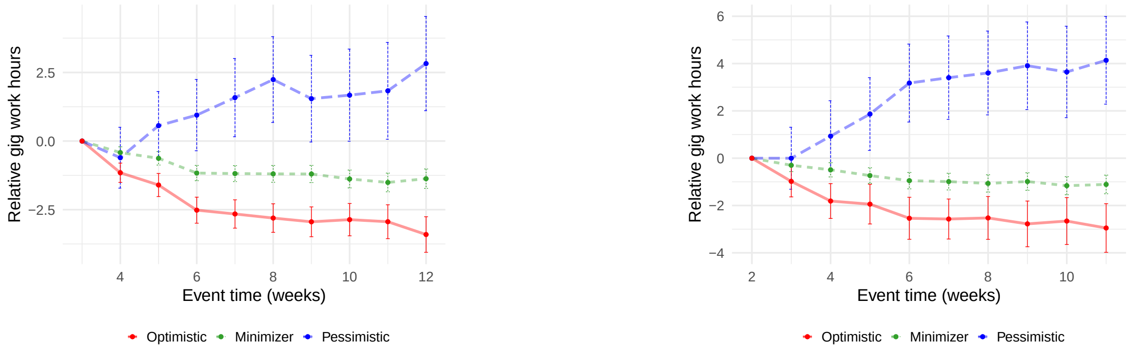

(g) Not balanced

(h) Third party only

Notes: These figures plot three sets of coefficients from three separate regressions, which are run on a balanced panel (unless otherwise specified) of each category of worker. Weekly hours are regressed on fixed effects and event time dummies, where $t = 2$ corresponds to their second month in the gig economy (*i.e.*, event time is tenure in the gig economy), as well as calendar time controls. Intuitively, each coefficient represents the difference in hours at a given point in time from their hours in their first month in the gig economy. SEs are clustered at the worker level with a HC3 weighting scheme.

Table 1.C.4: Worker Covariate Summary Statistics

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Age	15,411	32.227	7.811	20	26	31	37	65
Gender	15,440	0.913	0.281	0	1	1	1	1
Cover	16,575	0.813	0.390	0	1	1	1	1
Licence	15,412	5.585	5.991	0	1	3	8	45

Notes: This figure shows summary statistics for worker-level covariates. The gender variable was received as binary and takes value one, if male, and zero otherwise. Cover is coded so that third party coverage takes value one, while fire and theft and comprehensive cover take value zero. The licence variable reports how long a worker has had their licence for.

Table 1.C.5: Covariates by Categories

Categories	Age	Licence	Cover	Gender
Minimizer	33.00	5.95	0.87	0.93
Optimistic	32.12	5.32	0.50	0.93
Pessimistic	36.88	6.64	1	0.93

Notes: This table shows the average level of worker covariates in each category.

Figure 1.C.5: Hours Worked Over Time by Category at a Monthly Frequency

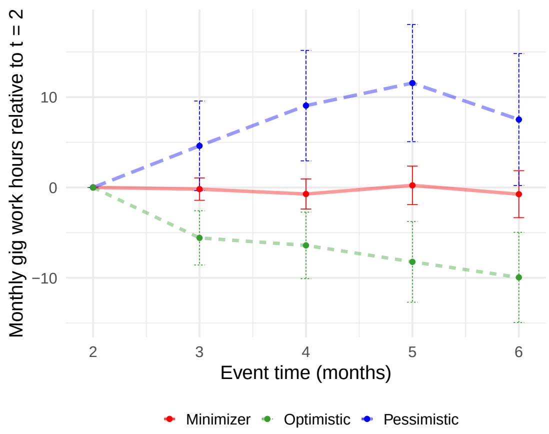

(a) By baseline categories

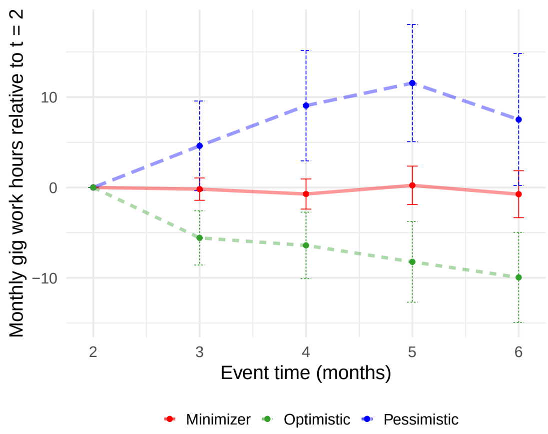

(b) By alternative categories: bill minimization

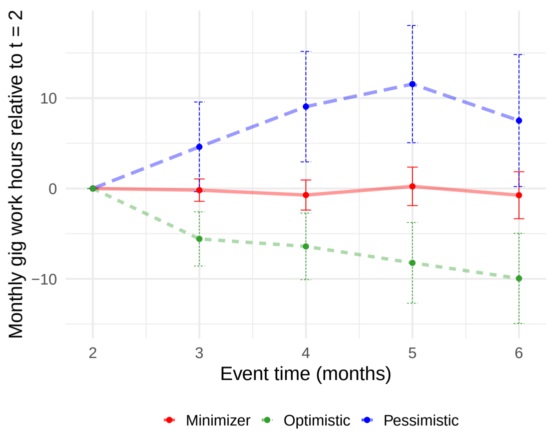

(c) By alternative categories: mean hours

Notes: These figures plot three sets of coefficients from three separate regressions, which are run on a balanced panel of each category of worker. Monthly hours are regressed on fixed effects and event time dummies, where $t = 2$ corresponds to their second month in the gig economy (*i.e.*, event time is tenure in the gig economy), as well as calendar time controls. Intuitively, each coefficient represents the difference in hours at a given point in time from their hours in their first month in the gig economy. SEs are clustered at the worker level with a HC3 weighting scheme.

Figure 1.C.6: Survey Responses by Category Relative to Minimizers

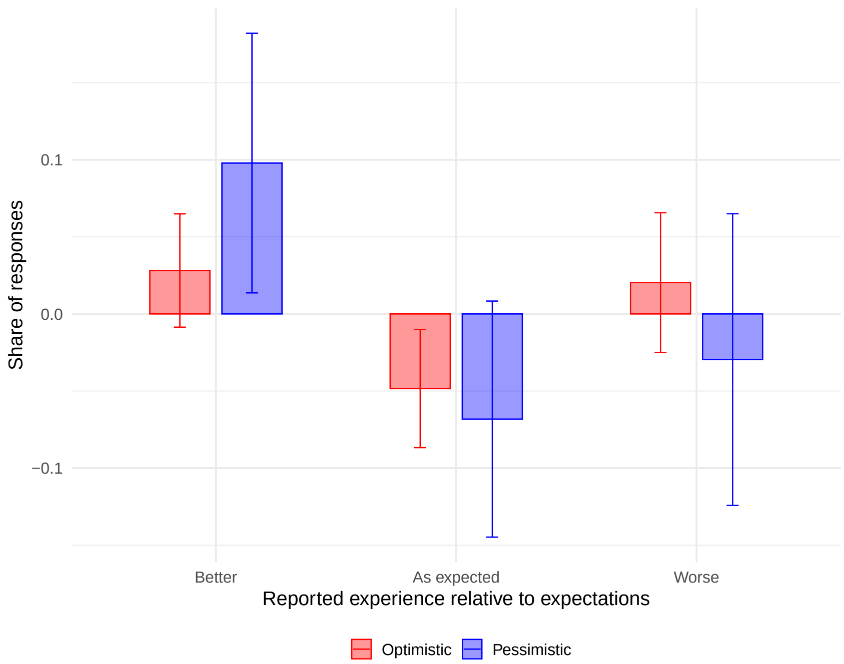

Notes: This figure shows the relative frequency of responses to an aggregate measure of workers' experiences in the gig economy.

Figure 1.C.7: Survey Responses by Category Relative to Minimizers

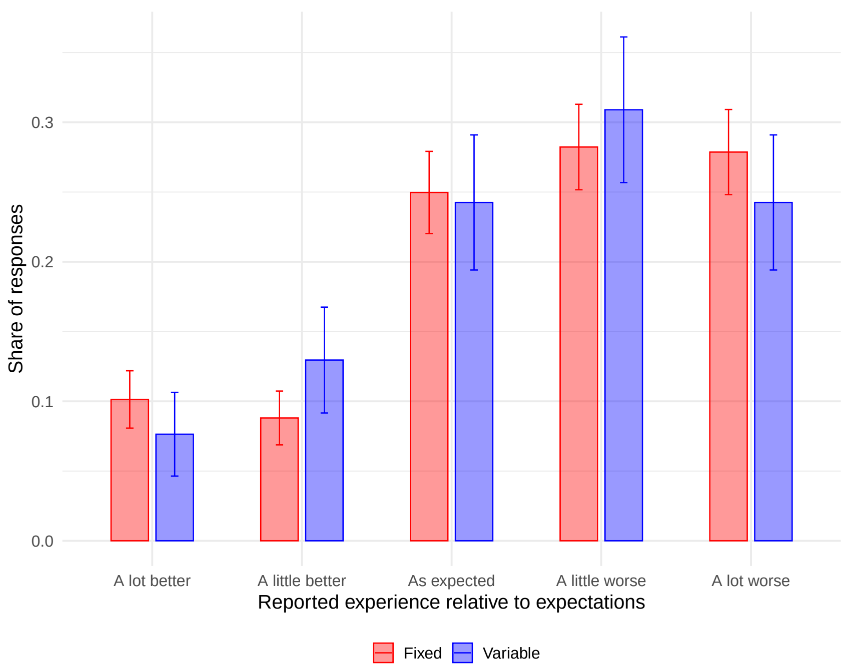

Notes: This figure shows the relative frequency of responses to an aggregate measure of workers' experiences in the gig economy.

Figure 1.C.8: Distribution of Break-Even Points

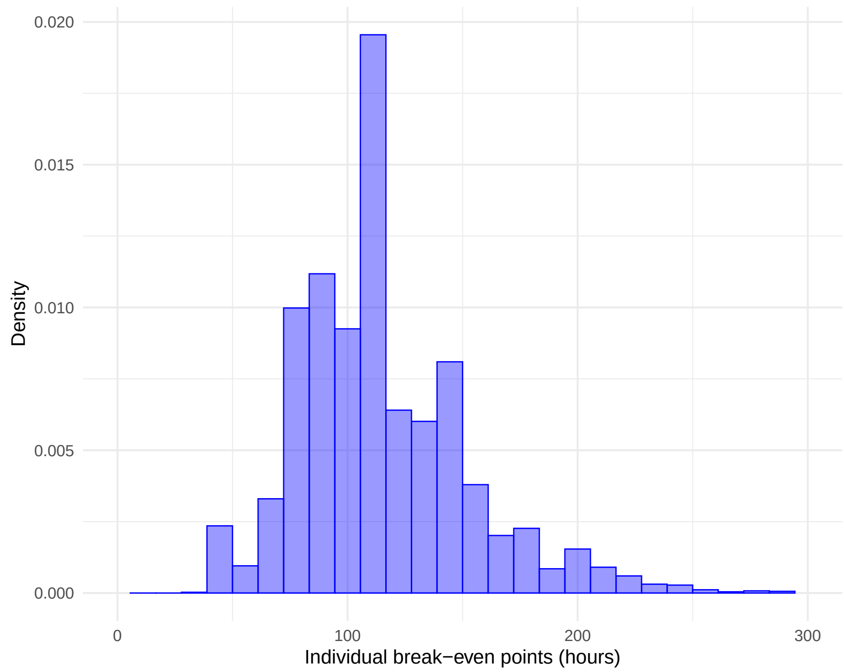

Notes: This figure shows the distribution of individual-level break-even points, which are constructed as a worker's quoted fixed premium divided by the variable premium.

1.D Bayesian Learning

In this appendix, I show that the specified functional form for how gig worker's learn over time is derived from a micro-founded learning process, where individuals update their priors in a Bayesian fashion as they learn from normally distributed signals. As an example, I do this in the context of taxi drivers who learn about their wage rate.

1.D.1 The Environment

Drivers receive a signal $\log(w_i) + \mu_{i,t}$ about their true wage rate $\log(w_i)$ after a weeks' driving is finished. So for week zero, drivers drive according to their prior $\log(w_{i,0})$. Note that this perceived wage rate also has an associated variance σ_0^2 which forms an exogenous, initial prior $\mathrm{N}(\log(w_{i,0}), \sigma_0^2)$ over the true wage rate $\log(w_i)$. Then, in week one, they update $\log(w_{i,0})$ to $\log(w_{i,1})$ using their prior, the signal, and the variance of the distribution from which the signal is drawn, where $\log(w_i) + \mu_{i,t} \sim \mathrm{N}\big(\log(w_i), \sigma_\mu^2\big)$.

Given homoskedastic variance across drivers' priors and signals, the perceived wage rate $\log(w_{i,t})$ and it's variance σ_t^2 acquires a convenient form

$$\log(w_{i,t}) = \frac{\sigma_\mu^2 + \sigma_0^2}{\sigma_\mu^2 + (t+1)\cdot\sigma_0^2}\cdot\log(w_{i,0}) + \frac{t\cdot\sigma_0^2}{\sigma_\mu^2 + (t+1)\cdot\sigma_0^2}\cdot\log(w_i) + \frac{\sigma_0^2}{\sigma_\mu^2 + (t+1)\cdot\sigma_0^2}\cdot\sum_{k=0}^{t}\mu_{i,t-k} \tag{1.18}$$

$$\sigma_t^2 = \frac{\sigma_0^2}{1 + t\cdot\frac{\sigma_0^2}{\sigma_\mu^2}}.$$

Note that the mean of beliefs is a homographic function in time, as is the learning process that I specify in section 1.4.

To see this, consider the variance of beliefs over time

$$\begin{aligned}
\sigma_1^2 &= \frac{\sigma_\mu^2\cdot\sigma_0^2}{\sigma_\mu^2+\sigma_0^2}\\
\sigma_2^2 &= \frac{\sigma_\mu^2\cdot\sigma_1^2}{\sigma_\mu^2+\sigma_1^2}\\
&= \frac{\sigma_\mu^2\cdot\left(\frac{\sigma_\mu^2\cdot\sigma_0^2}{\sigma_\mu^2+\sigma_0^2}\right)}{\sigma_\mu^2+\left(\frac{\sigma_\mu^2\cdot\sigma_0^2}{\sigma_\mu^2+\sigma_0^2}\right)}\\
&= \frac{\sigma_0^2}{1+2\cdot\frac{\sigma_0^2}{\sigma_\mu^2}}\\
\sigma_3^2 &= \frac{\sigma_\mu^2\cdot\left(\frac{\sigma_0^2}{1+2\cdot\frac{\sigma_0^2}{\sigma_\mu^2}}\right)}{\sigma_\mu^2+\left(\frac{\sigma_0^2}{1+2\cdot\frac{\sigma_0^2}{\sigma_\mu^2}}\right)}\\
&= \frac{\sigma_0^2}{1+3\cdot\frac{\sigma_0^2}{\sigma_\mu^2}}\\
&\vdots\\
\implies \sigma_t^2 &= \frac{\sigma_0^2}{1+t\cdot\frac{\sigma_0^2}{\sigma_\mu^2}}.
\end{aligned}$$

The variance of beliefs are used to weight signals, and so the mean of beliefs over time look like

$$\begin{aligned}
\log(w_{i,t}) &= \frac{\sigma_\mu^2 + t \cdot \sigma_0^2}{\sigma_\mu^2 + (t+1) \cdot \sigma_0^2} \cdot \log(w_{i,t-1}) + \frac{\sigma_0^2}{\sigma_\mu^2 + (t+1) \cdot \sigma_0^2} \cdot \big(\log(w_i) + \mu_{i,t}\big) \\
&= \frac{\sigma_\mu^2 + t \cdot \sigma_0^2}{\sigma_\mu^2 + (t+1) \cdot \sigma_0^2} \cdot \left(\frac{\sigma_\mu^2 + (t-1) \cdot \sigma_0^2}{\sigma_\mu^2 + t \cdot \sigma_0^2} \cdot \log(w_{i,t-2}) + \frac{\sigma_0^2}{\sigma_\mu^2 + t \cdot \sigma_0^2} \cdot \big(\log(w_i) + \mu_{i,t-1}\big) \right) \ldots \\
&\ldots + \frac{\sigma_0^2}{\sigma_\mu^2 + (t+1) \cdot \sigma_0^2} \cdot \big(\log(w_i) + \mu_{i,t}\big) \\
&\vdots \\
&= \frac{\sigma_\mu^2 + \sigma_0^2}{\sigma_\mu^2 + (t+1) \cdot \sigma_0^2} \cdot \log(w_{i,0}) + \frac{t \cdot \sigma_0^2}{\sigma_\mu^2 + (t+1) \cdot \sigma_0^2} \cdot \log(w_i) + \frac{\sigma_0^2}{\sigma_\mu^2 + (t+1) \cdot \sigma_0^2} \cdot \sum_{k=0}^{t} \mu_{i,t-k}.
\end{aligned}$$

1.D.2 Implications

Taking expectations of equation (1.18) with respect to signals yields

$$\mathbb{E}\big[\log(w_{i,t})\big] = \frac{\sigma_\mu^2 + \sigma_0^2}{\sigma_\mu^2 + (t+1) \cdot \sigma_0^2} \cdot \log(w_{i,0}) + \frac{t \cdot \sigma_0^2}{\sigma_\mu^2 + (t+1) \cdot \sigma_0^2} \cdot \log(w_i).$$

Note that this can be rewritten as

$$\mathbb{E}\big[\log(w_{i,t})\big] = \frac{\lambda}{\lambda + t} \cdot \log(w_{i,0}) + \frac{t}{\lambda + t} \cdot \log(w_i).$$

where $\lambda = \frac{\sigma_\mu^2 + \sigma_0^2}{\sigma_0^2}$. Thus, the speed of learning parameter λ reflects how noisy the signal is relative to initial aggregate uncertainty. Further, this analysis reveals that λ should be bounded from below by one. In practice, I find that λ is estimated to be close to one, which implies that the signals gig workers receive are precise relative to the variance in initial misperceptions.

Intuitively, the way that learning is modelled in equation (1.2) is similar to estimating a full model of Bayesian learning, where the agents' signals are simulated. This is because, if a sufficient number of agents were simulated, then averaging over agents with the same initial misperception would lead to (1.2). However, the model does deviate from Bayesian learning because agents perceive that they have an infinitely precise signal the value of gig work, rather than a noisy posterior, which may affect their behavior.

1.E Model Derivations and Extensions

This section provides some additional derivations for analysis conducted in the paper, and for extensions of the model presented in appendix 1.F.

1.E.1 Proposition 22 Hours Thresholds

The thresholds for workers' labor supply rule with hours-qualified benefits are given by

$$\rho_j = \frac{j^{1/\varepsilon} \cdot (p(\omega)+\kappa)}{\hat{\theta}_{i,t}} \text{ for } j = 15, 25,$$

$$\bar{\bar{\rho}}^{\varepsilon} \cdot \left(\frac{1}{\varepsilon-1} \cdot \frac{\hat{\theta}^{\varepsilon}_{i,t}}{(p(\omega)+\kappa)^{\varepsilon-1}}\right) - \bar{\bar{\rho}} \cdot \left(\hat{\theta}_{i,t} \cdot \frac{25^{1-1/\varepsilon}}{1-1/\varepsilon}\right) + 200 - 400 + (p(\omega)+\kappa) \cdot 25 = 0,$$

$$\bar{\rho}^{\varepsilon} \cdot \left(\frac{1}{\varepsilon-1} \cdot \frac{\hat{\theta}^{\varepsilon}_{i,t}}{(p(\omega)+\kappa)^{\varepsilon-1}}\right) - \bar{\rho} \cdot \left(\hat{\theta}_{i,t} \cdot \frac{15^{1-1/\varepsilon}}{1-1/\varepsilon}\right) - 200 + \big(p(\omega)+\kappa\big) \cdot 15 = 0.$$

And to find $\underline{\rho}$

$$\begin{aligned}
(1-\psi)\cdot\nu \le u(15,\omega;\rho_{15}) &\implies (1-\psi)\cdot\nu = u(h^*(\omega),\omega;\underline{\rho}) \\
u(15,\omega;\rho_{15}) < (1-\psi)\cdot\nu \le u(15,\omega;\rho_{15}) + 200 &\implies \underline{\rho} = \rho_{15} \\
u(15,\omega;\rho_{15}) + 200 < (1-\psi)\cdot\nu \le u(25,\omega;\rho_{25}) + 200 &\implies (1-\psi)\cdot\nu = u(h^*(\omega),\omega;\underline{\rho}) + 200 \\
u(25,\omega;\rho_{25}) + 200 < (1-\psi)\cdot\nu \le u(25,\omega;\rho_{25}) + 400 &\implies \underline{\rho} = \rho_{25} \\
u(25,\omega;\rho_{25}) + 400 < (1-\psi)\cdot\nu &\implies (1-\psi)\cdot\nu = u(h^*(\omega),\omega;\underline{\rho}) + 400.
\end{aligned}$$

1.E.2 Switching

To introduce endogenous switching to the model, I allow workers to switch policy, as well as exiting, when their perceived valuation changes. The endogenous exit rule is now altered such that a gig worker exits if and only if

$$\tilde{V}_i\big(\omega^*_i;\hat{\theta}_{i,t}\big) \le \frac{\nu_i}{1-\beta} \text{ and } \tilde{V}_i\big(\Omega/\omega^*_i;\hat{\theta}_{i,t}\big) - \tau \le \frac{\nu_i}{1-\beta}, \tag{1.19}$$

where τ is a new parameter: the hassle cost of switching policy. Therefore, workers will switch policy if and only if

$$\tilde{V}_i\big(\Omega/\omega_i^*;\hat{\theta}_{i,t}\big)-\tau>\tilde{V}_i\big(\omega_i^*;\hat{\theta}_{i,t}\big)\ \text{ and }\ \tilde{V}_i\big(\Omega/\omega_i^*;\hat{\theta}_{i,t}\big)-\tau>\frac{\nu_i}{1-\beta}. \tag{1.20}$$

If neither of these conditions are satisfied, then policyholders will remain on their current policy.

1.E.3 Cholesky Decomposition of Variance-Covariance Matrix

$$\begin{pmatrix}\sigma_\theta^2 & \sigma_{\theta,\phi} & \sigma_{\theta,\nu}\\ \sigma_{\phi,\theta} & \sigma_\phi^2 & \sigma_{\phi,\nu}\\ \sigma_{\nu,\theta} & \sigma_{\nu,\phi} & \sigma_\nu^2\end{pmatrix}=\begin{pmatrix}l_{1,1} & 0 & 0\\ l_{2,1} & l_{2,2} & 0\\ l_{3,1} & l_{3,2} & l_{3,3}\end{pmatrix}\times\begin{pmatrix}l_{1,1} & l_{2,1} & l_{3,1}\\ 0 & l_{2,2} & l_{3,2}\\ 0 & 0 & l_{3,3}\end{pmatrix}$$

$$\begin{aligned}\sigma_\theta^2 &= l_{1,1}^2\\ \sigma_\phi^2 &= l_{2,1}^2+l_{2,2}^2\\ \sigma_\nu^2 &= l_{3,1}^2+l_{3,2}^2+l_{3,3}^2\\ \sigma_{\theta,\phi} &= l_{1,1}\cdot l_{2,1}\\ \sigma_{\theta,\nu} &= l_{1,1}\cdot l_{3,1}\\ \sigma_{\phi,\nu} &= l_{2,1}\cdot l_{3,1}+l_{2,2}\cdot l_{3,2},\end{aligned}$$

where $l_{2,1}$ and $l_{3,2}$ is set equal to zero in order to ensure $\sigma_{\phi,\theta}=0$ and $\sigma_{\phi,\nu}=0$.

1.F Structural Model Extensions (***In Progress***)

Figure 1.F.1 shows that, for the small minority of gig workers who opt for the annual policy, their monthly intensive margin labor supply is roughly the same as fixed policy workers.

1.F.1 Variance in κ

In this subsection, I re-estimate the model allowing for independent variation in the linear cost parameter κ. Tables 1.F.1, 1.F.2, and 1.F.3 show the resulting parameter estimates, population characteristics, and model fit. The model finds allowing for a stan-

Figure 1.F.1: Hours Worked Comparison with Annual Policyholders

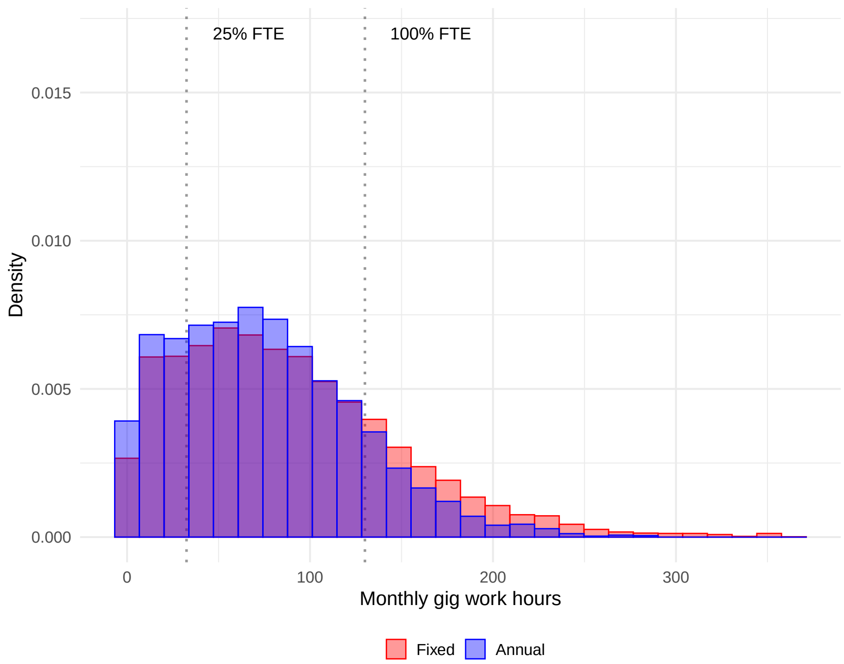

Notes: This figure plots the histogram of monthly hours worked for fixed customers in blue, and the analogue for annual customers in red.

dard deviation of £1.79 helps fit the empirical moments.

Table 1.F.1: Parameter Estimates

$\zeta_{1:6}$	$\hat{\zeta}_{1:6}$	$\zeta_{6:12}$	$\hat{\zeta}_{6:12}$
μ_θ	3.69	μ_ν	12.66
	(0.05)		(0.33)
σ^2_θ	0.31	σ^2_ν	4.53
	(0.02)		(0.49)
$\sigma_{\theta,\nu}$	-1.09	σ^2_ϕ	0.14
	(0.05)		(0.01)
λ	1.00	κ	$-$
	(0.09)		$(-)$
σ^2_ρ	0.03	ψ	0.27
	(0.00)		(0.00)
μ_κ	3.67	σ^2_κ	0.00
	(0.02)		(0.00)
μ_P	101.82	σ_P	24.08
p	0.94	β	$0.95^{1/12}$
η	0.07	ε	$\Delta \widehat{\log}(h) \cdot \log(1 + p/\hat{\kappa})$

Notes: The top panel of this table presents estimates of the structural parameters from the model. Standard errors are contained in the parentheses, which are estimated as described in subsection 1.5.1. The second panel shows the variable policy premium and the mean of the quoted fixed premium distribution used in the estimation, which correspond to the empirical averages of these parameters, as well as the set discount factor.

Table 1.F.2: Simulated Population Characteristics

Statistic	Population	Participants
Mean valuation θ	46.50	210.86
SD valuation θ	27.93	48.65
Mean misperception ϕ	1.00	1.04
SD misperception ϕ	0.39	0.11
Mean outside option ν	9,745	672
SD outside option ν	1,255	487
Mean linear cost κ	39	39
SD linear cost κ	2	2
Correlation $\rho_{\theta,\nu}$	-0.62	-0.31
Correlation $\rho_{\theta,\phi}$	0.00	-0.27
Correlation $\rho_{\nu,\phi}$	0.00	0.52
Correlation $\rho_{\theta,\kappa}$	0.00	-0.01
Correlation $\rho_{\kappa,\phi}$	0.00	0.01
Correlation $\rho_{\nu,\kappa}$	0.00	0.04

Notes: This table presents statistics that describe the simulated population. The first column shows these statistics for the entire population, while the second column conditions on participating in the gig economy.

Table 1.F.3: Model Fit

Moment	Data	Model
Labor market share (%)	3.6	3.7
Variable policy share (%)	66.8	69.0
Mean hours per month, variable policy	53.4	57.7
SD hours per month, variable policy	48.4	37.5
Mean hours per month, fixed policy	95.3	105.5
SD hours per month, fixed policy	68.6	76.1
Share non-cost-minimizing (%), variable policy	11.6	9.2
Share non-cost-minimizing (%), fixed policy	60.6	61.7
Mean hours per month from cost-minimizing, variable policy	40.1	41.2
SD hours per month from cost-minimizing, variable policy	36.6	48.1
Mean hours per month from cost-minimizing, fixed policy	55.6	40.3
SD hours per month from cost-minimizing, fixed policy	41.7	30.8
Hazard rate for cost-minimizers (%), variable policy	28.9	23.5
Hazard rate for cost-minimizers (%), fixed policy	20.7	17.4
Hazard rate for non-cost-minimizers (%), fixed policy	29.2	38.0
Decline in hours per month for non-cost-minimizers, variable policy	-10.2	-9.3
Decline in hours per month for non-cost-minimizers, fixed policy	2.6	2.7
Mean quoted fixed premium (£), variable policy	97.0	95.3
Mean quoted fixed premium (£), fixed policy	106.5	104.4
SD quoted fixed premium (£), variable policy	25.1	20.1
SD quoted fixed premium (£), fixed policy	24.3	26.6
Share of zero hours months (%)	5.5	5.8
Mean within worker SD hours per month	31.0	37.6

Notes: This table presents the targeted empirical moments alongside their model implied counterparts. The first column contains the empirical moments and the second column contains the model analogue.

Chapter 2

The Cost of Labor Supply Biases

2.1 Introduction

Self-employment is an important form of work around the world. Across OECD countries, the median self-employment rate is 15 percent and the rise of the gig economy has precipitated an increase in the number of people who engage in self-employment as a source of extra income (Collins et al., 2019; Katz and Krueger, 2019). Furthermore, most workers in low- and middle-income countries are self-employed (Fields, 2013).

A key feature of self-employment is "being your own boss" (Hamilton, 2000). As their own boss, the self-employed must organize their work but they may not necessarily do this efficiently. In this paper, I examine how self-employed workers manage a key input: their labor. Specifically, I provide evidence that intensive margin labor supply decisions deviate from an optimal benchmark and, most importantly, I develop a new theoretical framework to quantify the welfare effects of these deviations.

The results in this paper point towards an important aspect of the typical flexibility versus security trade-off used to frame self-employment. Namely, behavioral frictions that prevent workers from fully exploiting flexibility. For a group of self-employed workers who pick their own hours and face regular variation in their return to work, I find economically significant welfare losses of up to six percent of income which stem from suboptimal daily labor supply.

Intensive margin labor supply responses to wage changes are the focus of a vast literature in economics (Blundell and MaCurdy, 1999). The results from this literature are often used to calibrate parameters in models that help to inform normative topics, such as the efficiency of income taxes (Keane, 2011). This rests on a revealed preference logic. Individuals are willing to work ε percent more when the wage increases by 1 percent, so the ordinal relationship between utility and work is governed by parameters that are a function of ε.

Yet, in some settings, labor supply responses appear inconsistent with even the most unrestrictive of models. New York cabdrivers, for example, have been shown to exhibit negative daily Frisch elasticities (Camerer et al., 1997b; Doran, 2014; Schmidt, 2018), positive daily Frisch elasticities (Farber, 2005, 2008, 2015), and to be influenced by hourly income effects (Morgul and Ozbay, 2015; Thakral and Tô, 2017), as well as to demonstrate behavior consistent with hours targeting (Crawford and Meng, 2011). Moreover, these findings replicate in experimental and observational studies in other contexts (Agarwal et al., 2013a; Andersen et al., 2014; Chang and Gross, 2014; Chou, 2000; Dupas et al., 2015; Fehr and Goette, 2007; Nguyen and Leung, 2009).

Such results raise two key questions: Are they a consequence of behavioral biases, or are they benign deviations from normative assumptions made by researchers? And if they are the product of behavioral biases, how do we distinguish between preferences and behavior so as to infer the cost of biases? Similar questions are pervasive in behavioral economics (Allcott and Taubinsky, 2015; Bernheim and Rangel, 2009; Bernheim and Taubinsky, 2018; Goldin and Reck, 2017; Mullainathan et al., 2011; Nielsen and Rehbeck, 2022), but they remain unanswered in the labor supply context despite significant positive analysis.

This paper proposes and evinces a simple story about labor supply behavior in order to shed light on these questions. I argue that individuals vary daily labor supply substantially and in line with preferences when they face certain types of wage variation. In particular, salient and unusual variation that negates biases. However, workers rely on suboptimal labor supply rules otherwise. Therefore, labor supply responses to specific sources of wage variation can be used to infer preferences over consumption and leisure. But general behavior will deviate from optimal behavior—it will be biased.

I make three contributions to establish this narrative and to evaluate its normative impact using data on a set of self-employed taxi drivers in London. Firstly, I distinguish between different types of wage variation to separately infer preferences and behavior in the spirit of Chetty-Looney-Kroft (2009). London Tube strikes provide a source of wage variation that is both salient and special owing to its rarity and substantial coverage in the media. I posit that these qualities have the potential to override a wide range of biases. For example, media attention likely relaxes information constraints and the one-off nature of Tube strikes may alleviate self-control problems through framing effects. In the language of Allcott and Taubinsky (2015), I require Tube strikes to be a "pure nudge" in that they debias drivers without any confounding labor supply effects.

Conversely, exogenous wage variation stemming from the quasi-random allocation of jobs to drivers comprises a less notable feature of workers' environment. In particular, owing to institutional quirks, persistent differences in the distance of jobs given to drivers over the course of a shift lead to fluctuations in driver wages. Unlike Tube strikes, this source of variation is ever-present and inconspicuous and, as such, invokes standard labor supply heuristics.

Absent the factors discussed above, the labor supply responses to both types of changes in the wage rate should constitute the same Frisch elasticity because of their temporary nature. Instead, drivers exhibit strikingly different behavioral responses to these sources of wage variation. Tube strikes imply a large Frisch elasticity (0.80, s.e. 0.10), while wage fluctuations driven by mundane features of the institutional setting suggest a much smaller elasticity (0.12, s.e. $<$0.01). The Tube strike elasticity estimate is in line with Frisch elasticities from the labor supply literature in contexts where biases are not a concern and workers pick their hours (Angrist et al., 2021; Caldwell and Oehlsen, 2021). In contrast, the other elasticity estimate is at the very low end of Hicksian labor supply elasticity estimates which, theoretically, bound the Frisch elasticity from below (Keane, 2011).

Secondly, I confirm the smaller labor supply elasticity is not consistent with preferences under the mildest of neoclassical assumptions. A control function estimator, which also leverages the quasi-random allocation of jobs to identify exogenous wage variation, reveals a portion of drivers' labor supply function to be downward sloping in a way that is reminiscent of income targeting; hours increase with high and low wages but fall for wages around the mean. A twice non-monotonic labor supply function cannot be reconciled with unrestrictive neoclassical (and many other) models but, still, one may suspect that a downward-sloping labor supply curve is driven by income effects.

To rule out income effects, I exploit a permanent fare reform to estimate a Marshallian elasticity. This elasticity bounds all neoclassical labor supply responses from below since it incorporates the income effect of price changes. I estimate a Marshallian elasticity of -0.14 (s.e. 0.04), which is too small in magnitude to explain the downward-sloping daily labor supply from the control function estimator. Comfortingly, this result aligns closely with Ashenfelter et al. (2010), which follows a similar approach.

As further support for biased decision-making, I present survey evidence on how drivers determine their labor supply. Drivers readily subscribe to income and hours targeting but do not recognize the implications this has for the relationship between the wage rate and their hours. This is consistent with the results from the control function estimator, which is indicative of target-based behavior, as well as a cognitive dissonance between drivers' general behavior and their true preferences.

Thirdly, I derive a behavioral welfare expression (BWE) that approximates losses due to suboptimal behavior in continuous decision-making settings without taking a strong stance on the structure of workers' utility functions. The BWE reveals the average marginal bias (Allcott and Taubinsky, 2015) of suboptimal labor supply in terms of the wedge in the intratemporal optimality condition between the wage and the marginal rate of substitution between consumption and leisure (MRS). Crucially, the expression involves only the difference between workers' biased and optimal labor supply, and a small number of estimable and familiar sufficient statistics, which describe worker preferences.

To estimate the BWE, I use the Tube strikes Frisch elasticity and the permanent fare reform Marshallian elasticity to inform preferences. This treats labor supply during Tube strikes as an optimal benchmark and rests on the assumption that biased behavior satisfies the intratemporal optimality condition in expectation. Intuitively, this is necessary because these elasticities compare behavior in general, which is biased, with efficient behavior on particular occasions. The condition is also appealing since it grants a degree of sophistication to drivers and, thus, does not mechanically overstate welfare losses due to biases.

Elasticities describe the responsiveness of hours around some level, so I consider two ways to calibrate the optimal level of hours: first, to match the observed level of hours during Tube strikes and, second, to keep income constant under both the biased and optimal regimes. Quantitatively, these approaches yield similar results.

Lastly, welfare losses due to biases depend on the level of the wage because the difference between observed and efficient labor supply varies with the wage rate. Therefore, I calculate an expected welfare loss by averaging over losses using the empirical distribution of observed wage rates. This can be viewed as a partial equilibrium assumption since it assumes that, in a counterfactual with drivers behaving optimally, the market price will not change.

Bringing the BWE to the data reveals expected daily welfare losses that range from two to six percent of daily income (or, equivalently, £2.09 to £5.29) under a variety of assumptions about the optimal level of hours and the extent of biases in the population. When these daily losses are accumulated over the course of a standard driver-year they point towards losses of up to £1,000 for biased drivers. This is comparable to annual welfare losses found in other settings, for example, health insurance choices (Handel and Kolstad, 2015), which are considered large. Therefore, this paper suggests an economically significant impact of intensive margin labor supply biases on the welfare of self-employed workers.

Introducing heterogeneity over the extent of biases in the population dramatically raises welfare losses due to biased labor supply. Survey evidence and lower bounds on the magnitude of biased labor supply elasticities suggest that around one-third of drivers supply labor optimally. If this is the case, then welfare losses more than double relative to a scenario with uniform biases. Fewer biased drivers must explain the aggregate deviations of observed labor supply from the optimal benchmark, which implies more acutely biased labor supply. This increase in the severity of biases outweighs any reduction in the number of workers suffering biases. This is a local result but points to another mechanism by which heterogeneity in biases can exacerbate welfare losses (Taubinsky and Rees-Jones, 2018).

The composition of welfare losses suggests a limited role for self-control issues as the direct source of biases. When the level of optimal hours is set to match those on Tube strikes, rather than to keep income constant, then workers' average income falls under optimal labor supply. Since reductions in average income likely translate with a lag to reduced consumption, while larger reductions in the disutility of work are immediate, this pattern of losses is not consistent with workers suffering from a lack of self-control. Indeed, labor supply from the control function estimator is indicative of workers experiencing reference dependence instead of self-control problems. However, a lack of self-control, which can motivate goal setting and that manifests as reference dependence, may still be the ultimate source of biases (Hsiaw, 2013).

I conclude the paper by discussing the implications of these results for workers, organizations that contract with these workers, and government policy. For self-employed workers, there is a clear lesson: hours targeting dominates income targeting. While hours targeting imposes losses relative to an optimal benchmark, it does not introduce a costly covariance between the wedge in individuals' intratemporal optimality condition and the hours that they work, which is a feature of income targeting. This advice holds under all but the most extreme demands for income insurance.

Contracting companies have many potentially fruitful levers that they could pull to aid workers' optimization, but it is unclear whether they would prefer to ameliorate or exploit worker biases. Examples include the provision of training around the ramifications of target-based labor supply rules, placing greater emphasis on the going wage rather than earned income, as well as reductions in the volatility of wages which necessitates smaller labor supply adjustments.

A conflict of interests between firms and workers in terms of tackling labor supply biases may justify government intervention. Policymakers could, for example, mandate the provision of information to self-employed workers to prevent firms from withholding details that help workers override their biases. Moreover, if self-employment entails inefficiencies due to labor supply biases, then policies which either purposefully or unintentionally encourage self-employment should be scaled back, *ceteris paribus*.

Literature review. This research is connected to several literatures. Most related is part of the behavioral welfare economics literature that assesses the normative implications of policy interventions, such as commodity taxes and nudges (Thaler and Sunstein, 2009), via the estimation of reduced form statistics. This approach is explained in Mullainathan et al. (2011) and exemplified by Allcott and Taubinsky (2015), Berkouwer and Dean (2019), Chetty et al. (2009), and Spinnewijn (2015), among others.

Similarly, the BWE derived below is estimable with a few sufficient statistics and identification of biased behavior. Yet, it differs by directly considering a behavior change over a continuous choice variable, rather than a price-induced behavior change for a binary variable.[1] Moreover, the results in this paper suggest that welfare losses can increase on aggregate when a (locally) smaller proportion of individuals are biased, which provides another example of how heterogeneity in biases can exacerbate welfare losses (Taubinsky and Rees-Jones, 2018).

The empirical aspects of this paper lean on the many studies of income targeting amongst New York cabdrivers, for example, Camerer et al. (1997b), Crawford and Meng (2011) , Farber (2005, 2008, 2015), and (Thakral and Tô, 2017), as well as the broader labor supply literature (Blundell and MaCurdy, 1999) and methodological work on control function estimators (Wooldridge, 2015). Crucially, in this paper, I move beyond a positive analysis of labor supply and the role of income targeting in order to understand the welfare impacts of behaviorally biased labor supply.

Finally, this paper touches upon recent work that studies alternative work arrangements. Two important papers in this literature are Mas and Pallais (2017) and Chen et al. (2017),[2] which focus on the value of flexible work. These papers use stated and revealed preference approaches, respectively, to find different results; individuals are not willing to pay much for flexibility but their behavior indicates that they value it greatly. Further, recent work by Lachowska et al. (2023) evinces welfare losses in more traditional forms of employment that are less flexible and prevent workers from trading off leisure and income optimally. This paper challenges the premise that workers are able to perfectly trade off wages and the disutility of work when they are in control of their hours.

The paper proceeds as follows. Section 2.2 discusses the institutional details of the empirical setting, and the data available. This is followed by an analysis of drivers' behavioral daily labor supply in section 2.3. In section 2.4, I argue and evince that these behavioral tendencies are not normative and generate welfare losses. In line with this argument, section 2.5 develops a theoretical apparatus to calculate this welfare loss and estimates the cost of labor supply biases. Section 2.7 concludes.

[1] Though I also derive an analogous BWE that considers price changes in appendix 2.D.

[2] See also Chen et al. (2020) who use a similar approach to Chen et al. (2017) but allow for fixed costs of working, and correlation between wages and the cost of working.

2.2 Institutional Details and Data

This section discusses institutional details and the data available. I also define the analysis sample and present some summary statistics to provide an introduction to the data.

Institutional Details. This paper uses data from a private hire taxi firm in London. The firm leases cars to its self-employed drivers,[3] who are allocated jobs to complete via an application on their mobile phone. These jobs are demanded by customers either on another application or over the phone. The car can be used for leisure purposes but may not be used for other professional reasons, such as serving other ride hailing businesses, so there is little risk of conflating intensive margin labor supply responses with switching between other work (Caldwell and Oehlsen, 2018).

Jobs are allocated by a central computer system; the system ranks a number of the closest cars according to how suitable they are for the job. The suitability of a car is determined by a number of factors which include the size of the car and whether the car is currently occupied. The top ranked car is then allocated to the job. While drivers can *de jure* turn down jobs, this is rare because they are disadvantaged in future job allocations if they do so. Moreover, the final destination of a job is not visible to drivers and this would likely be a key determinant of whether a driver would like to accept a job, or not. Nonetheless, drivers are able to determine when they finish work; at any point in time they can tell the firm that they are "Going Home" through the application. After this notification, the driver will receive no further jobs until they next log on or, if they are in the middle of a job, they will receive no further jobs after the completion of the current job.

Drivers are paid by the job and receive between 50 and 80 percent of the amount the customer pays. The amount received for a job is primarily determined by the distance of a job;[4] there is no separate compensation for the duration of a job. The amount also varies with the type of job,[5] the number of passengers, the number of stops, the time of

[3]The leases are normally for a 12 week period and, over the span of this data, the average cost of a lease is around £200.00 per week. The firm provides incentives for drivers, which are similar to Uber's Driver Rewards, that lead to discounts on leases. Unfortunately, these have not been well documented and I have not been able to exploit this variation. The cost of a lease includes various maintenance costs and insurance for the driver, though drivers are responsible for fuel costs. As is common in the literature, incomes and wages in this paper are gross of these costs since they are not observed.

[4]A linear regression of the driver's fare on a 5th order polynomial of distance interacted with a dummy for the fare schedule yields an R-squared of 0.85.

[5]The type of job is affected by the type of car requested and supplied (e.g. Ford Galaxy or Toyota Prius)

day, and the location of the job. For some jobs, a value-added tax of 20 percent of the total transaction value is payable; drivers must pay their share of this from the fare they receive, which I account for in the analysis below. Drivers' earnings are paid out on a weekly basis by the firm.

Data. I observe job-level data from January 2012 to December 2019, which is electronically recorded from drivers' phone applications. Thus, each observation is a job and has variables for the driver ID, the total transaction value, the driver's fare, the start and finish time of the job, the start and finish location, and the distance of the job, among others. The raw data contains over 60 million jobs, which I cleaned analogously to Haggag and Paci (2014). This process removes rides if they contain anomalous variables, or are missing key variables, or are cancelled.[6] After cleaning, around 50 million jobs remain. From the clean job-level data, I constructed shift-level data by allocating jobs to the same shift, if they are completed by the same driver and are within six hours of one another (Farber, 2015; Thakral and Tô, 2017). If there is a break of more than one hour between two jobs, I deduct the excess duration of the break from the length of the shift in order to get a more accurate picture of actual labor supply. The shift wage is then calculated as the total income earned over a shift divided by the shift length. Finally, I cleaned these shifts in an analogous way to jobs, which leaves just under seven million shifts.

Analysis Sample. This paper focuses on intensive margin decisions made by drivers, so I use a sub-sample of drivers for whom I observe sufficient wage variation in their shifts. Specifically, I ensure that I observe a driver during at least one Tube strike and for 50 shifts either side of a fare reform that I use to estimate a Marshallian labor supply elasticity. This leaves me with around 3.5 million shifts driven by 2,600 drivers; I refer to this as the *balanced* sample. All analysis is conducted on this sample, unless otherwise stated. I also construct a *robust* sample, which imposes the same restrictions as the balanced sample but requires drivers to be observed during a minimum of four Tube strikes in order to check whether my selection along the extensive margin has important

and the customer (e.g. corporate client or individual customer).

[6]The cleaning of job- and shift-level data is explained further in appendix 2.A.

Table 2.2.1: Job Summary Statistics

Statistic	Mean	St. Dev.	Pctl(25)	Pctl(75)
Job time (minutes)	33.41	15.60	21.00	42.00
Job distance (kilometers)	10.69	8.85	4.53	13.69
Driver fare (£)	14.89	9.34	8.00	18.67

Notes: This table presents summary statistics of job-level variables from the balanced sample, which requires that drivers are observed during tube strikes and sufficiently often on either side of the fare reform.

consequences for my results.[7] While extensive margins are certainly of interest more generally, this empirical setting is not best suited to investigate such questions because drivers in the sample essentially work full time.

Summary Statistics. I report summary statistics at the job and shift level. Table 2.2.1 shows job summary statistics from the balanced sample. The mean ride duration is just over half an hour, and 50 percent of all rides in my sample take between 21 and 42 minutes. In general, the rides are longer than those in the New York taxi context. The average job distance is 10.69 kilometers, and a driver receives £14.89 on average for a job.

Table 2.2.2 reports shift summary statistics. On average, a shift is comprised of just over six jobs, which translates to a working shift length of six and a half hours. If breaks exceeding an hour are included, this rises to eight hours. Shift income averages £91.62. The hourly wage, which is constructed as shift income divided by shift length, averages £14.32 with a standard deviation of £3.80. Table 2.2.3 shows sample sizes, and that shift variable means do not vary much between the full, balanced, and robust samples.

An examination of shift variable distributions is also informative of drivers' behaviors. Figures 2.2.1 and 2.2.2 illustrate the distribution of shift start and end times, respectively. The figures suggest a driver's typical day begins early in the morning and ends in the mid-afternoon. It is also common for drivers to begin shifts after lunchtime in anticipation of the afternoon rush, and to end shifts around midnight. These different shift patterns lead to a distribution of shift lengths that are displayed in figure 2.2.3; approximately half of all shifts fall between five and ten hours in length.

[7] I do not have any driver characteristics to compare my sub-sample with the full sample of drivers, but in table 2.2.3 I report shift variables means to show that they are very similar.

Table 2.2.2: Shift Summary Statistics

Statistic	Mean	St. Dev.	Pctl(25)	Pctl(75)
Number of jobs	6.15	2.33	4	8
Shift length w/o breaks (hours)	6.56	2.41	4.64	8.27
Shift length w/ breaks (hours)	8.14	3.08	5.79	10.31
Shift income (£)	91.62	35.91	64.95	114.35
Shift wage (£/hour)	14.32	3.80	11.58	16.61

Notes: This table presents summary statistics of shift-level variables from the balanced sample, which requires that drivers are observed during tube strikes and sufficiently often on either side of the fare reform.

Table 2.2.3: Comparison of Samples' Size and Shift Means

Sample	Full	Balanced	Robust
Sample size	6,871,701	3,627,711	3,154,271
Number of jobs	6.12	6.15	6.18
Shift length w/o breaks (hours)	6.47	6.56	6.59
Shift length w/ breaks (hours)	8.04	8.14	8.16
Shift income (£)	89.3	91.62	91.24
Shift wage (£/hour)	14.14	14.32	14.21

Notes: This table compares means of shift-level variables from different samples: the full sample, the balanced sample, and the robust sample, as defined in section 2.2.

Figure 2.2.1: Shift Start Time Density

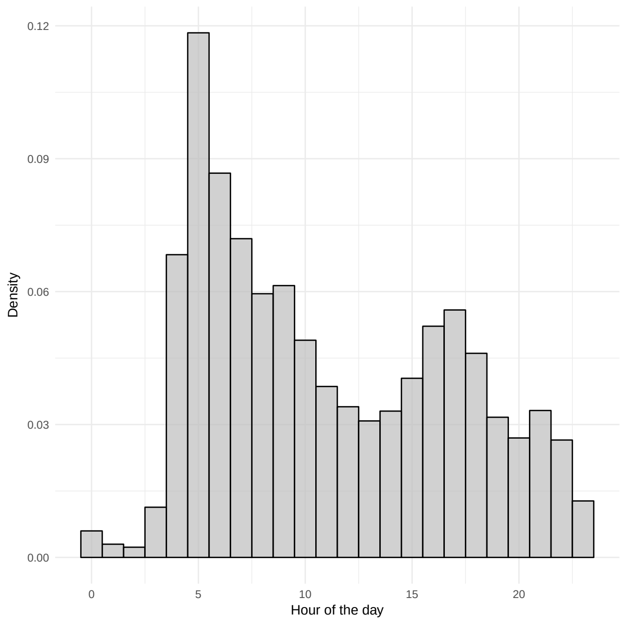

Notes: This figure presents the empirical density histogram of the start times of shifts for the balanced sample, where each bin is an hour of the day.

Figure 2.2.2: Shift End Time Density

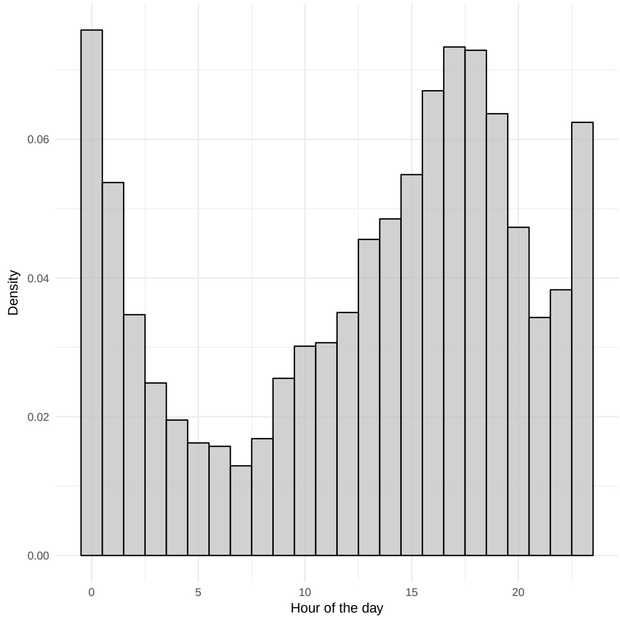

Notes: This figure presents the empirical density histogram of the end times of shifts for the balanced sample, where each bin is an hour of the day.

Figure 2.2.3: Shift Length Density

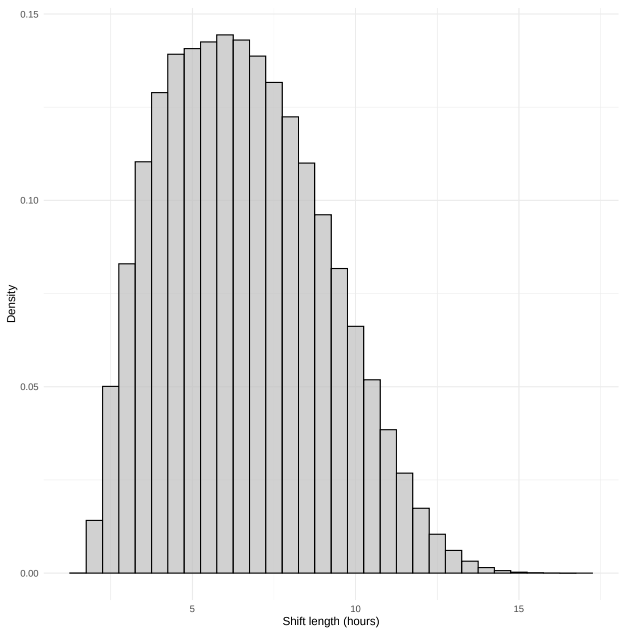

Notes: This figures presents the empirical density histogram of shift length for the balanced sample, where each bin is 30 minutes.

2.3 Empirical Motivation

In this section, I focus on intensive margin labor supply responses to daily wage changes via the estimation of two Frisch elasticities and a labor supply function. There are two primary reasons to focus on the daily level. First, wage rates *between* shifts are not correlated in an economically meaningful way,[8] so labor supply should be determined daily with reference to the wage relative to its usual level. Secondly, preempting the results below, it is hard to reconcile a downward sloping labor supply function with any other time horizon, as argued by Camerer et al. (1997b).

I leverage two variables that induce wage variation which is temporary and exogenous: London Tube strikes and variation in the mean distance of jobs within a shift.[9] While these sources of variation induce theoretically equivalent variation, Tube strikes are seldom events whereas the mean distance of jobs varies regularly.

With this exogenous variation, I use different empirical frameworks to document contradictory labor supply responses to wage fluctuations, all in one setting. Firstly, I use Tube strikes in order to estimate a large and positive daily Frisch elasticity (0.80, s.e. 0.10). I refer to this as the *True* Frisch elasticity because, later in the paper, this statistic will be used to determine preferences. Secondly, I find a much smaller Frisch elasticity (0.12, s.e. <0.01) when I exploit variation in mean job distances within a shift. I refer to this latter estimate as the *Behavioral* Frisch elasticity because, as I will argue later, it reflects behavioral biases. Thirdly, I use a control function model to fully trace out the shape of the labor supply function, which reveals a negatively sloped portion of the labor supply curve. Then, I discuss and reject issues that could explain these different results and, in doing so, present estimates of strong hourly income effects from a probability stopping model (Thakral and Tô, 2017).

2.3.1 True Frisch Elasticity

London Tube strikes serve as a natural experiment to estimate a Frisch elasticity; they cause a significant change in the average wage, and are neither long enough nor severe enough to affect the marginal utility of income. Data on the dates and types of Tube strikes have been provided online by TfL thanks to a Freedom of Information request, and I verified these dates by checking coverage of media outlets at the time. There are two types of strikes: network wide and line specific. A network wide strike affects the capacity of all lines in the Tube network, while line specific strikes affect the running of

[8] A £1 increase in today's shift wage increases the next shift's wage by £0.02 (s.e. <0.01) on average.

[9] As a result of the temporary nature of this variation, the marginal utility of income is likely fixed, which is required to estimate a Frisch elasticity.

Table 2.3.1: True Frisch Estimation First Stage

	Dependent variable:
	log(Shift wage)
Network strike 04/02-06/02	0.007* (0.004)
Network strike 28/04-30/04	0.055*** (0.004)
Network strike 09/05-10/05	0.010* (0.005)
Network strike 13/06-14/06	−0.002 (0.005)
Network strike 01/07-09/07	0.033*** (0.003)
Network strike 09/07-15/07	0.014*** (0.003)
Line strike 22/08	0.001 (0.006)
Line strike 01/12-02/12	−0.005 (0.006)
First stage F-statistic	29.99
Observations	654,045
R^2	0.046
	*p<0.1; **p<0.05; ***p<0.01

Notes: This table presents coefficients on different London Tube strikes from the regression of driver wages on this variables, as well as driver fixed effects, a factor variable for start time of the shift, Ramadan, bank holidays, days of the week, months, and a 6th order polynomial time trend. The regression uses data from 01/01/2014 to 01/05/2015. Standard errors are clustered at the driver level.

only one or two lines. I restrict my analysis in this subsection to the period 1st January 2014 to 1st May 2015, which spans one fare schedule. This is for two interrelated reasons. First, this period contains the majority of network wide strikes, which causes the most systematic variation in the wage rate. Secondly, results are robust to time controls over this period and are also replicated for the full time horizon.

Given the data on Tube strikes, I investigate whether they are informative of wage rates at the driver-shift level. To do so, I regress the wage rate on a separate dummy for each strike, driver fixed effects, and controls for: start time of the shift, Ramadan,[10] bank holidays, days of the week, months and a flexible time trend. These are the main set of controls in the analysis. The results are shown in table 2.3.1. Five out of eight of the Tube strikes cause a statistically significant increase in the wage rate, and these are all network strikes. The largest effect was for the strike in April 2014, which caused a rise of five percent in the mean wage. These effects contribute to a first stage F-statistic of 29.99 although, for the results, the first and second stages are jointly estimated below. The main conclusion of this investigation is that Tube strikes significantly affect drivers wages.

[10] A large proportion of drivers are Muslim so controlling for Ramadan is logical though, in practice, it does not significantly alter results.

Therefore, the validity of my instrumental variable analysis rests on the exogeneity of Tube strikes, namely, that Tube strikes do not affect labor supply other than through their effect on the wage rate. Commonly cited reasons for Tube strikes include: insufficient pay, poor working conditions, and unfair dismissal of staff. There is no clear reason why the emergence of such concerns should be related to drivers' labor supply. However, the determinants of the specific dates of strikes may be different from their fundamental cause. It is often argued that the precise dates of strikes are set in order to cause maximum disruption.[11] This could imply that Tube strikes occur on dates when drivers are least able to expand their labor supply to meet the additional demand caused by a Tube strike. However, *ex post* this effect seems unlikely given the sizeable Frisch elasticity I estimate. Another concern is that Tube strikes cause major traffic. Obviously, traffic has implications for drivers' wages but a problem would only arise if drivers prefer more or less traffic aside from its impact on wages. I have no evidence of a systematic preference, but this concern is a valid caveat.

Before estimating a daily Frisch elasticity, I note that Tube strikes generally last two days. This is not problematic if drivers still make labor supply decisions one day at a time, which is likely given the labor supply function I estimate and the small sums of income at the daily level. However, if drivers do optimize over the course of a Tube strike, this could bias my daily Frisch elasticity estimate downwards. This is clear from a simple example: rather than driving a long shift during a Tube strike, a driver may decide to drive two normal shifts, which would cause me to estimate a zero response of hours to wages—despite a clear increase in labor supply. Again, this does not seem to be a significant problem given that my Frisch elasticity estimate is in the upper range of previously estimated elasticities.

Intuitively, identification rests on temporal variation within drivers; I compare the hours worked on a day with a Tube strike versus the hours worked on a day without a Tube strike, and their associated wages. In order to estimate the True Frisch elasticity, I implement a full information estimation of the structural equations (2.1) and (2.2) with driver fixed effects,[12] which is asymptotically equivalent to generalized 2SLS. The

[11] See, for example, https://www.telegraph.co.uk/news/2017/01/08/government-accuses-union-bosses-co-ordinating-transport-strikes/.

[12] Therefore, this regression identifies a positively weighted treatment effect for "changers", i.e. those drivers who experienced a wage change.

equations are as follows,

$$h_{i,s} = \alpha_i + \beta \cdot w_{i,s} + \Gamma \cdot X_{i,s} + u_{i,s}, \tag{2.1}$$

$$w_{i,s} = \delta_i + \Theta \cdot T_{i,s} + \Lambda \cdot X_{i,s} + v_{i,s}, \tag{2.2}$$

where subscripts i denotes a driver and s denotes the shifts, h is hours worked, w is the wage rate, T is a vector of length equal to the number of strikes with each element a dummy to indicate whether the shift took place during the respective strike, and X is a vector of controls that have been previously mentioned and are noted in the results table.

The estimates of β from equation (2.1) are shown in table 2.3.2, where standard errors are clustered at the driver level. The first column shows the OLS estimate, which is significantly negative because of endogeneity that is brought on for two reasons. Firstly, due to division bias because of the construction of the wage rate (Borjas, 1980) and, secondly, because demand and supply are conflated. The main estimate in the second column uses all strikes in 2014 and shows a significantly positive response of hours worked in a shift to the wage rate: a Frisch elasticity of 0.80 (s.e. 0.10). This estimate changes only marginally when restricting to network strikes or using the robust sample of drivers, which is shown in columns (3) and (4), respectively. When I use the full time horizon, which spans 2012 to 2019, I get a smaller point estimate as shown in column (5) but the 95 percent confidence intervals overlap considerably and the estimates are not statistically significantly different. The same is true when I vary the flexibility of the time trend, see columns (6) and (7), except this time the point estimates increase.

In summary, Tube strikes cause an increase in the wage rate, which significantly raises the length of drivers' shifts.

2.3.2 Behavioral Frisch Elasticity

Shifts in the composition of consumer demand act as a further instrument with which to estimate a second Frisch elasticity. Precisely, I use the mean distance of jobs in a driver's shift as an instrument for the wage rate. As discussed in section 2.2, a driver's fare for a job is primarily determined by the distance of the job; longer distance jobs have higher driver fares. Column (1) in table 2.3.3 shows that this effect maps to wages

Table 2.3.2: True Frisch Elasticity Results

	Dependent variable:						
	log(Shift length)						
	OLS	All strikes	Network only	Robust sample	All schedules	4th order poly.	8th order poly.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log(Shift wage)	-0.136^{***}	0.799^{***}	0.808^{***}	0.814^{***}	0.625^{***}	0.954^{***}	0.912^{***}
	(0.004)	(0.103)	(0.103)	(0.105)	(0.086)	(0.132)	(0.107)
Shift type dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ramadan dummy	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bank holiday dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DOTW dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Schedule dummies	No	No	No	No	Yes	No	No
Time trend poly. order	6th	6th	6th	6th	6th	4th	8th
Observations	654,519	654,519	654,519	613,941	3,466,056	654,519	654,519
R^2	0.114	0.013	0.013	0.012	0.013	0.008	0.009

$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Notes: This table presents coefficients on wages from a series of regressions of shift length on this variable, as well as driver fixed effects, a factor variables for start time of the shift, Ramadan, bank holidays, days of the week, months, fare schedule, and a time trend of varying flexibility. An indicator for the fare schedule is also included in column 5. The wage variables is instrumented for with London Tube strikes. The regressions uses data from various periods and drivers. Standard errors are clustered at the driver level.

Table 2.3.3: Behavioral Frisch Estimation First Stage

	Dependent variable:		
	log(Shift wage)		
	Mean job dist.	LOM wage	Both
	(1)	(2)	(3)
Mean job distance	0.028^{***} (0.0001)		0.028^{***} (0.0001)
Leave-out mean wage		0.059^{***} (0.0004)	0.047^{***} (0.0003)
First stage F-statistic	$1.659\times e^6$	$1.638\times e^4$	$8.304\times e^5$
Observations	3,466,056	3,466,055	3,466,055
R^2	0.552	0.290	0.571
			$^*p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Notes: This table presents coefficients on alternative instruments to the wage from a series of first stage regressions that control for driver fixed effects, a factor variable for start time of the shift, Ramadan, bank holidays, days of the week, months, fare schedule, and a 6th order polynomial time trend. The regression uses data from the full time horizon and balanced sample. Standard errors are clustered at the driver level.

for shifts, so that shifts with longer distance jobs have higher wages. Since the distance of jobs is determined by the pick-up and drop-off that customers demand, different mean job distances across shifts reflect differences in the composition of demand from customers.

This alternative instrument is valid if the mean distance of jobs in a shift does not affect the supply of hours aside from its effect on wages. I highlight two threats to this condition. First, longer jobs may systematically cause drivers more or less disutility; for example, longer jobs may be more fatiguing for drivers. There is no clear direction to this effect. Indeed, this instrument identifies a non-monotonic labor supply function in subsection 2.3.3, which is at odds with a preference for shorter, or longer, jobs driving the results. Second, if driver supply differentially effects the surplus for customers of different distance rides, the distance of rides may be correlated with labor supply separately from the wage. Broadly, this could work in two ways. If longer distance jobs initially have a higher value to customers then, when driver supply is low and waiting times are higher, only longer distance jobs will remain worthwhile. Or if the value of different distance jobs to customers is differentially affected by labor supply, and only long distance jobs remain beneficial. These concerns are alleviated by the fact that the firm prices according to distance—and so discriminates precisely on this variable.

I estimate the Behavioral Frisch elasticity analogously to Tube strikes, but I replace the vector of Tube strike dummies T in equation (2.2) with the scalar mean job distance

Table 2.3.4: Behavioral Frisch Elasticity Results

	Dependent variable:			
	log(Shift length)			
	Strikes IV	Mean job dist. IV	LOM wage IV	Both IV
	(1)	(2)	(3)	(4)
log(Shift wage)	0.801^{***}	0.117^{***}	0.114^{***}	0.117^{***}
	(0.102)	(0.005)	(0.013)	(0.005)
Observations	654,045	3,466,056	3,466,055	3,466,055
R^2	0.013	0.081	0.082	0.081

*p<0.1; **p<0.05; ***p<0.01

Notes: This table presents coefficients from regressions of shift length on the wage. The wage is instrumented for with alternative, and the regression includes controls for driver fixed effects, a factor variable for start time of the shift, Ramadan, bank holidays, days of the week, months, fare schedule, and a 6th order polynomial time trend. The regression uses data from the full time horizon, except in the strikes column, and balanced sample. Standard errors are clustered at the driver level.

d. Therefore, equation (2.2) is replaced with,

$$w_{i,s} = \delta_i + \mu \cdot d_{i,s} + \Lambda \cdot X_{i,s} + v_{i,s}, \tag{2.3}$$

The resulting estimate is shown in column (2) of table 2.3.4 and, again, standard errors are clustered at the driver level. The estimate is less than a sixth of the True Frisch elasticity implied by the Tube strikes, which is shown in column (1) for comparison. The standard errors in column (2) are small because mean job distance explains vastly more of the variation in wages. The Behavioral Frisch elasticity estimate is very statistically significantly different from the True Frisch elasticity estimate. Moreover, the point estimate of 0.12 is at the lower bound of previous estimates of the *Hicksian* labor supply elasticity, which are typically in the range 0.1 to 0.3 (Keane, 2011). Columns (3) and (4) display estimates of the Frisch elasticity using the leave-out mean (LOM) wage and a combination of the LOM wage and mean job distance as instruments, respectively. The results barely change.

2.3.3 Control Function Model

In order to explore the difference between the True and Behavioral Frisch, I implement a control function model to trace out the full shape of the labor supply function using the mean job distance instrument. This requires two further assumptions, which are stronger than those necessary for instrumental variable estimation. Firstly, an independence assumption implies the mean shift distance contains no additional information about the shift length after conditioning on the wage. Secondly, a functional form assumption relates to how one controls for distance's effect on the wage. I state these assumptions more precisely as I outline the four steps that underlie my estimation procedure.

First, I residualize hours h, wages w, and mean job distances d with respect to all my controls; I denote these residualized variables with a dot $\dot{\bullet}$. Residualizing ensures that these variables are uncorrelated with the controls, but it does not guarantee independence. Therefore, this step is not without loss of generality because the identifying assumption in control function models revolves around independence. However, this simplification makes the analysis more convenient and, importantly, allows me to specify a more flexible control function.

Second, I specify a first stage where the residualized wage $\dot{w}$ is stated in terms of the residualized mean distance $\dot{d}$. There is a trade-off in how flexibly this relationship is formulated. On the one hand, making the relationship more flexible will make the independence assumption more plausible but, against this, it makes specifying a more flexible control function less feasible. With this balance in mind, I specify a fifth order polynomial for the first stage relationship,

$$\dot{w}_{i,s} = \zeta + \sum_{j=1}^{5} \eta_j \cdot \dot{d}_{i,s}^{j} + e_{i,s}. \tag{2.4}$$

Third, I specify a functional form for the object of interest $\mathbb{E}\left[\dot{h}|\dot{w}, \dot{d}\right]$. For this specification the aforementioned trade-off does not exist—the more flexible the better. I specify a tenth order polynomial for the relationship,

$$\dot{h}_{i,s} = \iota + \sum_{j=1}^{10} \kappa_j \cdot \dot{w}_{i,s}^{j} + \varepsilon_{i,s}. \tag{2.5}$$

Results are robust to this specification; the estimated labor supply function resembles a cubic function and so is not constrained in practical terms.

In the fourth and final step, I state and use the two assumptions that enable me to sketch out the labor supply function. Firstly, my approach requires an independence assumption: the joint distribution of the residuals from regressions (2.4) and (2.5) is independent of the residualized distance variable $(e, \varepsilon) \perp \dot{d}$. This allows me to write,

$$\mathbb{E}\left[\dot{h}\middle|\dot{w}, \dot{d}\right] = \iota + \sum_{j=1}^{10} \kappa_j \cdot \dot{w}^j + \mathbb{E}\left[\varepsilon\middle|e, \dot{d}\right] \tag{2.6}$$

$$= \iota + \sum_{j=1}^{10} \kappa_j \cdot \dot{w}^j + \mathbb{E}\left[\varepsilon\middle|e\right]. \tag{2.7}$$

The second assumption is the specification of the conditional expectation in equation (2.7). Often this is assumed to be linear, but given the size of my data I can be more flexible. As a baseline, I specify the control function as a quadratic, however, I check the robustness of my findings up to a fourth order polynomial. That is, I write,

$$\mathbb{E}\left[\varepsilon\middle|e\right] = \nu \cdot e + \xi \cdot e^2. \tag{2.8}$$

In order to implement this empirically, knowledge of $e_{i,s}$ and $e_{i,s}^2$ is required. I first construct the estimates $\hat{e}_{i,s}$ with the regression specified in equation (2.4). Then, I square both sides of equation (2.4) to yield a specification which can estimate $\hat{e}_{i,s}^2$. Note that this latter regression requires knowledge of $\hat{e}_{i,s}$. I proceed in similar steps when constructing the third and fourth order polynomial control functions. The number of parameters that are estimated in these latter regressions rises rapidly with the order of the polynomial, hence, the trade-off mentioned in the specification of (2.4).

The independence assumption required to implement a control function model is stronger than the exogeneity assumption required in an instrumental variables regression. This condition will not be perfectly met in practice, but I offer three defenses of the approach: (i) the output is consistent with results from the instrumental variable analysis, which suggests that (ii) although the assumption may be violated mildly, the results are still informative; and (iii) behavioral theories of labor supply, such as income targeting, predict a non-monotonic relationship between hours and wages, and this is the correct way to identify non-monotonicities.[13]

Given the steps above, I use the following regression to identify drivers' labor supply

[13] Farber (2015) estimates a Frisch elasticity for wages close to the mean because of suspected non-monotonicities. However, such an approach is akin to a "forbidden regression". Stronger assumptions, as used here, are necessary.

function,

$$\dot{h}_{i,s} = \iota + \sum_{j=1}^{10} \kappa_j \cdot \dot{w}_{i,s}^j + \nu \cdot \hat{e}_{i,s} + \xi \cdot \hat{e}_{i,s}^2 + r_{i,s}. \tag{2.9}$$

The estimated labor supply function is illustrated by the blue line in figure 2.3.1. The grey bars show the distribution of residualized wage rates $\dot{w}$ used in the analysis and the dashed red line marks the mean wage rate. The function is striking; it is steeply rising for many wage rates—in line with the findings from the Tube strike analysis—but is *negatively* sloped for wage rates just above the mean wage. This result is robust to the specification of the control function, as shown in figure 2.3.2 which illustrates labor supply functions for different polynomial orders for the control function.

The labor supply function reveals that the Behavioral Frisch elasticity masks significant non-monotonicity in drivers responses to changing wage rates. Indeed, while the latter is qualitatively compatible with a neoclassical model of labor supply, even if not quantitatively, the downward sloping portion of the labor supply function is incompatible with optimal decision making, which I explain further in section 2.4.

Figure 2.3.1: Labor Supply Function

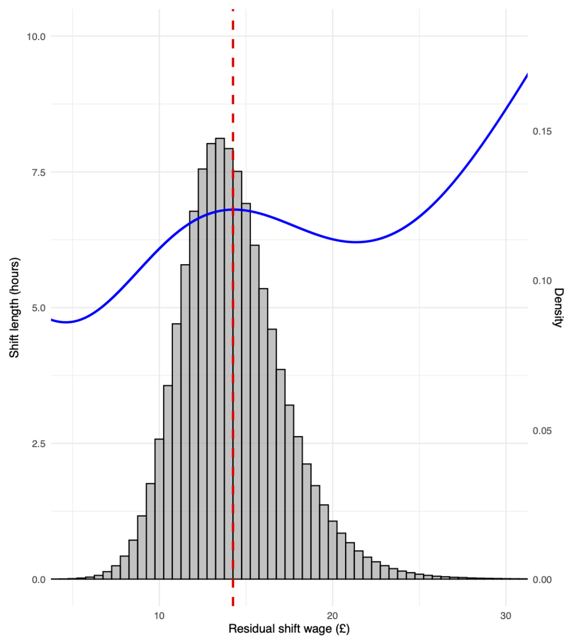

Notes: This figures presents the function identified by the control function model. The grey bays denote the empirical density histogram of the residualized wage rate, where each bin is an hour of the day, and the dashed red line marks the mean of that variable.

Figure 2.3.2: Robustness to Control Function Form

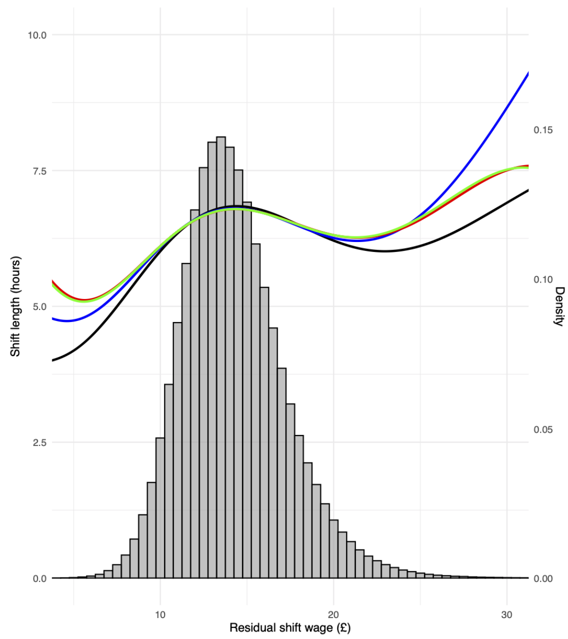

Notes: This figures presents the functions identified by the control function models which specify the control function as linear, quadratic, cubic, and quartic.

2.3.4 Discussion

The hours-wages nexus is a form of aggregation which lends itself to familiar concepts in labor economics, however, it may over simplify the economic environment that drivers face. I discuss the impact of non-constant wages, schedule rigidities, and expectations in order to ameliorate these concerns.

Non-constant wage. In reality, drivers do not face a constant wage. Figure 2.3.3 illustrates the median hourly wage for drivers throughout the day for weekdays and weekends, separately. It is evident that wages have a peak in the morning, which is more pronounced during weekdays, and then rise initially steadily and then rapidly from midday to midnight, before declining until the start of the morning peak. This is problematic because the wage a driver has experienced through their shift may be different from the wage they would have received had they continued their shift, so that regressions of hours on wages do not capture the actual trade-off which drivers face. I present three pieces of evidence which suggest this is not a significant problem.

First, fares are a key driver of wages and exhibit significant auto-correlation. I regress current fares on fixed effects for the shift and the previous fare. Given the dynamic panel structure, I estimate this equation in first differences and instrument for the previous fare with its own preceding fare. A one standard deviation in the current fare level leads to an increase of the next fare of £2.76 (s.e. 0.13) in addition to level effects between shifts. This suggests there is strong persistence in the wage rate caused by auto-correlation in fares.

Figure 2.3.3: Wage by Hour of the Day

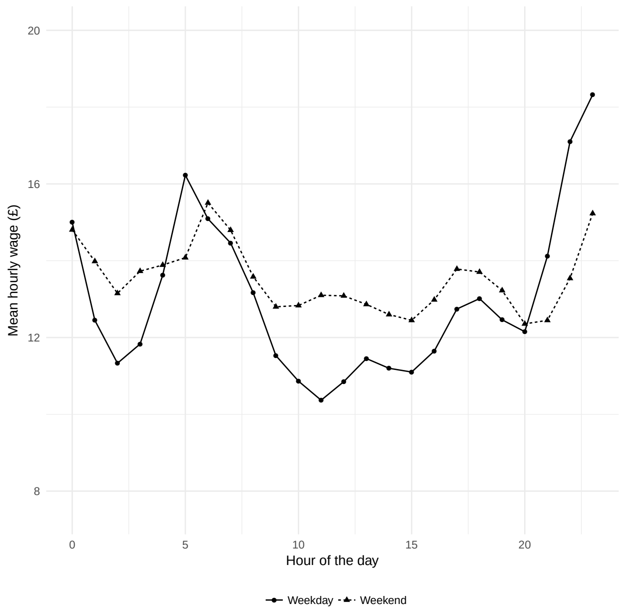

Notes: This figure presents the mean wage at every hour of the day, which is constructed by allocating earned fares to hours of the day proportional to the hours in which the job took place.

Second, alternative approaches that do not rely on constant wages produce results that are consistent with the findings here. For example, Farber (2015) and Thakral and Tô (2017) use probability stopping models in order to test if labor supply behavior deviates from the neoclassical benchmark. The focus of this approach is to test whether accumulated income in a shift is predictive of a driver ending their shift. In the neoclassical model, the amounts of income under consideration are not sufficient to affect the marginal utility of income, therefore, there should be no relationship between accumulated income and the probability of ending a shift.[14] I follow the approach of Thakral and Tô (2017), which non-parametrically controls for the disutility of hours worked, to assess the importance of income in determining the end of a shift. The approach regresses a dummy variable for ending the shift after a job on variables that summarize

[14]If wages are positively auto-correlated—as they are in the data—then accumulated income should lead to a lower probability of ending a shift, *i.e.*, a negative coefficient on accumulated income.

a drivers experience throughout the shift thus far, such as accumulated income and other job-level controls. A separate regression is run for different durations in a shift so that each regression is run on different data. These data bins are defined by 30 minute windows through a shift.[15] This method is akin to a local linear regression and flexibly controls for the disutility of work. The regressions are specified as,

$$\mathrm{q}_{i,j} = \pi_i + \xi \cdot y_{i,j} + \Upsilon \cdot Z + \rho_{i,j},$$

where the new subscript j denotes the job, the variable q is a dummy that takes value one if a driver ends their shift after the job, y is the logarithm of accumulated income, and Z is a rich vector of controls. The coefficients on log accumulated income ξ from each regression are shown by the black points in figure 2.3.4. All are significantly positive at durations through a shift, as in Thakral and Tô (2017).[16] Mean stopping probabilities are plotted with grey bars, which allows a naïve calculation of the elasticity of the probability of stopping with respect to income[17] in order to get a sense of the magnitude of these effects. At the average shift length duration, the elasticity is of the order 0.25, which indicates a significant behavioral response; this is consistent with a downward sloping labor supply function.

Third, a non-constant wage cannot easily explain the shape of the labor supply function that I uncover with the control function model. For example, drivers could perceive the wage to be mean reverting. In such a case, if a driver experienced a wage just above the mean, they would expect their wage to drop in future. As a result they may end their shift earlier on an above average wage because they anticipate lower wages in the future. This is incompatible with the fact that hours increase with wages either side of the downward sloping portion of the labor supply function.

[15] For example, the first regression uses job level data where the jobs occur in the first 2.5 to 3 hours of the shift. The second regression uses job level data where the jobs occur in the first 3 to 3.5 hours of the shift, and so on.

[16] I also replicate the timing effects of these authors; income later in the shift has a greater effect on the probability of stopping, as illustrated in figure 2.3.5, and income at one point in time becomes less and less influential. However, unlike Thakral and Tô (2017) I find that the influence of earlier income can be negative on the likelihood of stopping. This finding can be reconciled with the authors' adjusting reference point, if the wage is auto-correlated throughout the day, as I argue, and drivers update their targets in a Bayesian fashion.

[17] This is naïve because it does not account for the covariance between mean stopping probabilities and behavioral responses to income.

Figure 2.3.4: Stopping Model Income Effect Estimates

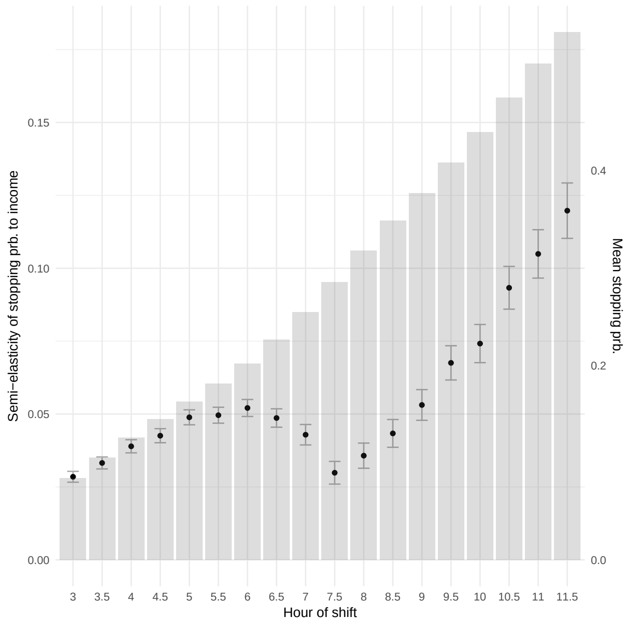

Notes: This figure presents the coefficients on accumulated income from the probability stopping model regressions as black dots with standard error bars, and the mean stopping probabilities as grey bars.

Schedule rigidities. Schedule rigidities can cause non-convexities in drivers' decision making environments, which confound the interpretation of labor supply elasticities. The existence of schedule rigidities should show up as a low response of hours to wages. If the rigidities are fixed, then the daily Frisch elasticity will be estimated as zero. However, stochastic rigidities in combination with a non-constant wage can lead to the spurious estimation of non-zero elasticities. The logic is highlighted by Thakral and Tô (2017); if the wage rate is falling over time and drivers face a stochastic schedule rigidity, then a regression of hours on wages will yield a negative coefficient because longer shifts—that are only longer because of a shock to the schedule rigidity—necessarily entail a lower wage. This does not appear to be a major concern in my setting for two reasons. Firstly, my results suggests that drivers adapt their hours very positively to variation in the wage rate due to Tube strikes and to low or very high wages. Secondly, most drivers end their shift between 10:00am and 11:00pm, which is a period of time where the wage rate is rising and continues to rise until midnight. There is still a significant mass of drivers who finish their shifts around midnight, which is around the time the wage falls substantially. I replicate the analysis without these shifts and the results are unchanged; for example, the Behavioral Frisch is not statistically significantly

Figure 2.3.5: Stopping Model Timing Effect Estimates

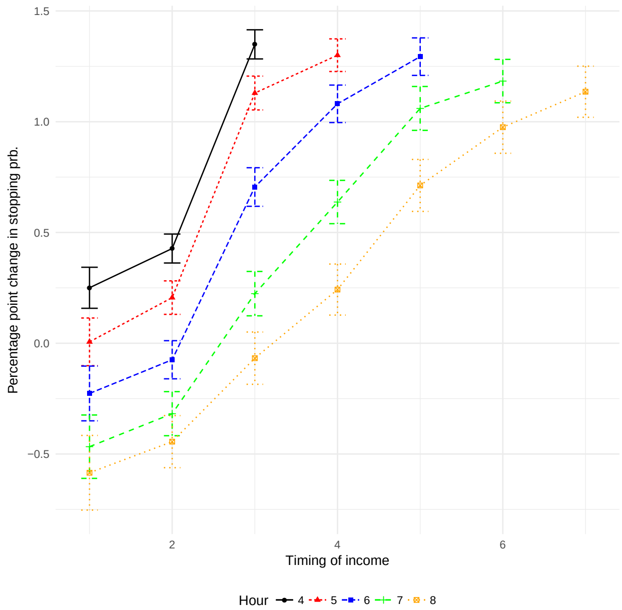

Notes: This figure presents the coefficients on income earned at different durations in shifts for shifts of different durations. For example, the blue line shows coefficients from a regression on shifts six hours through, and shows the influence of income earned in hour one, two, three, and so on, on the probability of stopping.

different.

Expectations. Another potential reason for the difference between the two Frisch elasticities is that one source of wage variation is expected, while the other is not. In particular, variation in the wage rate that comes from fluctuations in mean job distances may not be expected and, as a result, drivers do not have time to substitute labor intertemporally. This phenomenon would imply that hours do not respond significantly to variation in the wage rate. However, the labor supply function shows the shift length does respond strongly to wages, albeit in different directions depending on the level of the wage.

The inability of these factors to explain the disparate estimates of labor supply responses demands a more fundamental explanation, which I seek to give in section 2.4.

2.4 Behavioral Interpretation

I posit that there is a fundamental difference between the True Frisch and Behavioral Frisch: the True Frisch is the output of decision making when behavioral biases are attenuated, and the Behavioral Frisch and control function model identifies labor supply which is representative of general daily behavior—and subject to behavioral biases.

The Tube strikes and the mean job distance instruments cause wage variation that induces identical labor supply responses in typical models where drivers only optimize over hours worked and income, but in reality these sources of variation are very different in nature. Tube strikes are seldom events which receive substantial media coverage and, as a result, lead to salient wage variation with potential framing effects. These characteristics of Tube strikes attenuate driver biases, which allows them to behave closer to optimal.[18] Conversely, variation in mean job distance is very common and engages the standard heuristics used by drivers. This hypothesis simultaneously explains the non-neoclassical behavior evident in the labor supply function, as well as the probability stopping model, and the differences between the labor supply responses uncovered by the two instruments.

The behavioral nature of the labor supply function follows from its negative slope for a range of wages. The possibility that this is the result of any neoclassical income effects is ruled out by the estimation of a *less* negative Marshallian labor supply elasticity.[19] In a neoclassical model, this elasticity bounds all labor supply responses from below since it incorporates the income effect of price changes. Therefore, I can reject a neoclassical model of labor supply because the most negative elasticity from the labor supply function (-0.36) significantly exceeds the benchmark Marshallian estimate from table 2.B.2 (-0.14, s.e. 0.04). Further, the mean job distance instrument and other controls explain over half of all wage variation, which supports the interpretation of the labor supply function as indicative of general behavior. In contrast, the True Frisch elasticity

[18] Salience is the quality of being particularly noticeable, therefore, its role in debiasing drivers can be thought of as a form of information provision. Allcott and Taubinsky (2015) provide a clear and concise discussion of how information provision can eliminate some biases, including biased beliefs, exogenous inattention, costly information acquisition, and costly thinking models. Moreover, the potential framing effects, that is the sense in which days with Tube strikes feel somewhat "special", may counter other biases, such as self-control issues.

[19] This estimation leverages a permanent fare reform that raised driver wages by 10 percent, see table 2.B.1. A detailed description of the estimation that yields the results in table 2.B.2, which are close to Ashenfelter et al. (2010), is contained in appendix 2.B.

falls into the range of previously estimated Frisch elasticities (Reichling and Whalen, 2012), where behavioral biases are not perceived to be severe, and easily exceeds the Marhsallian elasticity.

Survey evidence from a different set of drivers, who also pick their hours and are paid by the job, supports the postulation of a divergence between driver behavior and preferences too. As part of a general survey,[20] we asked these drivers how they determine their shift length. The options we offer are shown in table 2.4.1 and we randomize the order in which they appear to respondents. The two rules in bold and the frequency with which they are selected are of note; both rules imply the same relationship between the wage and shift length, however, working to earn a certain amount each day is vastly more popular than working less if pay is higher. This is not surprising because the latter is patently suboptimal, but under the guise of an income target it can appear quite reasonable. Further, when asked explicitly about how they would respond to temporary versus permanent wage increases, all drivers' answers were consistent with the ordering of Frisch and Marshallian elasticities implied by the neoclassical model. A similar cognitive dissonance amongst the drivers I analyze could sustain the divergence between behavior and preferences that is revealed in the conflicting Frisch elasticities and labor supply function.

If Tube strikes reveal decisions designed to maximize utility and the labor supply function implies general behavior, then I can use revealed preference logic in order to uncover the parameters that govern the trade-off between hours worked and income and, in turn, estimate the implications of suboptimal behavior for welfare. This approach does not require a firm stance on the behavioral biases that are afflicting drivers; I only need to identify behavior and preferences. However, the source of biases remains important for three reasons. Firstly, the source of biases determines the extent to which the salience of Tube strikes removes drivers' biases. Secondly, it can affect the interpretation of the welfare cost that I will estimate. For example, if behavior is not optimal because of optimization costs, the welfare loss I estimate could provide a lower bound on the cost of optimization. And thirdly, the source of biases could have consequences for the optimal policy response; if optimization costs are responsible for suboptimal behavior, then forcing individuals to cognize would likely not improve welfare.

[20] More details of this survey are provided in appendix 2.C.

Table 2.4.1: Surveyed Labor Supply Rules

Labor Supply Rules	Frequency
Work for a certain number of hours	126
Work to earn a certain amount each day	**123**
Work as long as possible	96
Work longer if pay is higher and *vice versa*	88
Work shorter if pay is higher and *vice versa*	**43**

Notes: This table presents the frequency of responses to alternative labor supply rules from a sample of 476 drivers. The two rules in bold imply the same qualitative relationship between the wage and hours worked.

There are numerous potential biases that could affect drivers, including: limited attention to the wage rate (Gabaix and Laibson, 2006a; DellaVigna and Pollet, 2009; Bordalo et al., 2013), costly information acquisition since the going wage is not readily available (Gabaix et al., 2006; Sallee, 2014), costly thinking stemming from the cognitive burden of optimizing labor supply (Gabaix, 2014; Caplin and Dean, 2015; Sims, 2003), imperfect self-control (Gruber and Köszegi, 2001; Gruber and Kőszegi, 2004; Bernheim and Rangel, 2004), and biased beliefs (Spinnewijn, 2015). In the language of Allcott and Taubinsky (2015), I require Tube strikes to be a "pure nudge". That is, Tube strikes resolve drivers' biases without any confounding effects on labor supply. It is clear that salience will counter some biases more than others; information acquisition should be much less costly than otherwise since Tube strikes are very prominent to drivers. On the other hand, self-control biases may persist because the intertemporal properties of the costs and benefits to driving are unchanged.

Section 2.6, where I use the framework of Bernheim and Rangel (2009), formalizes how I leverage my empirical results to learn about preferences and behavior. But beforehand, in the next section, motivated by the idea of a divergence between behavior and preferences, I derive an estimable expression that quantifies the welfare cost of drivers' deviations from the optimal benchmark.

2.5 Behavioral Welfare Theory

Motivated by the empirical evidence presented in section 2.3 and the behavioral interpretation in section 2.4, I construct a simple static model to conduct a normative analysis of behaviorally biased labor supply. The model yields an object that captures an individual's change in utility, when labor supply moves from biased to optimal.[21] This object can be approximated by an expression which contains only a small number of sufficient statistics and is estimable given sufficient data. It reveals that the welfare cost of biased labor supply stems from the wedge in the intratemporal optimality condition, *i.e.* the difference between the wage and the marginal rate of substitution (MRS) between work and consumption. The model is static but the key dynamic effects, namely whether labor supply behavior translates to different consumption levels, can easily be captured in the empirical implementation.

I will use the language of drivers, consumption, and hours to match the empirical context of this paper but the potential applications of this expression are much more broad. Indeed, this formula is closely related to a continuous generalization of expressions for changes in consumer surplus stemming from a price change of a binary good (Allcott and Taubinsky, 2015).[22]

2.5.1 Behavioral Environment

Drivers derive utility from consumption c and disutility from hours worked h according to a utility function $U(c, h)$. They face a budget constraint defined by $c \leq w \cdot h + I$, where w is the wage rate and I is an exogenous, additional source of income. For convenience, I omit the latter variable from notation below since it is not important for what follows. Labor supply is suboptimal because drivers suffer from biases, which leads to two decision rules for consumption and hours, respectively,

$$\{\tilde{c}(w), \tilde{h}(w)\} \notin \underset{\{c,h\} \in \mathbb{B}(w)}{\arg\max}\ U(c, h),$$

where $\mathbb{B}(\cdot)$ is the choice set defined by the budget constraint inequality and non-negativity constraints on c and h. If drivers did behave optimally, then they would follow two optimal rules for consumption and hours, respectively,

$$\{c^*(w), h^*(w)\} \in \underset{\{c,h\} \in \mathbb{B}(w)}{\arg\max}\ U(c, h).$$

[21] Though the theory is adaptable to compare welfare under different types of behavior since it does not rely on any application of the envelope theorem.

[22] This generalization is presented in the appendix 2.D, alongside the proof of the BWE.

I am interested in a money-metric measure of the change in utility due to a change in labor supply from biased to optimal, namely,

$$\Delta(w) = \frac{U\big(c^*(w), h^*(w)\big) - U\big(\tilde{c}(w), \tilde{h}(w)\big)}{U_c\big(\tilde{c}(w), \tilde{h}(w)\big)}. \tag{2.10}$$

This quantity varies with the wage rate, which varies between days. Therefore, I treat the wage rate as a random variable so that the object of interest is the expected change in utility from a move to optimal behavior for one shift,

$$\mathbb{E}_w\left[\Delta(w)\right], \tag{2.11}$$

where $\mathbb{E}_w[\cdot]$ is the expectations operator that integrates with respect to the wage random variable. Higher order moments, such as the variance, can also be calculated and may be of interest; if a driver considers an investment to reduce the biases they suffer, they may also care about the variance of the return of that investment.

2.5.2 Sufficient Statistics Formula

While equation (2.10) is the precise quantity of interest, it provides little insight or practical use. In the theorem below, I derive an approximation of this expression that is estimable with a small number of sufficient statistics, which can be used to back out the necessary parameters. Moreover, the sufficient statistics are familiar labor supply elasticities that can be taken "off the shelf" if needed.

Theorem 1 (Behavioral Welfare Expression - BWE) *If the utility function is additively separable in consumption and hours, $\Delta(w)$ can be approximated by a second order Taylor series approximation and difference quotients in order to yield,*

$$\Delta(w) \approx \Delta c - \frac{1}{2} \cdot \eta(w) \cdot \frac{\Delta c^2}{\tilde{c}} - MRS \cdot \Delta h + \frac{1}{2} \cdot \Delta MRS \cdot \Delta h, \tag{2.12}$$

where,

$$\begin{aligned} MRS &= \tfrac{v'(\tilde{h}(w))}{u'(\tilde{c}(w))}, & \Delta h &= h^*(w) - \tilde{h}(w), \\ \Delta MRS &= -\tfrac{MRS}{\tilde{h}(w)/\gamma(w)} \cdot \Delta h, & \gamma(w) &= \tfrac{\tilde{h}(w) \cdot v''(\tilde{h}(w))}{v'\left(\tilde{h}(w)\right)}, \\ \eta(w) &= -\tfrac{\tilde{c}(w) \cdot u''\left(\tilde{c}(w)\right)}{u'(\tilde{c}(w))}, & \Delta c &= c^*(w) - \tilde{c}(w). \end{aligned}$$

See appendix 2.D for proof.[23]

When consumption is written as earned income, the approximation is made up of three intuitive components,

$$\Delta(w) \approx \underbrace{(w - \text{MRS}) \cdot \Delta h}_{\text{Wedge}} + \underbrace{\frac{1}{2} \cdot \Delta(w - \text{MRS}) \cdot \Delta h}_{\text{Change in wedge}} - \underbrace{\frac{1}{2} \cdot \frac{\eta(w)}{\tilde{c}} \cdot (w \cdot \Delta h)^2}_{\text{Income effect}},$$

which highlight the source of the welfare loss: a wedge in the FOC adjusted for diminishing marginal utility. This is easily explained with an example. Take a hypothetical day where the wage rate was high but the driver worked too little to equate the wage rate and their MRS between labor and consumption. Then, the *Wedge* term captures foregone income corrected for the cost of earning this income. The *Change in wedge* term corrects the Wedge term for the fact that the cost of earning income increases with hours already worked, and the *Income effect* corrects the Wedge term again for diminishing marginal utility. Conversely, if the wage was low and the driver worked too much, the formula captures foregone utility that stems from working for too low a wage with the analogous corrections.

In order to think about how this plays out in expectation, it is useful to consider a simple labor supply heuristic, such as perfect income targeting. Under this heuristic, the driver always works until they earn a set income target so that the elasticity of hours with respect to wages is -1. If we ignore the second order terms in equation (2.12) and take expectations, this yields,

$$\mathbb{E}_w \left[(w - \text{MRS}) \cdot \Delta h)\right] = \mathbb{E}_w \left[(w - \text{MRS})\right] \cdot \mathbb{E}_w \left[\Delta h\right] + \mathbb{C}_w \left[(w - \text{MRS}), \Delta h)\right],$$

[23] I leave the curvature objects $\eta(\bullet)$ and $\gamma(\bullet)$ as explicit functions of the wage, while suppressing this notation for other terms, in order to highlight a further approximation I use in the empirical implementation. Further, this notation also emphasises that I assume consumption is constant in both the biased or optimal paradigm.

where $\mathbb{C}_w[\cdot]$ is the covariance operator integrating with respect to the wage random variable. For simplicity, assume that the income target is set such that $\mathbb{E}_w\left[(w - \text{MRS})\right] \approx 0$, then the welfare consequences of this labor supply bias depend only on $\mathbb{C}_w\left[(w - \text{MRS}), \Delta h)\right]$, which is positive. Why? When the wage is high, drivers reach their income target quickly and work too little relative to the optimum—Δh is positive and the wedge between the wage and MRS is large. While when the wage is low, the driver who follows an income targeting labor supply bias works too much, so that Δh is negative, and the MRS exceeds the wage so that the wedge is negative. This highlights the inefficiency of income targeting more generally; it causes a negative covariance between hours worked and the wedge in the intratemporal optimality condition.

The accuracy of the approximation embodied in the BWE depends on the size of higher order derivatives of the utility function—if they are small, the approximation is good. In practical terms, the second order approximation corrects for changes in the cost of labor, as measured by the MRS, in a linear way. I find that the MRS is close to linear in hours and so third order and higher terms are not of concern. In situations where the cost of labor (or benefit of a good) is more convex this may not be the case.

The BWE succinctly captures drivers' welfare losses due to suboptimal behavior, and is amenable to estimation. I lay out and implement a roadmap to estimation in section 2.6.

2.6 Welfare Analysis

The empirical evidence from section 2.3 and the theory from section 2.5 can be brought together in order to conduct a welfare analysis of drivers' labor supply decisions. Here, I discuss an approximation and an assumption to facilitate this, before outlining and implementing a roadmap to the estimation of welfare of losses due to labor supply biases. I conclude the section by presenting and discussing the results.

2.6.1 An Approximation and an Assumption

The BWE contains several objects that are related to driver preferences. In this subsection, I explain the steps necessary to identify these with familiar labor supply elasticities.

The approximation. To reduce the burden of estimation I impose that $\eta(w)$ and $\gamma(w)$ are constant rather than functions of the wage rate. I note that the BWE is linear in these terms, and that the MRS is also linear in these terms up to a first order approximation around some level of hours. Therefore, this imposition yields a further approximation of $\Delta(w)$. It also has clear implications for the shape of the utility function; approximating $\Delta(w)$ with $\eta(w)$ and $\gamma(w)$ as constants is equivalent to assuming utility is constant relative risk aversion (CRRA)—*conditional on using the BWE*—because,

$$\{\eta(w), \gamma(w)\} = \{\eta, \gamma\} \iff U(c,h) = \frac{c^{1-\eta}}{1-\eta} - \theta \cdot \frac{h^{1+\gamma}}{1+\gamma}.$$

The CRRA equivalence offers two convenient corollaries. First, the shape of the MRS is known without any further loss of generality,

$$\text{MRS} = \theta \cdot c^{\eta} \cdot h^{\gamma}, \tag{2.13}$$

and so optimal labor supply is implied by $w = \text{MRS}$. Second, there is a simple mapping from the True Frisch elasticity and Marhsallian elasticity estimates to the curvature parameters,

$$\varepsilon^{\text{Frisch}} = \frac{1}{\gamma}, \tag{2.14}$$

$$\varepsilon^{\text{Marsh.}} = \frac{1-\eta}{\gamma+\eta}, \tag{2.15}$$

where the latter equality assumes that drivers have no source of capital income.

The assumption. A Frisch elasticity that is estimated with data from always optimizing drivers maps one-to-one with the curvature of the disutility of labor γ according to equation (2.14). However, the True Frisch elasticity implicitly compares behavior when optimizing under Tube strikes with biased behavior otherwise, and the associated wage differences. To see this clearly, consider that the Tube strikes instrument is a single binary instrument and that this instrument, hours, and wages have been residualized with the respect to all controls. Then I could estimate my True Frisch elasticity with a

Wald estimator,

$$\hat{\beta}^{\text{Wald}} = \frac{\bar{h}_1 - \bar{h}_0}{\bar{w}_1 - \bar{w}_0},$$

where the bar notation $\bar{\bullet}$ indicates the empirical average, and the subscripts $\{0, 1\}$ denote whether the average refers to outcomes during or not during Tube strikes, respectively. If $h_{i,t}$ is always defined by $\log(w_{i,t}) = \log(\theta) + \gamma \cdot \log(h_{i,t})$ it is simple to show that $\hat{\beta}^{\text{Wald}} = 1/\gamma$, where I have omitted the influence of consumption for simplicity. However, when labor supply is biased outside of Tube strikes shift length does not necessarily satisfy the first order condition and so it is not clear that anything about preferences is revealed. But, if biases satisfy,

$$\mathbb{E}\left[\log(w_{i,t})\right] = \log(\theta) + \gamma \cdot \mathbb{E}\left[\log(h_{i,t})\right]. \tag{2.16}$$

then γ is still identified. Therefore, I require biased labor supply to obey the equality in equation (2.16): the logarithm of drivers' first order conditions holds in expectation. Note that this is a restriction on the average level of hours, but does not constrain the shape of the biased labor supply function. Beyond its convenience, this condition is attractive for three reasons. Firstly, it grants drivers a degree of sophistication and so does not mechanically overstate the welfare cost of behavioral biases. Secondly, the condition would hold if drivers aim to ensure their intratemporal condition holds in expectation despite their biases, but neglect Jensen's inequality. Lastly, the condition would hold without this neglect if the MRS is linear in hours, which is roughly the case with my results. Given (2.16), the curvature of consumption utility can also be backed out of equation (2.15) with the Marshallian elasticity, and no other condition on the nature of biases is necessary.

2.6.2 Estimation

The estimation of the BWE in expectation requires a number of ingredients; table 2.6.1 lists these components alongside how they are estimated and a graphical, mathematical, or numerical representation.

First, I estimate the distribution of wages with a kernel density estimator which is run on the wages observed in my sample. This gives a clear interpretation to my results:

Table 2.6.1: Ingredients for BWE in Expectation

Function/Parameter	Estimation	Reference
$f_w(\bullet)$	Kernel density estimator	Figure 2.6.2
$\tilde{h}(w)$	Control function model	Figure 2.3.1
$h^*(w)$	Intratemporal optimality condition	Figure 2.6.1
$\eta(w) = \eta$	True Frisch + Marshallian + eq. (2.15)	Table 2.6.2, row one
$\gamma(w) = \gamma$	True Frisch + eq. (2.14)	Table 2.6.2, row two
$\theta \cdot \tilde{c}^{\eta}$	Hour levels during Tube strikes	$\mathbb{E}_{w\mid\text{strikes}}\left[\frac{w}{h^{\gamma}}\right]$
	Calibrated to income level	$\mathbb{E}_w\left[w \cdot \tilde{h}(w)\right] = \mathbb{E}_w\left[w \cdot h^*(w)\right]$
$\tilde{c}$	Average shift income	$\mathbb{E}_w\left[w \cdot \tilde{h}(w)\right]$
Δc	Change in average shift income	$\mathbb{E}_w\left[w \cdot \left(\tilde{h}(w) - h^*(w)\right)\right]$
ν	% income + hours targeters in survey	$\approx$0.66

Notes: This table presents the different objects required to calculate the expectation of the BWE, how these objects are estimated, and where they are illustrated in the paper.

Table 2.6.2: Curvature Parameter Estimates

Parameter	Estimate	Std. error
η	1.37	0.02
γ	1.25	0.03

Notes: This table presents estimates for the parameters η and γ. The estimates are produced using equations 2.15 and 2.14, and the coefficients from column (1) of table 2.B.2 and column (2) of table 2.3.2, respectively. Standard errors are constructed using the delta method.

the expected welfare loss due to suboptimal labor supply, if one were draw a random shift from the sample.[24] Labor supply under biases is specified as the function implied by the control function model, which is estimated with the mean job distance instrument. Optimal labor supply is characterized by the intratemporal optimality condition, which equates the wage and the MRS. The MRS is specified as in equation (2.13), which requires the estimation of γ and $\theta \cdot \tilde{c}^{\eta}$, where $\tilde{c}$ is the constant consumption level under baised labor supply, and η is also required separately for the income effect. The curvature parameters η and γ are backed out from the True Frisch elasticity and Marshallian elasticity estimates using equations (2.14) and (2.15); these estimates are presented in table 2.6.2, where standard errors are calculated with the delta method.

I consider two approaches to parameterize $\theta \cdot \tilde{c}^{\eta}$, which relates the optimal level of hours supplied. Firstly, I set it to match the observed level of hours during Tube strikes

[24]Instead, if one were to take the distribution of wages a particular driver receives, then the interpretation would be the expected welfare loss on any given shift for that driver.

when drivers are hypothesized to be behaving optimally. Given this hypothesis, there are a number of ways to derive the level parameter, which make different implicit assumptions about the nature of random (not behavioral) errors in driver behavior. This is discussed in appendix 2.E, and I use the specification in row 6 of table 2.6.1. Secondly, I calibrate the level of optimal hours to ensure the daily level of income is unchanged when behavior moves from bias to optimal. Both of these approaches assume constant consumption, either when behaving with biases or optimally since the exercise considers daily fluctuation in wages, which drivers are able to insure against. Naturally, this has ramifications for the estimation of $\tilde{c}$. I assume that constant consumption under biased behavior equals average daily income that is generated by $\tilde{h}(w)$ and the distribution of wages $f_w(\bullet)$. This parameter only mediates the strength of the income effect in the BWE and so it has no effect on the results when $\theta \cdot \tilde{c}^{\eta}$ is calibrated to maintain constant consumption. The change in consumption between the optimal and behavioral regime is set as the difference in expected income generated by the two labor supply schedules.

It is plausible that only a proportion of drivers ν are responsible for the deviation from optimal behavior. That is,

$$\tilde{h}(w) = \nu \cdot \hat{h}(w) + (1 - \nu) \cdot h^{*}(w),$$

where $\hat{h}(w)$ is the labor supply of the purely biased drivers, such that $\tilde{h}(w)$ is now viewed as the confluence of optimal and suboptimal labor supply by drivers. Given ν, it is easy to infer $\hat{h}(w)$, however, ν is unknown. The smaller ν is, the more severe biased behavior deviates from the optimal. I calibrate ν to equal the proportion of income and hours targeters in the survey presented in table 2.4.1. This calibration is also appealing because if biases were concentrated amongst an even smaller group of drivers, the most negative elasticity of $\hat{h}(w)$ would fall below -1, which is hard to reconcile with any behavioral theories. Note that two welfare loss estimates can be derived from this exercise: one for the biased individuals, and an average for all individuals, *i.e.* the former estimate scaled by ν.

With two alternative calibrations of $\theta \cdot \tilde{c}^{\eta}$ and homogeneous and heterogeneous behavior, I present four results for the expected daily welfare loss due to labor supply bi-

Table 2.6.3: Expected Daily Welfare Losses for Biased Drivers

		$\theta \cdot \tilde{c}^{\eta}$	
		Tube strike level	Expected income level
ν	Homog.	2.32 (p5, p95) $\Delta u(c)$: -9.01 $\Delta u(h)$: +11.32	2.09 (p5, p95)
	Heterog.	5.29 (p5, p95) $\Delta u(c)$: -13.78 $\Delta u(h)$: +19.07	4.72 (p5, p95)

Notes: This table presents the expected welfare losses in pounds (£) for the 2×2 variations under considerations. When the level parameter is not calibrated to keep expected income equal under biased and optimal behavior, I report how the welfare loss is composed of changes in utility due to consumption and labor supply behavior. 5^{th} and 95^{th} percentiles from 500 bootstraps are presented in the parentheses. The heterogeneity scenario considers $\nu = 0.66$.

ases. I also present a further fifth result, where I strictly assume CRRA preferences in combination with homogeneous behavior and an optimal level of hours during Tube strikes, in order to check that the BWE yields plausible results.

2.6.3 Results

The results of the estimation are summarized in table 2.6.3. The cell in the first row and column considers that all drivers are equally biased, and derives the level parameter from the hours observed during Tube strikes. This reveals an expected daily welfare loss of £2.32. Notably, optimal behavior implies a significant loss in consumption utility of £9.01 (approximately 10 percent of average daily income), which is exceeded by savings in disutility from work. The overall welfare loss only falls slightly, to £2.09, when the level parameter is calibrated to keep expected daily income, and thus consumption, constant. Both these numbers are relatively close to £2.44, which is the loss implied when CRRA preferences are assumed without any approximation.

The expected daily welfare loss rises sharply when a fraction of drivers behave optimally. If two thirds of drivers are responsible for biased behavior then the welfare loss more than doubles. If the level parameter is set to match hours during Tube strikes, the expected loss equals £5.29 for biased drivers. This result must be scaled by $1/\nu$ in order to determine the unconditional, expected welfare loss, which leads to a loss of £3.49. Again, this loss is composed of a consumption loss and an overwhelming gain due to more efficient labor supply. Moreover, this welfare loss does not change markedly when the level parameter is calibrated to keep expected income constant. In this scenario, the loss for biased individuals equals £4.72 and is £3.11 across all drivers.

Discussion. In order to gain a sense of the economic importance of these losses, it is necessary to account for how regularly they are incurred. Driver-weeks most commonly contain five shifts and the mean number of shifts in a week is four. Annually, this regularity of work combined with holidays leads to an average of 200 shifts. Therefore, given preferences that are additively separable across time and negligible discounting over the course of a year, the results imply losses ranging from £418.00 to £1058.00 per annum. For context, Handel and Kolstad (2015) find average losses due to suboptimal health insurance choices, which are made annually, of approximately £1,200.00.[25] Such losses are considered large and are of a comparable magnitude to the results presented here. Accordingly, the losses in this paper point towards significant welfare losses for drivers due to behavioral frictions in exploiting flexibility.

The results also have implications for the role of heterogeneity in the cost of biases and the nature of biases. Firstly, aggregate welfare losses increase even as biases are concentrated amongst a smaller proportion of drivers. This is because welfare losses for biased drivers increase quickly as they become more responsible for the aggregate deviation from optimal behavior, such that aggregate welfare losses increase locally as a smaller proportion of drivers are biased. More precisely, the elasticity of $\mathbb{E}_w\left[\Delta(w)\right]$ with respect to ν is greater than one. It is important to recognize that this non-unitary elasticity arises because drivers make different decisions, which imply different size internalities, for different wages, and they have to make these decisions repeatedly. This mechanism contributes another channel to those identified by Taubinsky and Rees-Jones (2018) as to how heterogeneity in biases can accentuate welfare losses.

Secondly, the results suggest that suboptimal labor supply is not likely to be driven by time-inconsistent decision making, but they are in line with a reference-dependent labor supply hypothesis. Column 1 of table 2.6.3 implies drivers spurn a saving in the expenditure of effort—an immediate benefit—at the price of lower income—a delayed cost—which is opposite to the expected behavior of an individual with self-control is-

[25] Handel and Kolstad (2015) model information frictions and hassle costs as the source of suboptimal choice.

sues. Moreover, if a driver were to deviate to the optimal labor supply schedule for just one day, the cost of lower income would be negligible since consumption can be smoothed with (dis)saving, which reinforces the argument that self-control is not immediately at play. However, it is plausible that sophisticated drivers with self-control issues may use goal-setting, which is consistent with the results here. Indeed, figure 2.3.1 resembles a labor supply schedule for drivers with reference dependence,[26] which has been used to motivate goal-setting amongst time inconsistent individuals (Hsiaw, 2013). Recent work by Reck and Seibold (2020) explains how decreases in the reference point generally improve welfare because they reduce over consumption, or in this case over work, that takes place to reach a reference point. This is what I observe; drivers work too much at low wages and this is the main cause of welfare losses.

These lessons suggest some remedies for the welfare losses experienced by drivers. Making wages more salient may help purge drivers of their biases in a way similar to Tube strikes. Further, drivers should be trained to avoid using income targets since this causes a negative correlation between the wedge in their intratemporal condition and their labor supply. Hours targets are less harmful, if targets are necessary as a way to avoid self-control issues. Lastly, if over working for low wages due to reference-dependent labor supply is a broader feature of self-employment, this could form part of the argument in favor of a minimum wage for the self-employed in some contexts.

2.7 Conclusion

The approach in this paper builds on previous work that has established non-neoclassical labor supply behavior in many settings, especially amongst New York taxi drivers, in order to ask whether such observations indicate a deviation from optimal behavior and are important for welfare; the answer is yes to both. I characterize preferences and behavior for a group of self-employed workers using salient and common wage variation, respectively, which resolves any normative ambiguity. Typical labor supply is non-monotonic in the wage rate and, for some wage rates, the elasticity of labor supply is more negative than an estimate of the Marshallian elasticity, which indicates labor supply is generally biased. Conversely, salient wage variation due to Tube strikes causes

[26]See figure 1 in DellaVigna (2009).

Figure 2.6.1: Biased Versus Optimal Labor Supply

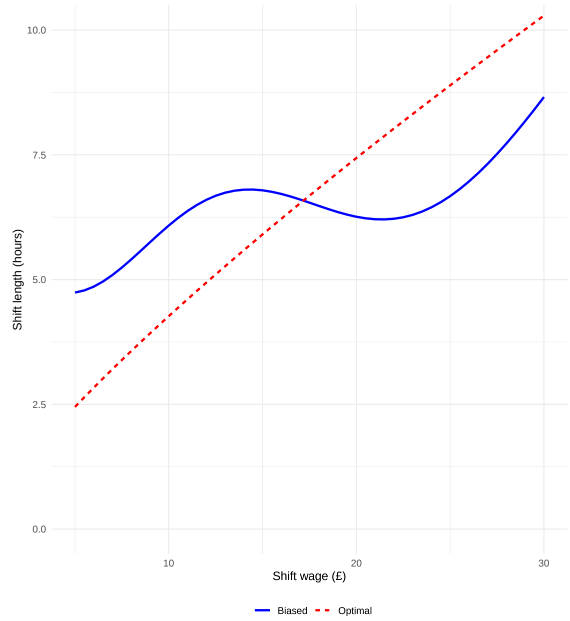

Notes: This figure plots the labor supply function implied by the quadratic control function model in contrast with the labor supply function implied by satisfying the intratemporal optimality condition, where the level parameter is specified by the level of hours during Tube strikes.

a large increase in hours worked that is consistent with Frisch elasticity estimates from the literature. I derive and estimate a theoretical expression to quantify the welfare losses due to behavioral biases, where responses to salient wage variation characterize preferences. This exercise reveals a significant welfare loss for these individuals over the course of a working year. The size of the welfare loss rises steeply as a smaller proportion of drivers are assumed to be biased, and the estimated labor supply function is suggestive of reference-dependence.

Given these findings, it is important to pin down more precisely what biases are afflicting drivers so that it is possible to facilitate choice making which is in line with individuals desires. I see this as a valuable area of future research, which requires experimentation to understand workers' incentives and surveys to understand their motivations. However, even without this knowledge, there are actionable points now. This paper suggests salience is important in affecting optimization, so organizations that contract with self-employed workers could help make predicted wages salient. Moreover, labor supply "training" could inform self-employed workers of the acute inefficiencies due to heuristics such as income targeting. However, it is not clear whether it is in the interest of these organizations to debias drivers. In this regard, this research

Figure 2.6.2: Shift Wage Density

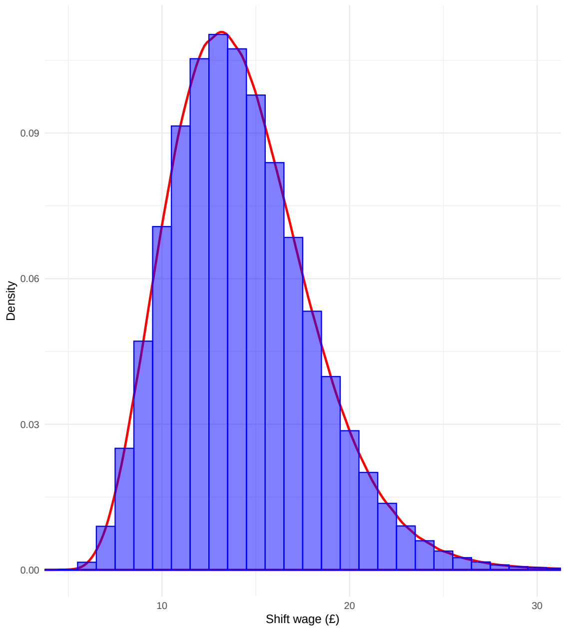

Notes: This figure plots the empirical density of wages in the sample with a fitted kernel density, which is used to integrate over welfare losses.

could provide part of the rationale for a minimum wage for the self-employed with the aim to prevent individuals working for too long at low wages.

Further, this paper has abstracted from extensive margin labor supply decisions because drivers in this setting tend to work a full working week. But in many settings, such as Uber drivers, extensive margin labor supply decisions are much more important and have been used as a key ingredient to determine the value of flexible work. Therefore, better understanding biases at the extensive margin level, and their interaction with intensive margin biases, is an important area of future research.

Appendices

2.A Data Cleaning

New: completed time seems the most reliable measure, so job times are computed by deducting journey time or, if that is not possible, difference with minimum pick up arrival.

At the ride level, I drop observations if,

- The ride was cancelled,
- The start time of the ride is not observed,
- The fare or total transaction value is not observed,
- The ride distance is not observed,
- The duration of the ride is not observed.

Then I trim the data if rides fall above the 2.5 percentile according to the following variables,

- Speed,
- Driving wage,
- Journey distance,
- Ride duration,
- Fare and total transaction value,
- Waiting time,
- Driver extras, *e.g.*, additional fees for toll gates.

I also drop rides if,

- The average speed was below 5 kilometres per hour,
- The driving wage fell below £3.00,
- The journey time was less than one minute,
- The fare or total transaction was less than £1.00.

From these rides I construct shifts by allocating rides that are within six hours of one another to the same shift. At the shift level, I ensure that shifts,

- Contain at least three jobs,
- Do not have more than five jobs an hour,
- Are not shorter than two hours and not longer than 18 hours.

Over the course of this project, I have tried variations on these restrictions and none have significantly impacted the results.

2.B Marshallian Labor Supply Elasticity

I estimate a Marshallian labor supply elasticity using a permanent fare reform that affected drivers' wage rates. The permanent nature of this reform implies that the resulting elasticity incorporates income effects. In the absence of any cross-sectional variation in the application of the reform, identification rests on an exogenous change in the wage that is not confounded by other factors which affect labor supply.

The fare reform took place on June 6^{th} 2016 and its effect on the wage rate is illustrated in figure 2.B.1. After residualizing the wage with respect to controls, as shown in figure 2.B.2, it is clear the wage rose by approximately £1.75, from £14.00 to £15.75. Table 2.B.1 estimates this as a percentage increase of around 10 percent.

Figure 2.B.3 demonstrates how shift hours evolve over the same time horizon. A drop in hours is apparent but, like the wage rate, trends over time complicate inference. Residualizing shift length on controls, as in figure 2.B.4, reveals a clear though small drop in shift length of around five minutes.

The results from the formal estimation of the Marshallian elasticity are shown in table 2.B.2. In this analysis, I regress shift length on the wage rate and controls, and instrument for the wage rate with the fare reform. Controls are analogous, where applicable, to the Tube strike analysis; they include dummies for Tube strikes, factor variables for when the shift was started, a dummy for Ramadan, dummies for bank holidays, and time controls. Given the significant role of trends over time, I test robustness to different specifications of the the time trend polynomial and the length of the window on which the estimation is run. The resulting estimates range from zero to slightly negative—as in Ashenfelter et al. (2010). My preferred estimate, which falls in the midrange of the estimates, is a Marshallian elasticity of -0.142 (s.e., 0.04). When a smaller window is considered, an elasticity much closer to zero is estimated, which is not surprising given the declining trajectory of shift length seen in the raw data. Adjusting the polynomial order of the time trend leads larger and smaller Marshallian elasticities, but none of these alternative estimates are statistically different.

Figure 2.B.1: Raw Marshallian First Stage

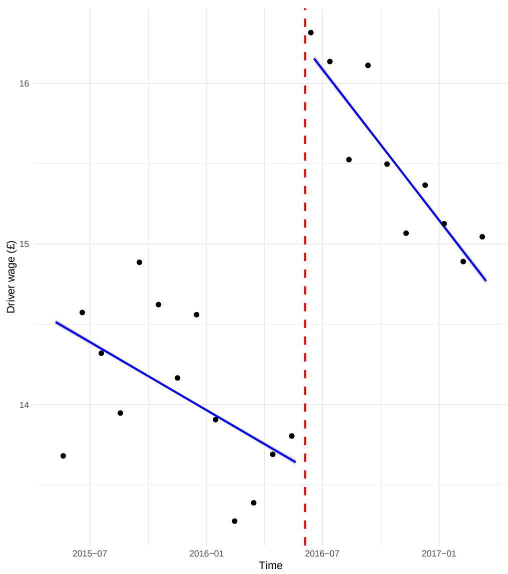

Notes: This figure plots mean monthly wages with a blue line representing a fitted linear trend either side of a fare reform, which is marked with a dashed red line.

2.C Survey Details

Table 2.4.1 presents results from a survey conducted on a separate group of self-employed workers who are predominantly based in London. These individuals work for a variety of ridesharing and food delivery platforms, and the survey was conducted via a firm which provides hire and reward vehicle insurance to these individuals. In order to incentivize completion of the survey, individuals were entered into a lottery for a £50.00 Amazon voucher upon completion. The survey contained 12 questions, including the

Table 2.B.1: True Marshallian Elasticity First Stage

	Dependent variable:						
	log(Shift wage)						
	Full window	3 mo. window	6 mo. window	4th order poly.	5th order poly.	7th order poly.	8th order poly.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Reform	0.114^{***}	0.102^{***}	0.101^{***}	0.105^{***}	0.115^{***}	0.090^{***}	0.082^{***}
	(0.003)	(0.004)	(0.004)	(0.002)	(0.002)	(0.003)	(0.003)
Strike dummies	Yes	No	Yes	Yes	Yes	Yes	Yes
Shift type dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ramadan dummy	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bank holiday dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DOTW dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time trend poly. order	6th	1st	1st	4th	5th	7th	8th
Observations	994,486	140,958	266,279	994,486	994,486	994,486	994,486
R^2	0.114	0.156	0.172	0.113	0.114	0.114	0.115

$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Notes: This table presents first stage results from the influence of a fare reform on drivers wages. The regressions control for driver fixed effects, a factor variables for start time of the shift, Ramadan, bank holidays, days of the week, months, and a time trend of varying flexibility. The regressions either use data from different windows around the reform, or allow for different polynomial time trends. Standard errors are clustered at the driver level.

Table 2.B.2: True Marshallian Elasticity Results

	Dependent variable:						
	log(Shift length)						
	Full window	3 mo. window	6 mo. window	4th order poly.	5th order poly.	7th order poly.	8th order poly.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log(Shift wage)	-0.142^{***}	0.013	0.076	-0.234^{***}	-0.073^{**}	-0.212^{***}	-0.240^{***}
	(0.040)	(0.063)	(0.061)	(0.038)	(0.036)	(0.053)	(0.059)
Strike dummies	Yes	No	Yes	Yes	Yes	Yes	Yes
Shift type dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ramadan dummy	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bank holiday dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DOTW dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time trend poly. order	6th	1st	1st	4th	5th	7th	8th
Observations	994,486	140,958	266,279	994,486	994,486	994,486	994,486
R^2	0.092	0.054	0.044	0.097	0.084	0.097	0.098

$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Notes: This table presents coefficients on wages, which are instrumented with a permanent fare reform, from a series of regressions of shift length on this variable. The regressions control for driver fixed effects, a factor variables for start time of the shift, Ramadan, bank holidays, days of the week, months, and a time trend of varying flexibility. The regressions either use data from different windows around the reform, or allow for different polynomial time trends. Standard errors are clustered at the driver level.

Figure 2.B.2: Residual Marshallian First Stage

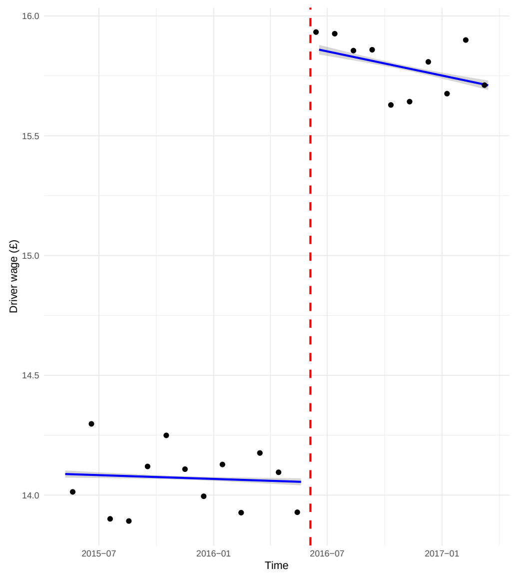

Notes: This figure plots mean residualized monthly wages with a blue line representing a fitted linear trend either side of a fare reform, which is marked with a dashed red line.

question presented here, on general working patterns, wages, and costs. Participants were invited to select one of the responses displayed in table 2.4.1 to the question "How do you decide the number of hours to work in a day?". The ordering of the possible responses was randomized. For the temporary versus permanent wage increase question, drivers were asked "How would you change your daily hours in response to a temporary [permanent] wage increase of 10 percent?". Drivers had the option of reducing their hours, keeping their hours the same, or increasing them. In total, 476 individuals completed this question.

2.D Behavioral Welfare Expression Derivation

In this part of the appendix, I prove theorem 1 and derive an analogous expression for changes in consumer surplus caused by price changes.

2.D.1 Proof of Theorem 1

In order to derive $\Delta(w)$ in terms of sufficient statistics, I make use of the fact that utility is assumed to be additively separable in consumption and hours worked. To start, I will work with consumption utility. I consider a change in consumption induced by

Figure 2.B.3: Raw Marshallian Second Stage

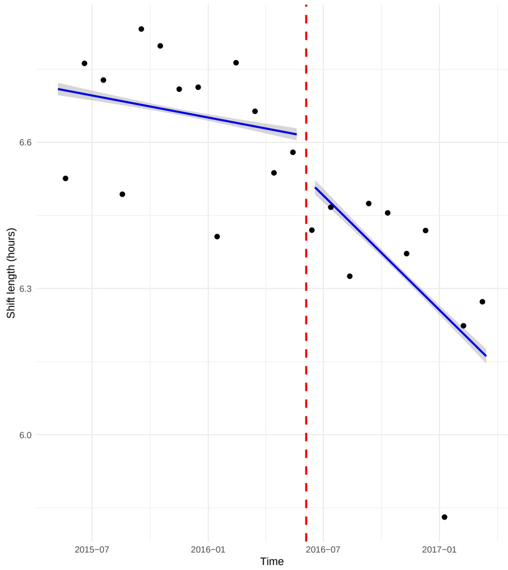

Notes: This figure plots mean monthly shift lengths with a blue line representing a fitted linear trend either side of a fare reform, which is marked with a dashed red line.

switching from a biased consumption rule to an optimal consumption rule at a given wage rate. Mathematically I use a second order Taylor series approximation to show,

$$\begin{aligned} u(c^*(w)) &\approx u\big(\tilde{c}(w)\big) + u'\big(\tilde{c}(w)\big) \cdot \big(c^*(w) - \tilde{c}(w)\big) + \frac{1}{2} \cdot u''\big(\tilde{c}(w)\big) \cdot (c^*(w) - \tilde{c}(w))^2 \\ \iff \frac{u(c^*(w)) - u(\tilde{c}(w))}{u'(\tilde{c}(w))} &\approx c^*(w) - \tilde{c}(w) + \frac{u''(\tilde{c}(w))}{u'(\tilde{c}(w))} \cdot \frac{(c^*(w) - \tilde{c}(w))^2}{2}. \end{aligned}$$

The same operations with hours disutility yield,

$$\frac{v(h^*(w)) - v(\tilde{h}(w))}{u'(\tilde{c}(w))} \approx \frac{v'(\tilde{h}(w)}{u'(\tilde{c}(w))} \cdot (h^*(w) - \tilde{h}(w)) + \frac{v''(\tilde{h}(w))}{u'(\tilde{c}(w))} \cdot \frac{(h^*(w) - \tilde{h}(w))^2}{2}.$$

Figure 2.B.4: Residual Marshallian Second Stage

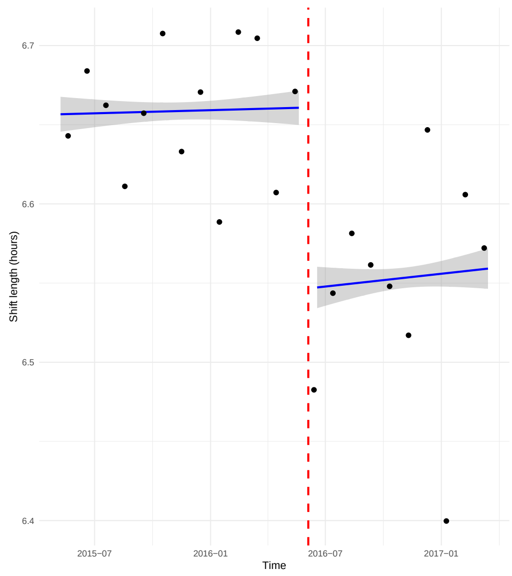

Notes: This figure plots mean residualized monthly shift lengths with a blue line representing a fitted linear trend either side of a fare reform, which is marked with a dashed red line.

Combining these terms gives,

$$\begin{aligned}
\frac{U\big(c^*(w), h^*(w)\big) - U\big(\tilde{c}(w), \tilde{h}(w)\big)}{u'(\tilde{c}(w))} \approx & c^*(w) - \tilde{c}(w) + \frac{u''(\tilde{c}(w))}{u'(\tilde{c}(w))} \cdot \frac{(c^*(w) - \tilde{c}(w))^2}{2} + \dots \\
& \dots \frac{v'(\tilde{h}(w)}{u'(\tilde{c}(w))} \cdot (h^*(w) - \tilde{h}(w)) + \frac{v''(\tilde{h}(w))}{u'(\tilde{c}(w))} \cdot \frac{(h^*(w) - \tilde{h}(w))^2}{2} \\
\approx & c^*(w) - \tilde{c}(w) + \frac{1}{2} \cdot \frac{\tilde{c}(w) \cdot u''(\tilde{c}(w))}{u'(\tilde{c}(w))} \cdot \frac{(c^*(w) - \tilde{c}(w))^2}{\tilde{c}(w)} + \dots \\
& \dots \frac{v'(\tilde{h}(w)}{u'(\tilde{c}(w))} \cdot (h^*(w) - \tilde{h}(w)) + \frac{1}{2} \cdot \frac{v''(\tilde{h}(w))}{u'(\tilde{c}(w))} \cdot (h^*(w) - \tilde{h}(w))^2 \\
\approx & c^*(w) - \tilde{c}(w) + \frac{1}{2} \cdot \frac{\tilde{c}(w) \cdot u''(\tilde{c}(w))}{u'(\tilde{c}(w))} \cdot \frac{(c^*(w) - \tilde{c}(w))^2}{\tilde{c}(w)} + \dots \\
& \dots \frac{v'(\tilde{h}(w)}{u'(\tilde{c}(w))} \cdot (h^*(w) - \tilde{h}(w)) \dots \\
& \dots + \frac{1}{2} \cdot \frac{\tilde{h}(w) \cdot v''(\tilde{h}(w))}{v'(\tilde{c}(w))} \cdot \frac{v'(\tilde{h}(w))}{u'(\tilde{c}(w))} \cdot \frac{(h^*(w) - \tilde{h}(w))^2}{\tilde{h}(w)},
\end{aligned}$$

which is in the same form as equation (2.12) without the simplified notation.

2.D.2 Change in Consumer Surplus Due to Price Change

In this subsection of the appendix, I consider the change in consumer surplus due to a wage change from w' to w but keep the biased policy rules that define behavior constant. Take a second order Taylor approximation of the value function at wage w around the wage w',

$$\begin{aligned} V(w) =& U\big(\tilde{c}(w), \tilde{h}(w)\big) \\ \approx& V(w') + \frac{dV(w')}{dw}(w-w') + \frac{1}{2}\frac{d^2V(w')}{dw^2}(w-w')^2. \end{aligned}$$

It is not possible to apply the Envelope theorem here because of the policy functions do not maximize utility, so,

$$\begin{aligned} \frac{dV(w')}{dw} =& U_c\bigg[\tilde{h}(w') + \bigg(w' + \frac{U_h}{U_c}\bigg)\frac{d\tilde{h}(w')}{dw}\bigg], \\ \frac{d^2V(w')}{dw^2} =& U_c\bigg[\frac{d\tilde{h}(w')}{dw} + \frac{d\tilde{h}(w')}{dw} + w'\frac{d^2\tilde{h}(w')}{dw^2}\bigg] + U_{cc}\bigg[\tilde{h}(w') + w'\frac{d\tilde{h}(w')}{dw}\bigg]^2 \dots \\ & \dots + U_h\frac{d^2\tilde{h}(w')}{dw^2} + U_{hh}\bigg[\frac{d\tilde{h}(w')}{dw}\bigg]^2. \end{aligned}$$

where I have assumed additively separable utility in consumption and hours. Rearranging the Taylor series approximation and substituting in the above yields:

$$\begin{aligned} \frac{V(w)-V(w')}{U_c} =& h\cdot\Delta w + \frac{1}{2}\cdot\Delta h\cdot\Delta w - \frac{1}{2}\cdot\frac{\eta}{c}\cdot(\Delta hw)^2 \dots \\ & \dots + (w - \text{MRS})\cdot\Delta h + \frac{1}{2}\cdot\Delta(w-\text{MRS})\cdot\Delta h, \end{aligned} \tag{2.17}$$

where $\eta = \frac{-cU_{cc}}{U_c}$, $\gamma = \frac{hU_{hh}}{U_h}$, $\text{MRS} = -\frac{U_h}{U_c}$, $h = \tilde{h}(w')$, $\Delta h = \tilde{h}(w) - \tilde{h}(w')$, $\Delta w = w - w'$, $c = \tilde{c}(w')$, $\Delta hw = \tilde{h}(w)w - \tilde{h}(w')w'$, and $\Delta(w - \text{MRS}) = \Delta w - \gamma\cdot\frac{\text{MRS}}{h}\cdot\Delta h$. Relative to the BWE, this expression incorporates the mechanical change due to the price change as well as changes in the internalities. Further if the consumption and hours policy functions were optimal rules, the wedge between the wage and the MRS would always be zero such that equation (2.17) would collapse to the first line.

2.E Parameterization of the Level Parameter

If drivers behave optimally during Tube strikes there are a number of theoretically equivalent ways to identify the level parameter $\theta\cdot c^{\eta}$. However, given the inclusion of econometric errors, which may not be i.i.d., this equivalence breaks down. Therefore, it is important to consider how best to identify the level parameter $\theta\cdot c^{\eta}$. In the below, I consider only observations during Tube strikes and start with the following model,

$$w_{i,t} = (\theta \cdot c^{\eta}) \cdot u_{i,t} \cdot h_{i,t}^{\gamma} + v_{i,t} \tag{2.18}$$

where there are two possible sources of deviation from the intratemporal optimality condition: an additive error $v_{i,t}$ and a multiplicative error $u_{i,t}$, where $\mathbb{E}[v_{i,t}] = 0$ and $\mathbb{E}[u_{i,t}] = 1$. If both $v_{i,t}$ and $u_{i,t}$ were i.i.d., then equation (2.18) would imply,

$$\mathbb{E}\left[\frac{w_{i,t}}{h_{i,t}^{\gamma}}\right] = \frac{\mathbb{E}[w_{i,t}]}{\mathbb{E}\left[h_{i,t}^{\gamma}\right]},$$

which is *not* the case empirically. Therefore, one or more of $v_{i,t}$ and $u_{i,t}$ is not i.i.d. If neither error term is i.i.d. then estimation is near impossible so, practically, the challenge is to determine which source of error is more severely not i.i.d. I argue that the multiplicative error $u_{i,t}$ term is most likely to be related to hours $h_{i,t}$ since $u_{i,t}$ intuitively captures idiosyncratic variation in consumption and the distribution of hours will be linked to the distribution of income, which is related to consumption c and, in turn, the level parameter $\theta \cdot c^{\eta}$. As a result, I use the following estimator,

$$\mathbb{E}\left[\frac{w_{i,t}}{h_{i,t}^{\gamma}}\right] = (\theta \cdot c^{\eta}) \cdot \underbrace{\mathbb{E}[u_{i,t}]}_{=1} + \underbrace{\mathbb{E}[v_{i,t} \cdot h_{i,t}^{-\gamma}]}_{=\mathbb{E}[v_{i,t}]\cdot\mathbb{E}[h_{i,t}^{-\gamma}]=0} .$$

I note that the estimate generated by this estimator is consistent with the assumption made on biases, *i.e.*, that the logarithm of drivers' first order condition holds in expectation, though the assumption itself cannot be used to estimate the level parameter $\theta \cdot c^{\eta}$ because of Jensen's inequality.

Chapter 3

Refinancing Cross-Subsidies in the Mortgage Market

3.1 Introduction

The inherent complexity of finance interacts with the behavioral tendencies of ordinary people so that they often pay insufficient attention to managing particular financial products once purchased. The resulting delays in promptly taking action in response to financial incentives can result in such customers unwittingly providing revenues to financial firms. In contrast, such products can be beneficial to more sophisticated customers who exploit opportunities to use them to their own advantage. This can result in regressive cross-subsidies in financial markets that flow from less sophisticated customers, who are often poorer and less educated, to those who are more sophisticated, wealthy, or educated. In this way, the design of household finance products can be a powerful contributor to wealth inequality.

In this paper, we provide a new approach to quantify household finance cross-subsidies, and to identify how they are distributed across the population. We take up this challenge in the setting of residential mortgage refinancing. Mortgages are typically the largest household financial liability (Campbell, 2006; Badarinza et al., 2016; Goetzmann et al., 2021), but despite their importance in household budgets, many households do not appropriately manage this component of their balance sheets. An important component of mortgage management is prompt refinancing in response to financial incentives to act, and evidence has built up that poorer and less-educated households fall short on this dimension (Agarwal et al., 2016; Keys et al., 2016; Andersen et al., 2020).

To assess cross-subsidies in mortgage refinancing, we build a structural model that we fit to high-quality administrative data on the stock of all outstanding mortgages in the United Kingdom in 2015. In the UK, as in many other countries, the dominant mortgage form is a "discounted rate" instrument with a relatively short initial fixation period. To fully take advantage of this discounted or "teaser" rate, it is imperative to promptly refinance at the point at which the initial fixation period ends, to avoid being rolled on to a significantly more expensive "reset rate."[1] Households who fail to promptly refinance pay these higher reset rates, and this contribution to lender profits allows in turn for discounted rates to be lower in equilibrium. Households who are swift to refinance take advantage of these lower discounted rates.

Our approach to uncovering cross-subsidies in this setting is to use the estimated structural parameters to consider a counterfactual scenario in which all households face a simpler contract, which pays a single rate until mortgage maturity, and requires no refinancing. We compare outcomes experienced by households of different income levels and those located in different regions of the UK in the prevailing mortgage contract design to those in the counterfactual simpler contract design. By doing so, we make the "invisible" cross-subsidies that exist in the current system visible. This is an approach that can be more widely applied in other market settings.

Our empirical analysis exploits rich data from the UK's Financial Conduct Authority (FCA). These data are particularly well-suited for our purposes for several reasons. First, they are both granular and comprehensive, tracking individual mortgages in the stock of all outstanding UK mortgage loans issued by all regulated financial institutions in the country at a semi-annual frequency between 2015H1 and 2017H2 (we mainly focus on the stock in 2015H1).[2] The granular nature of the data means that we observe household-level mortgage refinancing behavior; and the comprehensive coverage of the entire mortgage stock facilitates our cross-subsidy calculations. Second, the refinancing decision in the UK is a simple choice between getting back into a discounted rate contract versus paying the substantially higher reset rate, which, as we demonstrate, is a dominated choice given sensible parameter values, even when accounting

[1] This feature of the UK mortgage market has prompted prominent calls for reform which highlight the implicit cross-subsidy (Miles, 2004).

[2] In what follows, we denote the first and second observations in each year of our sample by H1 and H2 respectively to denote "half-years".

for any option value of waiting to lock in an advantageous rate.[3] Third, subject to collateral re-verification, many UK mortgages are portable, which reduces concerns about the effect of households' unobservable moving propensities, a common determinant of refinancing.

In 2015H1, the total stock of household mortgage debt in our sample equals £470 billion. The majority of this stock (69.8%) pays the discounted rate, but the remaining large share (30.2%) pays the reset rate, with an average rate spread of 52 basis points (bps). These statistics are revealing: the spread between reset rates and discounted rates, combined with the fractions of balances that pay these different rates jointly contribute to lender revenues in the "cross-subsidy equilibrium." We use these and other data moments to capture multiple dimensions of cross-sectional and time-series variation in borrower refinancing behavior.

We match these moments using a rich structural model of the UK mortgage market. The model assumes that households are heterogeneous along two dimensions. The first dimension is household valuation for owned housing (we model renting as an outside option). The second dimension is household inaction, which features a persistent household-specific component as well as a time-varying idiosyncratic shock. At the point of deciding whether to buy a house and, if they do so, their mortgage size, households have noisy information about their future costs to promptly refinance. We consider variants of the model in which households have different degrees of precision about their future refinancing costs at this decision point.

When the discounted rate expires, households in the model choose to refinance when the benefits of refinancing, primarily driven by the difference between the discounted and the reset rates as well as loan size, outweigh the fixed costs of refinancing. Larger loans are therefore more likely to pay discounted rates: in the cross-section, these loans correspond to households with greater valuation for housing; in the time-series, these loans correspond to households who recently originated their mortgages.

The model facilitates easy aggregation of loans, allowing us to write down intuitive expressions for aggregate mortgage loan balances on the discounted rate and on the reset rate. We estimate the model parameters assuming that the market is in steady state in 2015H1, and match the data well. Our estimates imply that average refinancing costs

[3]This feature eliminates reliance on an auxiliary model of the real option, such as that of Agarwal et al. (2013b) and allows us to recover households' implicit refinancing costs directly.

equal £$3,866$ among mortgage borrowers, with a standard deviation equal to £$6,641$.

The estimated parameters are the main input into our cross-subsidy calculations, which result from comparing the single-rate counterfactual economy with the dual-rate baseline economy—the single rate eliminates rate differences across households arising from their differential refinancing activities. In these counterfactual scenarios, households adjust both their individual loan sizes (intensive margin) as well as their participation in the housing and mortgage markets (extensive margin), in response to the different paths of mortgage rates and to the elimination of refinancing costs in the single-rate counterfactual economies relative to the baseline dual-rate economy. The size of these adjustments depends on how well-informed households are about their own persistent refinancing costs (i.e., their tendency towards inaction). The less "self-knowledge" that households have, the smaller their mortgage size and mortgage market participation adjustments tend to be. We find that even with an appreciable degree of noise in the information about persistent refinancing costs there are significant adjustments in the counterfactuals.

In the counterfactual single-rate economy with a rate equal to the weighted-average rate in the sample,[4] the total number of mortgages increases by 6.49 percent, because high-refinancing-cost households, who no longer pay either the punitive reset rate or refinancing costs, are more likely to enter the market. However, the mean initial loan balance falls in the counterfactual equilibrium by 2.37 percent of the baseline average loan size, because the composition of borrowers changes: marginal households who enter the mortgage market in the single-rate economy have smaller loan sizes than inframarginal households whose participation does not change.

We also estimate two extended versions of the model to assess how cross-subsidies are distributed across the population of mortgage borrowers. The first version estimates parameters separately for different geographical regions in the UK, to capture regional heterogeneity in preferences. The second estimates parameters separately for 12 income groups (bottom-eight income deciles, and the top-two deciles each additionally split into two sub-groups). These extended models continue to match the aggregate moments very well, but also feature considerable differences in refinancing costs across regions and income groups. This is a harbinger of significant regional variation as well

[4] We consider a range of single rates in the counterfactual economy, mainly focusing on a rate equal to the balance-weighted average rate in the 2015H1 sample.

as variation across the income-distribution in the cross-subsidies that are paid and received.

More specifically, we find clear evidence that higher-income households, and households in the richer South-West of the country would pay higher rates under the single-rate structure, and households in the relatively poorer North-East and North-West of the country would pay lower mortgage rates under the counterfactual single-rate scenario than they do in the dual-rate economy.

Finally, we find interesting variations in the endogenous response of different groups' mortgage takeup and mortgage sizes in the counterfactual single-rate economy. There is a shrinking of average mortgage debt for higher-income groups and wealthier UK regions in response to the higher single rate they pay in the counterfactual, since they no longer have access to the discounted rate. In contrast, the counterfactual single-interest rate scenario, which does not require refinancing, induces lower-income households to enter the mortgage market because they expect to pay lower rates and incur no refinancing costs. This is evident in increases in the home-ownership rate, mainly driven by low-income households. This "democratization" of mortgage takeup under the counterfactual is another important indicator of the effect of regressive cross-subsidies in the dual-rate economy.

3.1.1 Related Literature

Our paper contributes to several strands of the literature. First, our work complements many empirical papers that document switching costs, inertia, and inattention in insurance and household finance markets, such as health insurance (e.g., Handel, 2013), car insurance (e.g., Honka, 2014), retirement plans (e.g., Luco, 2019; Illanes, 2016), credit cards (e.g., Ausubel, 1991; Stango and Zinman, 2016; Nelson, 2022), pension contributions (e.g., Choi et al., 2002), and portfolio rebalancing (e.g., Brunnermeier and Nagel, 2008) among others.[5] However, none of these papers focuses on documenting regressive cross-subsidies, though this possibility has been raised in theory (e.g., Gabaix and Laibson, 2006b; Armstrong and Vickers, 2012).

The papers that document inaction and frictions in mortgage refinancing (e.g., Agar-

[5]Farrell and Klemperer (2007) present a survey of the literature on switching costs, with a theoretical focus; Heidhues and Kőszegi (2018) survey the literature on behavioral industrial organization, and Gavazza and Lizzeri (2021) the literature on markets with frictions.

wal et al., 2016; Keys et al., 2016; Scharfstein and Sunderam, 2016; DeFusco and Mondragon, 2020; Byrne et al., 2023; Berger et al., 2023) are more directly related to our work. We advance this literature, developing a novel framework for refinancing inaction that allows us to quantify the magnitudes of cross-subsidies across households through counterfactual analysis. This is a different approach to Andersen et al. (2020), who model a fixed refinancing cost ("state-dependent inaction"), but with intervals of "time-dependent inaction" where refinancing is not possible, using a periodic "Calvo" shock to borrowers, and Berger et al. (2021), who adopt a similar approach in their analysis of US refinancing behaviour. These approaches imply that the costs of refinancing are always higher than the benefits during periods of time-dependent inaction, but do not quantify these costs. In contrast, our model features a household-specific fixed refinancing cost with a time-varying shock, meaning that our structural estimation recovers the full distribution of the costs of inaction across households and over time. Apart from the differences in setting, this different modelling approach explains why the average refinancing costs that we estimate are modestly higher than those in Andersen et al. (2020) and Berger et al. (2021).[6]

Second, our paper is connected to a growing body of work on the design of mortgage markets around the world (Campbell, 2013; Piskorski and Seru, 2018). For example, several mortgage markets also feature rates that are fixed for a shorter interval than the maturity of the mortgage, and Allen and Li (2020) study borrower refinancing and lender mortgage pricing in this setting; similarly, Thiel (2021) studies a ban on price discrimination between new and existing customers in the Dutch mortgage market. We focus on implicit cross-subsidies across borrowers in the cross-section, whereas Allen and Li (2020) and Thiel (2021) focus on intertemporal price discrimination within borrowers in the Canadian and Dutch markets, respectively.[7]

Finally, our structural model provides a money-metric assessment of cross-subsidies in an important household finance market, and shows that these cross-subsidies are regressive. This showcases how the design of the financial system can contribute to in-

[6] Andersen et al. (2020) estimate an average total psychological plus fixed refinancing cost of £$1,852$ in the Danish mortgage market. Berger et al. (2021) estimate an average refinancing cost of $1,934 in the US mortgage market. These are slightly lower than our estimate of the average cost across both refinancing and non-refinancing borrowers, which equals $\pounds 3,866$.

[7] Our paper also contributes to the growing literature on UK mortgage markets, including Benetton (2021); Robles-Garcia (2022); Cloyne et al. (2019); Best et al. (2020); Belgibayeva et al. (2020); Benetton et al. (2021); Liu (2022). Most of these studies focus on the flow of newly originated mortgages, whereas we focus on the stock of mortgages.

equality, connecting our work to the growing literature on wealth inequality (Alvaredo et al., 2017; Benhabib and Bisin, 2018; Fagereng et al., 2020; Hubmer et al., 2020), and more specifically, those that contribute to inequality in financial wealth (Campbell et al., 2019; Greenwald et al., 2021). In the process, we document that cross-subsidies vary across regions and devolved administrations of the UK, which shows that regional redistribution can occur directly as a result of differential efficiency in the use of financial products. These results speak to the literature on regional redistribution in housing and mortgage markets (Hurst et al., 2016; Beraja et al., 2019).

3.2 Data and Institutional Setting

Our primary data source is the FCA, which comprehensively tracks the stock of outstanding mortgage loans issued by all regulated financial institutions in the UK. The specific FCA dataset that we use is the Product Sales Database 007 (henceforth PSD007), which reports information about the stock of mortgage loans between June 2015 (henceforth 2015H1), and December 2017 (2017H2) at a semi-annual frequency.[8]

At each reporting date, PSD007 records the original loan amount, outstanding balance, original loan term, remaining term to maturity, current interest rate, current monthly payment, and performance status (i.e., whether the loan is in arrears and if so, for how long this has been the case) for each outstanding mortgage. The database also includes information on the property location at the most granular level in the UK (6-digit postcode), and borrower characteristics such as date of birth and the opening date for the bank account associated with the mortgage. Table 3.A.1 in the appendix provides more detailed descriptions of the main variables from the PSD007 dataset used in this paper.

The PSD007 dataset does not include information on the evolution of borrower incomes, as these are typically reported at origination. We therefore merge borrowers in the stock data with comprehensive loan-level data on borrower characteristics shared with lenders at the time of loan origination. We also measure the current loan-to-value (LTV) ratio on each outstanding loan following a common approach in the literature, di-

[8] Regulated financial institutions in the UK are legally required to report these details within 30 working days following the end of each calendar half-year. The group of regulated financial institutions in the UK includes deposit-taking institutions (including building societies), as well as some non-bank financial institutions. Our sample focuses on the owner-occupier segment of the mortgage borrowing population, and excludes "buy-to-let" mortgages which are issued mainly to landlords on rental properties.

viding the outstanding loan balance by the scaled house price at mortgage origination, using Local Authority district-level house price indices. Appendix 3.A.2 provides details of the procedure used to merge borrower and house characteristics at loan origination to our stock data.

We further complement the PSD007 dataset with data on UK homeownership rates sourced from the Office for National Statistics (ONS) dataset *Dwelling stock by tenure*. These homeownership data allow us to measure households' extensive margin decision of whether to buy a house and take a mortgage, or rent.

Using rich data on the stock of mortgages offers several advantages over using the flow of originations. Notably, the stock allows us to accurately capture refinancing behavior across all mortgage maturities, including mortgages that were originated in the past. Moreover, the parameters of a structural model estimated using the stock of mortgages rather than using the flow depend less on changes in refinancing behavior or refinancing waves over short periods of time. Finally, using the mortgage stock facilitates computing average mortgage rates and aggregate lender revenues, which proves useful in our counterfactual analyses.

3.2.1 UK Mortgage Market: Institutional Features

Our work exploits a few key features of UK mortgage markets apparent in our PSD dataset.

First, the UK mortgage market features posted prices at the national level, with no variation across regions, as Cloyne et al. (2019), Benetton (2021), Robles-Garcia (2022), and Benetton et al. (2021) (among others) document. Borrower-specific pricing, common in the US mortgage market, is virtually non-existent in the UK market.

Second (and crucial for our purposes), the vast majority of UK mortgages are issued with discounted interest rates which are fixed for a set time period, usually between one and five years (the modal fixation period is two years), depending on the contract chosen by the borrower. During the discounted period, households typically incur substantial prepayment penalties (between 3-5% of the loan balance), which means that households typically refinance after the end of the fixed period (Cloyne et al., 2019; Belgibayeva et al., 2020). At the end of the discounted period, the mortgage rate automatically rolls over into a higher reset rate known as the "standard variable rate," unless borrowers choose to refinance the mortgage into another discounted rate (for a detailed treatment of the characteristics of the UK mortgage market see Miles, 2004).[9]

This "dual-rate" structure is a feature of many mortgage systems, including Canada, Australia, India, Ireland, Germany, and Spain, meaning that our study is more broadly applicable around the world.[10] We do not study the origins of this rate structure, which likely reflects mortgage lenders' funding structures and price-discrimination strategies between active and inactive borrowers (Ellison, 2005; Gabaix and Laibson, 2006b; DellaVigna and Malmendier, 2006), and we focus instead on its implications for borrowers' refinancing. That said, in Appendix 3.D, we follow and extend the analysis of Cloyne et al. (2019), who perform a thorough comparison between borrowers who pay the discounted and the reset rates, suggesting that the dual-rate structure does not seem designed for lenders to screen borrowers based on their default risk.[11]

This dual-rate contract structure provides strong incentives for households to refinance at the expiration of the fixation period. UK households are free to take advantage of these incentives to refinance, as there are no further credit checks when households refinance with their existing lender, and any upfront fees can be rolled into the loan balance, meaning that liquidity constraints do not inhibit refinancing (Best et al., 2020). In Appendix 3.E, we also rule out the possibility that borrowers rationally stay on the reset rate to exploit the real option of timing their refinancing to coincide with interest rate declines.

Third, an additional feature of the UK setting is that mortgages are portable, meaning that households can retain their existing mortgage contract when they move, subject to the new collateral being re-verified.[12] This feature stands in contrast with the US, where the lack of portability means that moving probabilities are a potentially more important driver of both prepayment/refinancing and contract choice (Stanton, 1995; Stanton and Wallace, 1998; Zhang, 2022).

[9]There is a third type of interest rate known as a tracker rate, paid on around 15% of all mortgages outstanding, which is a floating rate linked to the Bank of England base rate. We exclude such mortgages from our analysis since such mortgages are subject to rate fluctuations and there are rarely transitions from the rest and discounted rate category into this category. Further details in appendix 3.A.1.

[10]Badarinza et al. (2018) provide information on mortgage interest-rate fixation periods across a broad set of countries and show that many large economies have similar average mortgage-rate fixation periods to the UK.

[11]This pricing structure with a discount for new or active customers is quite common in many other retail markets, including electricity, telecoms, and magazines, in which default concerns play a negligible role.

[12]Among other countries, Australia, Canada, and Germany share this feature (Lea, 2010).

3.2.2 Borrowers Ineligible to Refinance

One potential challenge to our empirical analysis and our cross-subsidy calculations is to distinguish between households who can refinance, but do not do so promptly, from households who are constrained and unable to take advantage of refinancing opportunities. To address this potential confounding effect, we filter our data to remove borrowers who are potentially ineligible for refinancing—i.e., borrowers who are "involuntarily" on the reset rate, but who would potentially like to switch if they were allowed to do so.

To identify these ineligible borrowers, we follow studies by the FCA (Financial Conduct Authority, 2019b, 2021) and a 2018 industry agreement that unified and codified refinancing eligibility criteria across major UK lenders (65 lenders, with a market share of around 95%). Passing these eligibility criteria means that a mortgage borrower can refinance into a new contract with their lender, without any affordability assessment, meaning no additional credit or income checks.[13] The criteria are that the borrowers are first-charge owner-occupiers that are existing borrowers of an active lender, up to date with their payments, with a minimum remaining term of 2 years, and a minimum outstanding balance of £10,000 (Financial Conduct Authority, 2019b). We broaden out these eligibility criteria to filter out borrowers that are potentially ineligible for refinancing, under the assumption that the 2018 agreement ratified pre-existing practice that was prevalent in the 2015H1 stock.

Appendix Table 3.A.6 shows the exact proportion of loans that are potentially ineligible for refinancing using these criteria as well as broader definitions of ineligibility. Borrowers who have very high LTVs greater than 95% comprise approximately 2% of the sample. These borrowers may find it difficult to refinance, even though they are strictly eligible under the industry agreement if they fulfill all the other criteria. Borrowers with small remaining loan balances (loans smaller than £30,000) constitute approximately 6% of the total sample. And about 5% of loans are non-performing (in arrears, or under forbearance or possession orders). Applying these filters together removes around 14% of the mortgage stock in 2015H1-2017H2.

[13] The UK is somewhere between the US and Denmark in this respect. In the US, a credit check is triggered at the point of refinancing (Keys et al., 2016), whereas in Denmark, even delinquent borrowers are able to refinance as long as there is no cash out (Andersen et al., 2020). The UK system does not trigger a credit check at the point of refinancing as long as the borrower satisfies the eligibility criteria.

We note here that we estimated our model on both unfiltered and filtered samples. Filtering does not materially affect our main qualitative results on the regressive nature of cross-subsidies, for two main reasons. First, in the filtered sample the share of mortgage debt paying the reset rate is still quite large, and lower-income borrowers are more likely to pay the reset rate than higher-income borrowers. Second, the largest fraction of excluded borrowers are those with small loan balances, for whom refinancing benefits are small, because the refinancing benefit is proportional to the loan balance. Appendix 3.A.4 provides more information on these filtered borrowers.[14]

3.2.3 Summary Statistics of the Mortgage Stock

Our analysis focuses on the 2015H1 mortgage stock, which comprises 3.59 million mortgages of borrowers eligible to refinance, and for whom we have estimates of current income.[15]

Table 3.2.1 shows summary statistics for selected variables in this filtered 2015H1 sample. On average, the mean outstanding balance equal £130,871 and a mean loan balance at origination of £142,333 (the difference is attributable mainly to amortization). This aggregates to a total stock of outstanding mortgage debt of £470 billion.

Taking an equal-weighted average across all mortgages, Table 3.2.1 shows that they pay an average interest rate of 3.46% at the end of 2015H1, at a spread of 2.79% over maturity-matched UK Treasuries, and have a remaining term to maturity of 19.3 years on average.[16] 65.0% of the 3.59 million mortgages pay discounted rates in this 2015H1 sample, with an average equal-weighted remaining discounted period of 2.1 years.

[14] A 2018 FCA report of the mortgage market (Financial Conduct Authority, 2019b) studied 2 million reset rate mortgages using the same data that we employ, and concluded that only approximately 30,000 of these mortgages were unable to switch despite being up to date with payments. The report finds that two-thirds of these mortgages were associated with an inactive, failed lender (e.g., Northern Rock, famously subject to a run during the financial crisis); and the remainder were either interest-only mortgages that were subject to changes in lending standards following the financial crisis, or in negative home equity. We expect that our filters catch many of these mortgages.

[15] The main statistics of the mortgage stock are quite stable between 2015H1 and 2017H2, consistent with the idea that short-run changes have small effects on the stock of long-term debt contracts. Appendix 3.B describes the evolution of the mortgage stock between 2015H1 and 2017H2, which exhibits two main patterns: (1) the fraction of mortgage debt paying the reset rate decreases by 2017H2, and (2) the spread between the average reset rate and the discounted rate increases over the same period. While the first pattern should decrease the magnitude of cross-subsidies across borrowers, the second one should increase them, with a small net effect.

[16] Mortgage spreads are computed with respect to the yield on a nominal zero coupon UK Treasury with maturity matched to the mortgage interest rate fixation period. We use the short-term interest rate for mortgages paying the reset rate. For instance, for a mortgage with t years of fixation, the spread is calculated by subtracting off the spot rate for a UK treasury maturing in t years as at the reporting date.

Table 3.2.1: Summary Statistics for the Mortgage Stock in 2015H1

	MEAN	SD	P25	P50	P75
CURRENT LOAN BALANCE (£)	130,871	97,858	71,450	106,967	158,495
CURRENT INTEREST RATE (IN PP)	3.46	0.98	2.53	3.38	4.00
SPREAD TO T-BILL (IN PP)	2.79	1.05	2.05	2.53	3.54
ORIGINAL LOAN BALANCE (£)	142,333	100,661	82,000	118,399	170,000
ORIGINAL TERM (IN YEARS)	23.32	7.07	19.00	25.00	28.00
REMAINING TERM (IN YEARS)	19.27	7.67	13.92	19.00	24.33
REMAINING DISCOUNTED PERIOD	2.11	1.52	1.00	1.83	3.08
BORROWER AGE	41.97	10.02	34.00	41.00	49.00

Notes: The table above shows summary statistics of mortgages from the stock data reported in 2015H1. The sample includes mortgages in two categories, namely, those paying discounted interest rates, and those paying the Standard Variable Rate. The total sample comprises around 3.59 million mortgages, of which 65.0% are discounted rate mortgages at this point in time. Appendix Table 3.A.1 contains a description of the underlying variables.

Table 3.2.1 also reveals considerable cross-sectional variation in these variables. Such variation is particularly evident in the outstanding loan balance and the remaining mortgage term. When the outstanding loan balance and/or the remaining term are low, borrowers should be less likely to refinance given the lower financial incentive from any interest rate reduction associated with doing so. The mortgages also vary in terms of the overall interest rate they pay, as well as their spread over the maturity-matched UK treasury rate. There is also demographic variation evident in the age of borrowers, with both relatively young borrowers, aged 34 at the 25th percentile of the cross-sectional distribution, and older borrowers, aged 49 at the 75th percentile of the cross-sectional distribution represented in the mortgage stock.

Two additional statistics not reported in Table 3.2.1 constitute important targets for our model. First, the ONS dwelling data report that 63% of households are homeowners in 2015. Second, rate types are highly persistent. Figure B.2 in Appendix 3.B displays the transition probabilities between discounted and reset rates in the PSD007 dataset. Households on a reset or discounted rate are much more likely to stay on the same rate type over the next 24 months than to switch—i.e., 76.2% of 2015H1 borrowers pay the same rate type in 2017H1. However, some borrowers do switch over time. Switches from the lower discounted to the higher reset rate may reflect inattention and inertia, or more generally, refinancing costs, while refinancing benefits decline as the loans amortize. Against this backdrop, switches from the reset to the discounted rate suggest that these costs vary over time within households. Taken together, these switching patterns suggest that a combination of household-specific fixed refinancing costs and

Table 3.2.2: Summary Statistics for the Mortgage Stock in 2015H1, by Income Quantiles

QUANTILES	INCOME (£)	HOMEOWNERS (%)	BALANCE (£)	DISCOUNTED (%)	SPREAD (RESET-DISCOUNTED)
0-10	24,604	0.50	60,144	0.66	0.53
10-20	29,483	0.61	73,839	0.64	0.45
20-30	34,564	0.64	84,721	0.64	0.42
30-40	39,581	0.68	94,547	0.64	0.40
40-50	44,986	0.72	104,950	0.64	0.39
50-60	51,327	0.75	116,473	0.64	0.39
60-70	59,412	0.80	130,123	0.64	0.39
70-80	71,261	0.82	149,041	0.66	0.40
80-85	80,290	0.84	169,791	0.66	0.44
85-90	94,142	0.86	190,849	0.67	0.49
90-95	122,708	0.91	227,788	0.68	0.55
95-100	214,486	0.96	345,904	0.69	0.64

Notes: The table above shows summary statistics of mortgages from the stock data reported in 2015H1, split by income quantiles of borrowers. Appendix Table 3.A.1 contains a description of the underlying variables.

time-varying stochastic shocks may capture the high persistence of rate types, as well as the occasional switches across rate types over time.

Table 3.2.2 shows summary statistics across quantiles of the income distribution of borrowers; this is an important dimension along which we later evaluate cross-subsidies. The third column of the table shows that the homeownership rate rises monotonically with the level of income—it equals 50% in lowest-income group and attains 96% in the highest-income group. The remaining columns refer to borrowers. Their loan balance increases with their income, as expected. More importantly, the share of mortgages on the discounted rate (fifth column) also tends to increase with borrower income, whereas the spread between reset and discounted rates (sixth column) is broadly of a similar magnitude across groups. These patterns document that lower-income borrowers are less likely to refinance than higher-income borrowers, hinting at the likely direction of cross-subsidies. Table C.1 in Appendix 3.C provides a similar table across UK regions, confirming that borrowers in higher-income regions are more likely to pay discounted rates than those in lower-income regions.

Overall, the summary statistics reported in Tables 3.2.1 and 3.2.2 document that the UK mortgage market comprises a mix of borrowers paying discounted rates and reset rates. While mortgages on discounted rates constitute the main share, a large fraction (30.2% by loan balance) of the outstanding mortgage stock pays the reset rate, at a spread of 52 bps over the equivalent discounted rate. Our dataset includes mortgages by two large lenders who offered to cap reset rates at 250 bp for mortgages issued up to and during the 2007-09 financial crisis. Excluding these lenders (around 900k observations) pushes up the average rate for reset rate mortgages substantially (with no change in the average rate for discounted mortgages), increasing the spread to 110 bp. We have kept mortgages by these two large lenders in our sample to provide conservative cross-subsidy estimates.

In the next Section, we develop a model that we map to these data features in our structural estimation; we use the model to quantitatively assess the magnitude of the cross-subsidy that the dual-rate structure embeds.

3.3 Model

We model a mortgage market in which a measure M of households enters in each period. When they enter the market, households choose whether to buy a house with a mortgage or rent a property. If a household i chooses to buy, they pay a one-time origination cost k_i^o and obtain per-period flow utility from their house equal to $v_i h_i^{\alpha} - m(l_i, r, T)$, where v_i is household i's per-period valuation for housing, h_i is the size of the house that the household i chooses, and $0 < \alpha < 1$ is a parameter governing the utility from housing. $m(l_i, r, s)$ is the per-period mortgage payment of a household with a mortgage with current loan balance l_i, generic interest rate r, and remaining term s, which follows from the amortization of the loan:

$$m(l_i, r, s) = l_i \frac{r(1+r)^s}{(1+r)^s - 1}. \tag{3.1}$$

Renting a property yields per-period utility $\bar{u}$, which we assume is common to all households and fixed over time. All households discount the future at the common rate β.

Mortgages are long-term contracts for T periods that pay a discounted rate r_d for an initial time interval T_d, and subsequently pay a reset rate $R > r_d$ following this interval, unless the household refinances back into the discounted rate. To simplify and facilitate evaluating counterfactuals, we take both rates as given constant values. We also assume that T/T_d is a (positive) integer and that households can only refinance at the point at which the discounted rate expires. In what follows, we normalize by the length of this initial fixation period, treating it as a single time unit, i.e., we assume $T_d = 1$ and $T = 15$, and all rates are computed over the period T_d. Moreover, we assume that households do not change their loan balance (i.e., we rule out "cash-out refinancing"), and rule out maturity extensions (i.e., households in the model do not change the maturity of their loan at the point of refinancing). Households receive the loan amount at time $t = 0$, but make the first repayment at $t = 1$, which is also the first refinancing period. Hence, the loan balance of a mortgage with interest rate r evolves over time as follows:

$$l_{i,t+1}(r, l_{i,t}) = l_{i,t}(1 + r) - m(l_{i,t}, r, s). \tag{3.2}$$

Mortgages are fully repaid after T periods. Thus, each household makes T payments over the life of the loan, the same as the duration of the mortgage contract.

At time $t = 0$, if they choose to buy a house, households choose the size of their mortgage loan $l_{i,0}$ to finance their house h_i, where $\omega_i = h_i/l_{i,0}$ denotes the inverse of the loan-to-value at origination. In each subsequent period, households can refinance their mortgage at the discounted rate r_d; to do so, they have to pay refinancing costs equal to $k_{i,t} = k_i \varepsilon_{i,t}$, where k_i is a persistent component of the refinancing cost for household i and $\varepsilon_{i,t}$ is a transitory component. We assume that $\varepsilon_{i,t}$ is a non-negative random variable, independent and identically distributed across households and over time, with mean equal to one, with cumulative distribution function $F(\varepsilon_{i,t})$ and density $f(\varepsilon_{i,t})$. Hence, each household's average refinancing costs equal their persistent component of refinancing costs, i.e., $E(k_{i,t}) = k_i$.

Households are heterogeneous in their per-period valuation for housing v_i (to capture the heterogeneity of initially chosen loan sizes seen in the data) and in their persistent component k_i of the cost of refinancing (to capture the household heterogeneity in refinancing for a given loan balance). We assume that, at the time of originating a mortgage, households perfectly know their valuation for owned housing v_i, but only receive a signal of their persistent component k_i of refinancing costs and thus of their average refinancing costs over time. Specifically, we assume that k_i is correlated with the origination cost k_i^o according to $k_i = k_i^o \varepsilon_{i,0}$, where $\varepsilon_{i,0}$ is a non-negative random variable that is realized after the origination of the mortgage and before the first refinancing opportunity. Thus, the precision of the signal negatively depends on the variance of $\varepsilon_{i,0}$. We assume that $\varepsilon_{i,0}$ is independent and identically distributed across households, with mean equal to one, with cumulative distribution function $F_0(\varepsilon_{i,0})$ and density $f_0(\varepsilon_{i,0})$.

Valuations and origination costs are distributed according to the cumulative joint distribution function $G_o(v_i, k_i^o)$ with density $g_o(v_i, k_i^o)$. Hence, the joint density of valuations and persistent refinancing costs equals $g(v_i, k_i) = \int_0^{+\infty} g_o(v_i, k_i/\varepsilon_{i,0}) f_0(\varepsilon_{i,0}) \frac{1}{\varepsilon_{i,0}} d\varepsilon_{i,0}$.

Intuitively, in the model, households learn about their persistent ongoing mortgage refinancing costs from the costs/hassle that they experience during the process of mortgage origination. The extent to which this initial origination process is informative about ongoing mortgage refinancing costs is dictated by the variance of $\varepsilon_{i,0}$. If this variance is zero, the initial process of mortgage origination perfectly informs households about the future persistent cost of refinancing. Alternatively if this variance is very high, households learn little about the future process of refinancing from their experience during origination, since k_i is likely quite different from k_i^o.

We now solve the model to determine two household choices: (1) whether or not to refinance at each opportunity; and (2) the optimal size of the initial loan $l_{i,0}^*(v_i, k_i^o)$.

3.3.1 Optimal Refinancing

Households refinance when their refinancing costs are below a threshold that depends on their loan size. Hence, households with larger loans are more likely to refinance. Similarly, because the loan is amortizing, each household's incentives to refinance decline over time as the outstanding balance decreases; notably, some households (almost) always refinance because they have a low value of the persistent component k_i of the cost of refinancing.

We solve for the optimal refinancing path by backward induction. Consider period T, which is the last refinancing period, and households with a beginning-of-period (i.e., before making a payment) loan balance l_i (we suppress the subscript t for simplicity). Such households refinance if their refinancing cost $k_{i,T}$ is below the benefit of refinancing $k_i^*(T)$:

$$\begin{aligned} k_i^*(T) =& m(l_i, R, 1) - m(l_i, r_d, 1) \\ =& l_i(R - r_d). \end{aligned}$$

The benefit of refinancing depends on the difference between the interest rates $R - r_d$, as well as on the loan balance l_i.

We can define the expected (i.e., prior to the realization of the transitory component $\varepsilon_{i,t}$) value function $V_T(k_i, l_i)$ of a household with persistent cost k_i as the expected payment:

$$\begin{aligned}
V_T(k_i, l_i) &= \int_0^{+\infty} \min\left(m(l_i, r_d, 1) + k_i\varepsilon_{i,T}, m(l_i, R, 1)\right) dF(\varepsilon_{i,T}) \\
&= \int_0^{k_i^*(T)/k_i} \left(m(l_i, r_d, 1) + k_i\varepsilon_{i,T}\right) dF(\varepsilon_{i,T}) + \int_{k_i^*(T)/k_i}^{+\infty} m(l_i, R, 1) dF(\varepsilon_{i,T}), \quad (3.3)
\end{aligned}$$

where $k_i^*(T)/k_i$ is the cutoff point in the distribution of the transitory component $\varepsilon_{i,t}$ that determines household refinancing.[17]

Similarly, in the previous period $T-1$, households' expected value function equals the discounted sum of expected future payments:

$$\begin{aligned}
V_{T-1}(k_i, l_i) = &\int_0^{+\infty} \min\left(m(l_i, r_d, 2) + k_i\varepsilon_{i,T-1} + \beta V_T\left(k_i, l_i(1+r_d) - m(l_i, r_d, 2)\right), \dots\right. \\
&\left. m(l_i, R, 2) + \beta V_T\left(k_i, l_i(1+R) - m(l_i, R, 2)\right)\right) dF(\varepsilon_{i,T-1}) \\
= &\int_0^{k_i^*(T-1)/k_i} \left(m(l_i, r_d, 2) + k_i\varepsilon_{i,T-1} + \beta V_T\left(k_i, l_i(1+r_d) - m(l_i, r_d, 2)\right)\right) dF(\varepsilon_{i,T-1}) + \\
&\int_{k_i^*(T-1)/k_i}^{+\infty} \left(m(l_i, R, 2) + \beta V_T\left(k_i, l_i(1+R) - m(l_i, R, 2)\right)\right) dF(\varepsilon_{i,T-1}),
\end{aligned}$$

where

$$\begin{aligned}
k_i^*(T-1) = &m(l_i, R, 2) + \beta V_T\left(k_i, l_i(1+R) - m(l_i, R, 2)\right) + \\
&- m(l_i, r_d, 2) - \beta V_T\left(k_i, l_i(1+r_d) - m(l_i, r_d, 2)\right)
\end{aligned}$$

defines the monetary benefits of refinancing, such that households with $k_{i,t} \leq k_i^*(T-1)$ refinance, and households with $k_{i,t} > k_i^*(T-1)$ do not.

In a generic period t, the expected value function equals:

$$\begin{aligned}
V_t(k_i, l_i) = &\int_0^{+\infty} \min\left(m(l_i, r_d, T-t+1) + k_i\varepsilon_{i,t} + \beta V_{t+1}\left(k_i, l_i(1+r_d) - m(l_i, r_d, T-t+1)\right), \dots\right. \\
&\left. m(l_i, R, T-t+1) + \beta V_{t+1}\left(k_i, l_i(1+R) - m(l_i, R, T-t+1)\right)\right) dF(\varepsilon_{i,t}),
\end{aligned}$$

[17]When $\varepsilon_{i,T}$ equals $k_i^*(T)/k_i$, then $k_i \cdot \varepsilon_{i,T}$ equals $k_i^*(T)$ and the cost is exactly equal to the benefit of refinancing.

and the benefits $k_i^*(t)$ determine the cutoff point in the cost distribution that characterizes household refinancing decisions.

Therefore, we can describe the optimal refinancing policy as follows:

$$r(l_i, k_{i,t}) = \begin{cases} r_d & \text{if } k_{i,t} \leq k_i^*(t) \\ R & \text{otherwise.} \end{cases} \tag{3.4}$$

Hence, households with a lower persistent component k_i are more likely to refinance and pay the discounted rate r_d than households with a higher k_i. Moreover, the refinancing behavior of each household varies over time depending on the realization of the transitory shock $\varepsilon_{i,t}$. Generally, because the transitory shock $\varepsilon_{i,t}$ is multiplicative, its realization has a smaller effect on the refinancing activity of borrowers with low k_i, and a larger effect on that of borrowers with high k_i.

3.3.2 Optimal Loan Size

Households choose the loan size that maximizes their value function at origination, given their valuation for housing v_i and origination cost k_i^o. The value at origination equals:

$$W_0(v_i, k_i^o) = \max_{l_{i,0}} \sum_{t=0}^{+\infty} \beta^t v_i(\omega_i l_{i,0})^\alpha - k_i^o - \beta \int_0^{+\infty} V_1(k_i^o \varepsilon_{i,0}, l_{i,0}) dF_0\left(\varepsilon_{i,0}\right), \tag{3.5}$$

where the loan-to-value at origination equals $l_{i,0}/h_i = 1/\omega_i$ and k_i^o is the mortgage origination cost described above. Households do not know the exact value of their future refinancing cost and thus they form their expectations based on the available signal, which is their origination cost.

The optimal loan size $l_{i,0}^*(v_i, k_i^o)$ satisfies the first-order condition

$$\frac{\alpha \omega_i v_i (\omega_i l_{i,0}^*)^{\alpha-1}}{1-\beta} - \beta \frac{\partial}{\partial l_{i,0}} \int_0^{+\infty} V_1(k_i^o \varepsilon_{i,0}, l_{i,0}^*) dF_0\left(\varepsilon_{i,0}\right) = 0. \tag{3.6}$$

Hence, the optimal loan size depends directly on the household valuation for housing v_i, and indirectly on the origination costs k_i^o, because it is correlated with the expected future mortgage payments through the optimal refinancing policy $r(l_{i,t-1}, k_{i,t})$ described above. The refinancing policy in (3.4) highlights that refinancing costs determine the extent to which households make mortgage payments at the higher reset rate rather than at the lower discounted rate. This is because obtaining the cheaper discounted rate in a greater number of periods requires incurring the refinancing cost $k_{i,t}$ across a greater expected number of refinancing opportunities.

Given the optimal loan size, we can define $v_i^*(k_i^o)$ as the valuation for housing of a household that is indifferent between buying a house and getting a mortgage or renting a property:

$$W_0(v_i^*, k_i^o) = \frac{\bar{u}}{1-\beta}, \tag{3.7}$$

where $\bar{u}$ is a per-period utility of the outside rental option. This extensive-margin condition determines whether or not households enter the housing market rather than rent: households with a high valuation v_i and a low cost k_i^o enter the housing and mortgage market.

The precision of information that households have about their future refinancing costs plays into both optimal loan size (the intensive margin described in equation (3.6)) and whether or not households enter the housing market in the first place (the extensive margin described in equation (3.7)). On the intensive margin, a higher k_i generates an incentive to scale back the size of the initial loan, and on the extensive margin, a higher k_i may be a deterrent to entering the mortgage market in the first place. Conditional on the other parameters including their housing valuation v_i, the extent to which this effect operates depends on the variance of $\varepsilon_{i,0}$. If this variance is small, households choose an initial loan size that is strongly correlated with the origination costs k_i^o and thus with the persistent component k_i of refinancing costs. If the variance of $\varepsilon_{i,0}$ is larger, households have less precise information at origination to evaluate their future mortgage costs. Hence, their initial loan size will be weakly correlated with the cost k_i. The variance of the transitory component $\varepsilon_{i,t}$ of refinancing costs similarly affects households' optimal initial loan size, because a larger variance of $\varepsilon_{i,t}$ makes it more difficult for households to predict their refinancing activity, and thus the rates of their future mortgage payments.

Equation (3.7) shows that origination costs k_i^o also capture any household constraints to becoming homeowners. Once again, the precision of households' information at origination, captured by the variance of $\varepsilon_{i,0}$, critically affects this adjustment. This condition will play an important role in our counterfactual analysis as it determines how initial homeownership and mortgage takeup change. We return to these issues in greater detail when evaluating counterfactuals.

3.3.3 Aggregation: Mortgage Stocks in Steady-State

We calculate the total stock of mortgages that pay the discounted rate and the reset rate, assuming that the economy is in steady state.

It is useful in this calculation to recursively define the endogenous cumulative distribution function $H_t(\cdot)$ and its associated density $h_t(\cdot)$ of loan balances t periods after origination, given the evolution of the loan balances in (3.2), and the refinancing policy described in (3.4). This distribution evolves as follows:

$$H_0(z) = \iint_{\{(v_i,k_i^o):v_i \geq v_i^*(k_i^o) \cap l_{i,0}^*(v_i,k_i^o) \leq z\}} g_o(v_i,k_i^o)dv_i dk_i^o,$$

$$H_t(z) = \int_{\{l_{i,t-1}:l_{i,t}(r,l_{i,t-1}) \leq z\}} h_{t-1}(l_{i,t-1})dl_{i,t-1}.$$

We next define three groups (0, 1, 2) of mortgages. Group 0 comprises the mortgages of households who took a mortgage of initial size $l_{i,0}^*(v_i,k_i^o)$ and are on their initial discount period. The aggregate number $N_0(r_d)$ and aggregate balance $Q_0(r_d)$ of mortgages of this group equal:

$$N_0(r_d) = M \int_{-\infty}^{+\infty} \int_{v_i^*(k_i^o)}^{+\infty} g(v_i,k_i^o)dv_i dk_i^o, \tag{3.8}$$

$$Q_0(r_d) = N_0(r_d) \int_0^{+\infty} z h_0(z)dz = M \int_{-\infty}^{+\infty} \int_{v_i^*(k_i^o)}^{+\infty} l_{i,0}^*(v_i,k_i^o) g_o(v_i,k_i^o)dv_i dk_i^o. \tag{3.9}$$

To gain an intuition for equation (3.8), recall that a mass M of households enters the market in each time period. (Discounted) mortgage takeup among these households is determined by whether or not they satisfy the extensive margin condition $v_i \geq v_i^*(k_i^o)$, with the outer integral integrating across the k_i^o distribution. Equation (3.9) follows by weighting these mortgages by their initial loan sizes.

The second group comprises the mortgages of all households who refinanced and pay the discounted rate. In each period $t \in \{1, \ldots, T-1\}$, the number $N_{1,t}(r_d)$ of mortgages in this group equals:

$$N_{1,t}(r_d) = N_0(r_d) \int\limits_{\{l_{i,t}:r(l_{i,t},k_{i,t})=r_d\}} h_t(l_{i,t})dl_{i,t} \tag{3.10}$$

Equation (3.10) combines all borrowers who have refinancing costs $k_{i,t}$ lower than the benefits $k_i^*(t+1)$, and thus have policy functions $r(l_{i,t}, k_{i,t}) = r_d$. Thus, the aggregate number $N_1(r_d)$ of mortgages of this group equals:

$$N_1(r_d) = \sum_{t=1}^{T-1} N_{1,t}(r_d). \tag{3.11}$$

The aggregate balance of this group is the sum of the balances of the different cohorts who pay the discounted rate r_d. The aggregate balances $Q_{1,t}(r_d)$ of these cohorts evolve as follows:

$$Q_{1,t}(r_d) = N_0(r_d) \int\limits_{\{l_{i,t}:r(l_{i,t},k_{i,t})=r_d\}} l_{i,t}h_t(l_{i,t})dl_{i,t}.$$

Thus, the aggregate balance equals $Q_1(r_d) = \sum_{t=1}^{T-1} Q_{1,t}(r_d)$.

The third group comprises the mortgages of all households who did not refinance, and pay the reset rate. In each period $t \in \{1, \dots, T-1\}$, the number $N_{2,t}(R)$ of mortgages in this group equals:

$$N_{2,t}(R) = N_0(r_d) \int\limits_{\{l_{i,t}:r(l_{i,t},k_{i,t})=R\}} h_t(l_{i,t})dl_{i,t}, \tag{3.12}$$

which is the set of borrowers who have refinancing costs above the benefits $k_i^*(t+1)$, and thus have policy functions $r(l_{i,t}, k_{i,t}) = R$. Thus, the aggregate number of households who pay the reset rate equals

$$N_2(R) = \sum_{t=1}^{T-1} N_{2,t}(R). \tag{3.13}$$

The aggregate balance $Q_2(R)$ of this group is the sum of the balances of the different cohorts who pay the reset rate R: $Q_2(R) = \sum_{t=2}^{T} Q_{2,t}(R)$, where $Q_{2,t}(R)$ evolves as follows:

$$Q_{2,t}(R) = N_0(r_d) \int\limits_{\{l_{i,t}:r(l_{i,t},k_{i,t})=R\}} l_{i,t} h_t(l_{i,t}) dl_{i,t}.$$

The above expressions can be directly mapped to the empirically observed stock of mortgages in each category, under the assumption that the market is in steady state.

3.3.4 Cross-Subsidy

To calculate the cross-subsidy across different households, we consider a benchmark case in which all mortgages have a constant interest rate r_c for their entire duration. In Section 3.5, we consider several values of this constant interest rate.

Under the constant interest rate r_c, households do not need to refinance and their mortgage payments are constant over time. Hence, their optimal loan size $l^{**}_{i,0}(v_i, k^o_i)$ maximizes the value function at origination (3.5) evaluated at $\varepsilon_{i,t} = 0$ for all $t > 0$, with a constant payment stream $m(l_{i,0}, r_c, T)$. The expression for optimal loan size simplifies to:

$$\begin{aligned} l^{**}_{i,0}(v_i, k^o_i) =& \frac{1}{\omega_i} \left(\frac{1-\beta}{\alpha \omega_i v_i} \left(\sum_{t=1}^{T} \beta^t \frac{\partial m(l_{i,0}, r_c, T)}{\partial l_{i,0}} \right) \right)^{\frac{1}{\alpha-1}} \\ =& \frac{1}{\omega_i} \left(\frac{\beta(1-\beta^T)}{\alpha \omega_i v_i} \frac{r_c(1+r_c)^T}{(1+r_c)^T - 1} \right)^{\frac{1}{\alpha-1}}. \end{aligned} \tag{3.14}$$

The aggregate number $N(r_c)$ and aggregate balance $Q(r_c)$ of mortgages in this scenario then equals:

$$\begin{aligned} N(r_c) &= MT \int_{-\infty}^{+\infty} \int_{v_i^{**}(k_i^o)}^{+\infty} g_o(v_i, k^o_i) dv_i dk^o_i, \\ Q(r_c) &= M \sum_{t=1}^{T} \gamma_{r_c}(t-1) \int_{-\infty}^{+\infty} \int_{v_i^{**}(k_i^o)}^{+\infty} l^{**}_{i,0}(v_i, k^o_i) g_o(v_i, k^o_i) dv_i dk^o_i, \end{aligned}$$

where we define

$$\gamma_{r_c}(t-1) = \frac{l_{i,t}(r_c, l_{i,0})}{l_{i,0}} = \frac{(1+r_c)^T - (1+r_c)^t}{(1+r_c)^T - 1},$$

as the beginning-of-period-t share of the initial loan still to be repaid, and $v_i^{**}(k_i^o)$ is the valuation of a household that is indifferent between buying a house and getting a mortgage, or renting a property in this constant rate scenario. Thus, households still face the origination cost k_i^o that, as we recount above, includes additional household constraints to homeownership, but no subsequent refinancing costs, i.e., $k_{i,t} = 0$ for $t > 0$.

Based on this counterfactual constant rate r_c, the estimated parameters of the model, and the observed discounted rate r_d and reset rate R, we can calculate the differences in mortgage market outcomes between the current and counterfactual scenarios for each household (v_i, k_i^o). These outcomes include differences in loan sizes and mortgage payments between current and counterfactual scenarios. They also include a measure of the lifetime cross-subsidy paid or received by the household. This can be measured as the household-level reduction or increase (when comparing current and counterfactual scenarios) in the "all-in" interest rate including any refinancing costs. These household-level calculations can be aggregated up at the group level using the baseline model, or indeed, using an extended version of the model in which we estimate group-specific parameters. We describe this extended model next.

3.3.5 Multiple Groups

The richness of our data allows us to calculate subsidies across different groups based on observable demographic characteristics. We focus on two specific household groupings. The first groups households by income, and the second looks at households located in different UK regions.

Understanding variation in the extent of cross-subsidies paid or received along the income distribution helps us to understand how the design of the financial system contributes to the inequality of financial wealth, to the extent that wealth and income are correlated. We also look at the extent of regional variation in mortgage cross-subsidies given the importance of regional re-distribution through the mortgage market.

We extend the model to accommodate and interpret such heterogeneity. Consider different groups based on observable characteristics and indexed by $j = 1, ..., J$. Let M_j and $G_{oj}(v_i, k_i^o)$ be the measure and the cumulative distribution function of household housing preferences v_i and origination costs k_i^o in group j, respectively. Following the analysis of previous subsections, we can define the variables $N_{0,j}(r_d), Q_{0,j}(r_d), \ldots, Q_{2,j}(R)$ for each group j, and proceed with our counterfactual comparisons as before using this extended model.

We next turn to acquire quantitative estimates of the model's parameters and an assessment of the model-implied cross-subsidy by mapping the model to the data.

3.4 Quantitative Analysis

The model does not admit an analytic solution for all endogenous outcomes. As a result, we choose the parameters that best match moments of the data with the corresponding moments computed from the numerical solution of the model in steady state. We then study the quantitative implications of the model evaluated at the estimated parameters.

3.4.1 Estimation

We fix a subset of parameters, often reading them directly from the data, and we estimate the remaining parameters of the model to best match key moments of the mortgage data.

Specifically, we set the unit of time in the model to be $T_d = 2$ years, which is the modal initial fixation period in the UK mortgage market over the sample period; we then set the mortgage maturity at $T = 15$ periods, to give us the modal mortgage origination maturity of 30 years. We set the discount rate at $\beta = 0.95^2 = 0.9025$ to correspond to our assumption on the unit of time.

We read the annual interest rates on discounted and reset rate mortgages directly from the underlying data, using value-weighted averages of the corresponding rates in the 2015H1 sample, and compound them to correspond to the two-year T_d. Annual average discounted and reset rates equal 330 bps and 375 bps in our sample, meaning $r_d = 650$ bps and $R = 759$ bps over two-years.

We set the loan-to-value ratio at origination common across households at 80 percent, close to the modal value in our data, so $\omega = 1.25$.

We read market size M from the data, as follows. The total number of mortgages in the model equals:

$$N_0(r_d) + N_1(r_d) + N_2(R) = MT \int_{-\infty}^{+\infty} \int_{v_i^*(k_i^o)}^{+\infty} g_o(v_i, k_i^o) dv_i dk_i^o. \tag{3.15}$$

Hence, we compute the market size M by dividing the total number of mortgages $N_0(r_d)+N_1(r_d)+N_2(R)$ by their maturity T and by the share of households who own a property $\int_{-\infty}^{+\infty} \int_{v_i^*(k_i^o)}^{+\infty} g_o(v_i, k_i^o) dv_i dk_i^o$.

We estimate all remaining parameters by applying some assumptions about distributions. We assume that households' valuation v_i follows a lognormal distribution, i.e., $\log(v_i)$ follows a normal distribution with mean μ_v and standard deviation σ_v. We assume that the origination cost k_i^o follows a mixture distribution of two lognormal distributions, which allows for a bimodal distribution of origination costs. We model the persistent component of refinancing costs as $k_i = k_i^o \varepsilon_{i,0}$, which is correlated with the origination cost. As a result, the distribution of refinancing costs could be bimodal as well, with some households with low refinancing costs and others with high refinancing costs, and heterogeneity within these two household groups. With probability η, $\log(k_i^o)$ follows a normal distribution with mean μ_{k1} and standard deviation σ_{k1}; with probability $1-\eta$, $\log(k_i^o)$ follows a normal distribution with mean μ_{k2} and standard deviation σ_{k2}. Without loss of generality, we denote type-1 as households whose average cost is lower than the average cost of type-2 households—i.e., $\exp\left(\mu_{k1} + \frac{\sigma_{k1}^2}{2}\right) \leq \exp\left(\mu_{k2} + \frac{\sigma_{k2}^2}{2}\right)$. We do not restrict the variances of these distributions, and thus the type-2 distribution does not necessarily first-order stochastically dominate the type-1 distribution. Clearly, the assumption of lognormality implies that some type-1 households have higher costs than some type-2 households. We set $\eta = 0.5$, and the correlation between v_i and k_i^o to zero, because the empirical moments that we employ in the estimation do not allow us to separately identify these parameters, as we explain in more detail below.

We further assume that $\varepsilon_{i,t}$ follows a lognormal distribution with parameters μ_ε and σ_ε. We set the mean of $\varepsilon_{i,t}$ equal to one, meaning that $E(k_{i,t}) = k_i$, and hence $\mu_\varepsilon = -\sigma_\varepsilon^2/2$.

We estimate two versions of the model under two different assumptions about the distribution of $\varepsilon_{i,0}$. Version 1 is simply a degenerate distribution, with $\varepsilon_{i,0} = 1$ with probability one. Here, households perfectly know the persistent component of their refinancing costs k_i at origination. Put differently, they know their average refinancing costs because $E(k_{i,t}) = k_i$. Of course, households still face ex-ante uncertainty about their future refinancing costs because they do not know the temporary component of refinancing costs which is governed by $\varepsilon_{i,t}$. Version 2 of the model assumes that $\varepsilon_{i,0}$ is governed by the same distribution as $\varepsilon_{i,t}$.[18] In this version of the model, households obtain a noisy signal of the persistent component of their refinancing costs (and thus of their average refinancing costs) at the point of mortgage origination. Therefore, their loan size and participation decisions will exhibit weaker correlation with their refinancing costs than in version 1. The two versions differ in the precision of borrowers' information at origination, but require the same number of parameters to be estimated.

Finally, our estimation recovers the parameter α of the utility function and the level of the outside option $\bar{u}$.

We search for the vector of 9 parameters $\psi = (\mu_v, \sigma_v, \mu_{k1}, \sigma_{k1}, \mu_{k2}, \sigma_{k2}, \sigma_\varepsilon, \alpha, \bar{u})$ that minimizes the distance between selected moments in the data and the corresponding moments of the model. More specifically, for each combination of these unknown parameters, we solve the model shown in Section 3.3 to find households' optimal policies, characterized by their choice between buying a house with a mortgage or renting a property, and, if they choose to participate in the mortgage market, their mortgage loans at origination $l^*_{i,0}(v_i, k^o_i)$ and their optimal sequence of refinancing. Based on these household policies, we compute the following aggregate moments:

1. the average loan balance for mortgages on the discounted rate;
2. the standard deviation of the loan balance of mortgages on the discounted rate;
3. the average loan balance for mortgages on the reset rate;
4. the standard deviation of the loan balance of mortgages on the reset rate;
5. the average remaining maturity of mortgages on the discounted rate;
6. the standard deviation of the remaining maturity of mortgages on the discounted rate;

[18] In other words, in this version of the model, if shocks to household attention/distraction (i.e., temporary refinancing costs shocks) are drawn from a high-variance distribution, this also means that households learn less from their mortgage origination costs about their persistent refinancing costs and vice versa.

7. the average remaining maturity of mortgages on the reset rate;

8. the standard deviation of the remaining maturity of mortgages on the reset rate;

9-14. the shares of mortgages on the discounted rate for the following partition of the loan balance distribution: $[0-5]$ percentile, $(5-25]$ percentile, $(25-50]$ percentile, $(50-75]$ percentile, $(75-95]$ percentile, and $(95-100]$ percentile;

15. the share of mortgages on the reset rate in 2015H1 that paid the discounted rate in 2017H1;

16. the share of homeowners, i.e., the fraction of households that enter the housing market and choose to purchase a house and take on a mortgage.

Section 3.2 outlines several filters that we apply to the data. One of these filters is that the outstanding mortgage balance exceeds £30,000, and for consistency, we apply the same filter when computing moments 1 to 15 in the model.

The minimum-distance estimator chooses the parameters that minimize the criterion function:

$$(\mathbf{m}(\psi) - \mathbf{m}_S)' \, \Omega \, (\mathbf{m}(\psi) - \mathbf{m}_S),$$

where $\mathbf{m}(\psi)$ is the vector of moments computed from the model at the parameter vector ψ and $\mathbf{m}_S$ is the vector of corresponding sample moments. Ω is a symmetric, positive-definite matrix; in practice, in order for the moments to have a similar scale, we use a diagonal matrix whose elements are those on the main diagonal of the inverse of the matrix $E(\mathbf{m}_S'\mathbf{m}_S)$.

We estimate six cases of the model. Two baseline versions follow from the two different assumptions about the distribution of $\varepsilon_{i,0}$ recounted above. In both cases, we pool together all mortgages in our data and assume that all households can be characterized by a single distribution $G_o(v_i, k_i^o)$, as well as common σ_ε, α, and $\bar{u}$ parameters. This entails estimating 9 parameters using the 16 moments listed above.

We also pursue the estimation in cases with richer borrower heterogeneity. The first one estimates the model separately for different income groups, and the second one estimates the model separately for different geographic areas of the UK. We estimate all these cases with richer heterogeneity in the two versions with the two different assumptions about the distribution of $\varepsilon_{i,0}$. In each case, we set group-specific market sizes M_j and to estimate group-specific parameters of the distributions $G_{oj}(v_i, k_i^o)$, as well as σ_{ε_j}, α_j, and $\bar{u}_j$ for each group j (denoting either income groups or geographical areas).[19] This gives us additional flexibility to capture heterogeneity across groups in preferences, costs, and ultimately refinancing activities. Of course, when we estimate these parameters, we do so using an expanded set of group-specific moments in each case.

We consider 12 income groups based on the following percentiles of the distribution of reported incomes in the PSD: 0-10, 10-20, 20-30, 30-40, 40-50, 50-60, 70-80, 80-85, 85-90, 90-95, and 95-100. We also consider 12 broad regions and devolved administrations of the UK, namely North-East, North-West, Yorkshire and the Humber, East Midlands, West Midlands, East of England, Greater London, South East, South West, Wales, Scotland, and Northern Ireland.[20]

Hence, we estimate a total of 108 parameters (9 parameters for each of the 12 groups) using a total of 192 moments (16 moments listed above for each of the 12 groups).

3.4.2 Sources of Identification

The model is highly nonlinear, so (almost) all parameters affect all outcomes. That said, the identification of certain parameters does rely more heavily on particular moments in the data.

More specifically, moments characterizing the distributions of loan sizes on the discounted and the reset rate, those characterizing the distributions of remaining maturities in each mortgage category, and the shares of mortgages in the two categories together identify the parameter α, and the parameters of the distributions of household preferences v_i and the persistent component of costs k_i. Notably, households' initial loan amounts—and, thus over time, their loan balances—depend on their housing preferences v_i, as well as their expected refinancing costs which on average equal k_i. Moreover, for every mortgage, the parameter α affects the sensitivity of the initial loan size to expected mortgage payments, and thus to interest rates, as equations (3.6) and

[19] Hence, in the case of version 2, we implicitly assume that the heterogeneity of refinancing costs across groups is (perfectly) positively correlated with the noise of their information at origination, i.e., groups with smaller temporary shocks to refinancing costs have more precise information about their future refinancing costs at origination than groups with larger temporary shocks to refinancing costs.

[20] These are the 12 NUTS-1 regions of the UK, where NUTS stands for Nomenclature of Territorial Units for Statistics.

(3.14) show.

If the cost $k_{i,t}$ was prohibitively high for all borrowers, almost all mortgages would be on the reset rate, and conversely, if $k_{i,t}$ was extremely low for all borrowers, all mortgages would be on the discounted rate. Hence, the shares of mortgages on the reset rate are informative about the parameters of the distribution of the refinancing cost $k_{i,t}$ and its components.

Given a value of $k_{i,t}$, borrowers have stronger financial incentives to refinance if they have a large loan balance, meaning that the share of mortgages on the discounted rate should be increasing in the loan balance. The rate of change of the share of mortgages on each rate as loan size changes is informative about the heterogeneity in $k_{i,t}$. The increase is fast if the heterogeneity across households is small, whereas it is slow if the heterogeneity is large. Our assumption that k_i follows a mixture distribution allows us to flexibly capture different rates of increase in the share of mortgages on the discounted rate at different percentiles of the loan balance distribution. This means that the change in the share in the two categories of mortgages at different levels of the loan balance contributes to the identification of the refinancing cost heterogeneity parameters σ_{k1} and σ_{k2} of the mixing distribution. Because we allow the support of the two k_i distributions to overlap, it is difficult to separately identify the mixing probability η; therefore, we set it to $\eta = 0.5$. Moreover, identifying any correlation between v_i and k_i would require rich within-borrower moments; our mapping to the stock means that all our moments are cross-sectional (except for the share of mortgages on the reset rate that later pay the discounted rate, moment 15 above), so we set this correlation to zero by assumption.

The share of mortgages that transition from paying the reset rate to paying the discounted rate is informative about the within-borrower heterogeneity in refinancing costs, and thus identifies the parameter σ_ε. If refinancing costs were fixed over time for each borrower, because loan balances decline over time, borrowers' optimal refinancing policy would be deterministic: it would be characterized by a borrower-specific cutoff date $T_{max}(v_i, k_i)$, such that a (v_i, k_i)-borrower always refinances before $T_{max}(v_i, k_i)$ and never does after $T_{max}(v_i, k_i)$. Transitions from the reset rate to the discounted rate violate this deterministic refinancing policy, and therefore identify the within-borrower, transitory variation in refinancing costs governed by σ_ε.

Moreover, our data does not allow us to identify households' information and beliefs at origination about their future refinancing costs, captured by the variance of $\varepsilon_{i,0}$. Hence, we set it to different values in the two versions, and the similarity of the results of these two versions will allow us to establish the robustness of our results to differences in households' information at origination.[21]

Finally, the share of owners versus renters identifies the level of outside option utility $\bar{u}$.

3.4.3 Parameters and Model Fit

Table 3.4.1 reports the parameters of the model for the six cases of the estimated model: aggregate, income group-specific, and geography-specific, each one with the two versions with different assumptions about the precision of households' information at origination. The top of the table reports the fixed parameters, which are common across cases and across groups.

The main body of the table reports the estimated parameters. Columns (1) and (2) report the parameter estimates for the baseline versions that use UK-wide moments, and their asymptotic standard errors in parentheses.[22] The model in column (1) assumes $\varepsilon_{i,0} = 1$ with probability one, and thus at origination households know their persistent component of refinancing costs; the model in column (2) assumes that at origination households receive a noisy signal of their future persistent component of refinancing costs only.

Columns (3) and (4) report the estimates for the case that uses separate moments for each income group, and columns (5) and (6) for the case that uses separate moments for each region and devolved administration. Columns (3) and (5) assume that households have precise information about their persistent component of refinancing costs; columns (4) and (6) assume that households receive noise signals only. In columns (3)-(6), we report the weighted averages of the parameters across groups, as well as the

[21]Of course, we could set the value of $\varepsilon_{i,0}$ to alternative, higher values than those that we choose.

[22]To obtain standard errors, we compute the covariance matrix of the moments $\mathbf{m}_S$ by bootstrapping. Specifically, for N_s bootstrap resamples of the data, the covariance matrix of the moments $\mathbf{m}_S$ equals

$$W_{N_s} = N_s^{-1} \sum_{n_s=1}^{N_s} (\mathbf{m}_{n_s} - \mathbf{m}_S)(\mathbf{m}_{n_s} - \mathbf{m}_S)', \tag{3.16}$$

where $\mathbf{m}_{n_s}$ is the vector of moments in resample n_s. We set the number of resamples N_s at 1,000.

Table 3.4.1: Parameters

r	650	R	759	T	15	β	0.902	ω	1.250	η	0.500

	UK-WIDE		INCOME GROUPS		REGIONS	
	(1)	(2)	(3)	(4)	(5)	(6)
μ_v	0.001	0.001	0.001	0.001	0.001	0.001
	(0.004)	(0.004)	(0.000)	(0.000)	(0.000)	(0.000)
σ_v	0.150	0.157	0.234	0.230	0.349	0.349
	(0.010)	(0.013)	(0.010)	(0.012)	(0.023)	(0.022)
μ_{k1}	4.883	5.763	4.755	4.960	4.925	5.004
	(1.066)	(0.706)	(0.595)	(0.709)	(0.424)	(0.530)
σ_{k1}	2.670	2.618	2.823	2.978	2.333	2.347
	(1.263)	(1.287)	(0.394)	(0.454)	(0.987)	(0.985)
μ_{k2}	9.164	9.196	9.088	9.250	9.188	9.197
	(0.634)	(1.583)	(0.184)	(0.093)	(0.041)	(0.025)
σ_{k2}	0.988	0.987	0.960	0.947	0.987	0.975
	(0.072)	(0.691)	(0.059)	(0.069)	(0.015)	(0.035)
σ_ϵ	1.048	0.926	0.992	0.944	1.041	0.891
	(0.180)	(0.297)	(0.033)	(0.036)	(0.056)	(0.160)
α	0.787	0.786	0.789	0.789	0.788	0.788
	(0.001)	(0.002)	(0.007)	(0.007)	(0.003)	(0.003)
$\bar{u}$	1,190	1,039	1,580	1,224	1,455	1,323
	(235)	(544)	(606)	(441)	(436)	(484)
M	379,145	379,145	27,850	27,850	31,894	31,894
			(10,407)	(10,407)	(15,298)	(15,298)

Notes: This table reports the estimated parameters. In columns (1) and (2), the numbers in parentheses refer to asymptotic standard errors of the parameter estimates. In columns (3)-(6), the numbers in parentheses refer to standard deviations of the parameter estimates across groups. Odd-numbered columns correspond to version 1 of the model, and even-numbered columns correspond to version 2 of the model.

weighted standard deviations of the parameters across groups (in parentheses), where the weights are the estimated market sizes M_j.

The bottom of Table 3.4.1 reports the calibrated market size M computed using equation (3.15); in columns (3)-(6), they correspond to the unweighted averages and standard deviations of M_j across groups. Note that several parameters are not easily comparable across columns. For example, the outside options $\bar{u}$ differ across groups in columns (3)-(6), and affect the estimated parameters of the valuation distribution. Other parameters, such as α, are more easily comparable across columns.

Baseline Models. We focus our discussions on the parameters in column (1) because the differences between those in columns (1) and (2) are small and the two versions of the model have quite similar implications for market outcomes. Nevertheless, we note the key differences between the parameters of the two versions, most notably in origination and refinancing costs, because the two versions differ mainly in these costs.

The estimated parameters in column (1) imply that households' valuation v_i has a median equal to 1.001, a mean equal to 1.012 and a standard deviation equal to 0.152 in the full population of borrowers (homeowners) and non-borrowers (renters). In the model, households with the lowest valuations are less likely to participate in the mortgage market, choosing instead to rent a property. This means that, among borrowers, valuations are higher, with median v_i equalling 1.066, mean 1.077, and standard deviation 0.128.

The estimate of the parameter $\alpha = 0.787$ implies modest concavity in household utility from housing. This value implies that a household with average v_i enjoys a utility flow of $v_i h^{\alpha}$, i.e., £$10,342$ over a two-year period from a house worth £$125,000$, for example. This translates into an annual yield of 4.054%, which is slightly lower than the average rental yield for the whole of the UK, but broadly in line with average rental yields reported for London in this period.[23]

In version 1 of the model, in which households know their persistent component of refinancing costs at origination, this persistent component k_i equals the origination costs k_i^o. Among homeowners/borrowers, the median origination cost/persistent component equals £634, its mean equals £$3,866$, and its standard deviation equals £$6,641$. However, households with the highest origination and refinancing costs are less likely to participate in the mortgage market and choose to rent a property. Because our moments do not report any information on households who do not borrow (except for their share in the population), we obtain the distribution of costs k_i (as well as that of preferences v_i) in the full population by extrapolating those of borrowers out of sample. This leads us to estimate the persistent component k_i of refinancing costs across all households, including those that do not borrow, with a median that equals £$3,003$, a mean of £$10,107$, and a standard deviation of £$117,213$ in the full population. It is worth noting that in the counterfactual exercises, we retain origination costs k_i^o, but remove refinancing costs $k_{i,t}$. This results in a relatively small effect on our calculations of the large $k_{i,t}$ values estimated for non-participants in the baseline dual-rate economy.

[23]See, for example, Savill's UK Report on Rents and Returns, 2015.

In version 2 of the model, in which households have noisy information about their persistent component of refinancing costs at origination, the persistent component k_i differs from the origination cost k_i^o. This version (whose parameters are displayed in column (2) of Table 3.4.1) displays slightly higher origination costs and persistent components of refinancing costs relative to version 1 in which households have precise information and these costs are equal. Among homeowners/borrowers, the median origination cost k_i^o equals £$1,465$, its mean equals £$4,460$, and its standard deviation equals £$7,354$; the median persistent component of refinancing cost k_i equals £836, its mean equals £$4,159$, and its standard deviation equals £$9,882$.

Interestingly, while we assume that preferences v_i and costs k_i are uncorrelated in the population of households, they are correlated among borrowers because of households' endogenous selection into the mortgage market—borrowers with high k_i enter the market only if their v_i is sufficiently high. The correlation coefficient among borrowers equals 0.340, suggesting that the effect of selection is appreciable.

The estimate of $\sigma_\varepsilon = 1.048$ in column (1) of Table 3.4.1 means that the standard deviation of $\varepsilon_{i,t}$ equals 1.413, which implies that the within-household variation in refinancing costs is non-trivial. This estimate of σ_ε means that the ratio $\frac{St.Dev.(k_i)}{St.Dev.(k_{i,t})}$ equals 0.58 in the population and 0.54 among borrowers—that is, the persistent household component k_i (cross-household variation) accounts for a slightly larger share of the standard deviation of the refinancing costs $k_{i,t}$ than the transitory component $\varepsilon_{i,t}$ (within-household variation). The estimate of $\sigma_\varepsilon = 0.926$ in column (2) of Table 3.4.1 implies that in version 2 the ratio $\frac{St.Dev.(k_i)}{St.Dev.(k_{i,t})}$ equals 0.65 in the population and 0.63 among borrowers. Hence, the transitory component accounts for a smaller share of the standard deviation of borrowers' total refinancing costs $k_{i,t}$ in version 2 than in version 1. Below, we provide more statistics on borrower refinancing costs $k_{i,t}$ and compare them to the benefits of refinancing in the two versions.

The value of the per-period outside option utility $\bar{u}$ equals £$1,190$, which implies an annual net utility from renting equal to $\frac{\bar{u}}{1+\beta^{1/2}}$=£$610$. Households with a net utility value (over and above all mortgage payments and refinancing costs) greater than this level from purchasing a house enter the mortgage market.

Table 3.4.2 presents a comparison between the empirical moments and the moments calculated from the model at the estimated parameters reported in columns (1)

Table 3.4.2: Model Fit

	Data	Version 1	Version 2
Mean Loan Balance, Discounted Rate	140,647	143,697	140,986
Standard Deviation Loan Balance, Discounted Rate	105,062	106,551	107,225
Mean Loan Balance, Reset Rate	112,692	113,741	110,869
Standard Deviation Loan Balance, Reset Rate	79,684	76,546	77,921
Mean Remaining Years, Discounted Rate	20.57	18.63	18.87
Standard Deviation Remaining Years, Discounted Rate	7.73	7.91	7.84
Mean Remaining Years, Reset Rate	16.84	15.56	15.54
Standard Deviation Remaining Years, Reset Rate	6.95	7.40	7.38
Share of Mortgages on Discounted Rate, 0-5 Percentile	52.72	52.82	51.84
Share of Mortgages on Discounted Rate, 5-25 Percentile	56.36	58.03	56.39
Share of Mortgages on Discounted Rate, 25-50 Percentile	61.48	60.12	59.20
Share of Mortgages on Discounted Rate, 50-75 Percentile	67.76	63.73	63.08
Share of Mortgages on Discounted Rate, 75-95 Percentile	73.77	72.10	71.32
Share of Mortgages on Discounted Rate, 95-100 Percentile	81.19	83.66	81.60
Transition from Reset Rate to Discounted Rate	16.52	16.42	17.09
Share of Owners	63.13	64.50	63.25
Criterion Function		0.0287	0.0244

Notes: This table reports the values of the empirical moments and of the moments calculated at the estimated parameters reported in columns (1) and (2) of Table 3.4.1.

and (2) of Table 3.4.1, respectively. Overall, both versions of the model fit the data very well, with version 2 performing slightly better than version 1. Critically, both versions match two features of the data that underscore refinancing incentives across households and over time: on average, mortgages on the discounted rate have higher balances and are closer to issuance (have greater remaining maturity) than those on the reset rate.

Similarly, Figure 3.4.1 displays the comparison between the model-implied shares of mortgages paying the discounted rate and their empirical analogs. Notably, both versions are well-able to capture the concave relationship between the two variables, with a faster rate of increase in the share of mortgages on the discounted rate at low balances and a slower rate of increase at high balances.

Multiple Groups. Columns (3) and (4) in Table 3.4.1 report parameters for the model estimated on the 12 different income groups, and Columns (5) and (6) for the model estimated on the 12 UK regions and devolved administrations. As discussed above, some parameters are not easily comparable between Columns (3)-(6) and Columns (1)-(2), though many are similar in magnitude to those reported in Columns (1) and (2) for the

Figure 3.4.1: Share of Loans on Discounted Rate

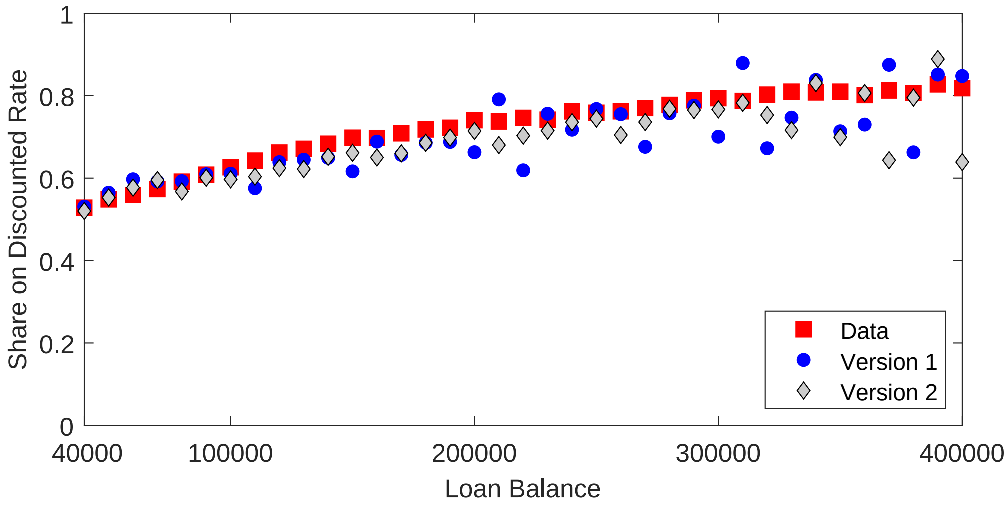

Notes: This figure displays the share of loans paying the discounted rate as a function of its loan balance in the data (red squares) and in the model evaluated at the estimated parameters reported in column (1) (blue dots) and column (2) (gray diamonds) of Table 3.4.1.

UK-wide case. The parameters in Columns (3) and (4) for the different income groups exhibit some differences from those in Columns (5) and (6) because the heterogeneity across and within income groups differs from the heterogeneity across and within regional groups, which in turn affects the average and the standard deviations of some of the parameters.

The parameters that exhibit the most meaningful heterogeneity in the population are those characterizing the distribution of the origination costs k_i^o and thus of the persistent component of refinancing costs k_i. This fact is particularly interesting for our purposes because this heterogeneity affects heterogeneity in refinancing activity across groups, and thus contributes directly to our quantitative assessment of the cross-subsidy across groups.

Moreover, the outside option $\bar{u}$ also displays significant heterogeneity in the population. This parameter is a key input into the "extensive margin" decision of households, i.e., whether or not they enter the mortgage market. The heterogeneity in this parameter across groups means that there are different sensitivities across groups of this extensive margin decision to changes in interest rates. This factor contributes to differ-

Figure 3.4.2: Distribution of Net Benefits of Refinancing

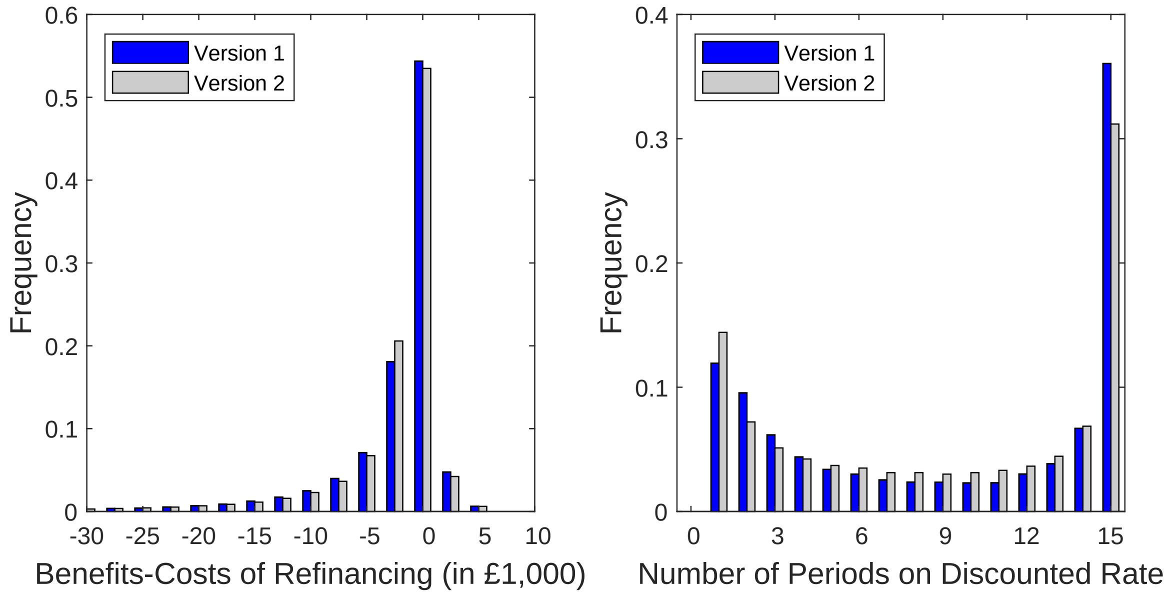

Notes: The left panel displays the histogram of the net benefits of refinancing. The right panel displays the histograms of the number of periods in which borrowers pay the discounted rate.

ences between the sensitivity of household participation decisions to interest rates in the multiple-group model and that in the baseline model.

While we do not report measures of goodness-of-fit across groups, we note that the model fits the group-specific moments well. This is perhaps not surprising given that Table 3.4.2 shows that the UK-wide model fits the aggregate data well; the same model might therefore be expected to fit as well or better at a lower level of aggregation.

3.4.4 Refinancing: Benefits and Costs

In this subsection, we discuss borrowers' refinancing behavior in the estimated UK-wide baseline model. We focus on version 1 of the model in which households know their persistent component of refinancing costs, because the qualitative patterns of refinancing behavior are quite similar between the two versions, noting the key differences between them.

The left panel of Figure 3.4.2 displays the full distribution of the net benefits of refinancing $k_i^*(t) - k_{i,t}$. The heterogeneity of net benefits is striking: in 60 percent of refinancing opportunities, borrowers have positive net benefits, thereby matching the aggregate share of mortgages on the discounted rate. The median estimated net benefit is positive (it equals $\pounds 304$), but the average net benefit is negative (it equals $-\pounds 2,619$), driven by the long left tail (the standard deviation equals $\pounds 12,240$). Some borrowers have extremely low measured net benefits, reflecting the fact that the model requires high costs to rationalize the non-refinancing behavior of a small group of borrowers with high loan balances and long maturities that would otherwise be expected to refinance. Moreover, we note that version 2 of the model exhibits slightly lower net benefits of refinancing than version 1: The median estimated net benefit equals $\pounds 237$ and the average net benefit equals $-\pounds 2,956$ (the standard deviation equals $\pounds 15,651$).

The distribution of gross benefits of borrowers who refinance has a median of $\pounds 1,039$, an average of $\pounds 1,357$, and a standard deviation of $\pounds 1,096$; their costs have a median of $\pounds 50$, an average of $\pounds 239$, and a standard deviation of $\pounds 466$. The corresponding distribution of gross benefits of borrowers who do not refinance has a median of $\pounds 885$, an average of $\pounds 1,084$, and a standard deviation of $\pounds 841$; the costs of these non-refinancing borrowers have a median of $\pounds 4,208$, an average of $\pounds 9,259$, and a standard deviation of $\pounds 18,030$. The comparison of these statistics between borrowers who refinance and borrowers who do not shows that the difference in their respective costs is larger than that in their benefits. Hence, heterogeneity in refinancing costs $k_{i,t}$ is the main driver in the model of the heterogeneity in refinancing behavior observed across borrowers.

The heterogeneity of refinancing behavior is also apparent in the right panel of Figure 3.4.2, which displays the distribution of the number of periods on the discounted rate across individuals. No borrower always pays the reset rate because all of them receive the discounted rate at origination, in period $t = 0$. Approximately 10 percent of borrowers never refinance thereafter, many borrowers refinance occasionally, and 36 percent of borrowers always refinance. This heterogeneous distribution obtains because borrowers with low values of their persistent component k_i of refinancing costs (almost) always refinance, whereas borrowers with high values of k_i refinance only when they receive a temporary shock $\varepsilon_{i,t}$ that is low enough. The distribution of the share of periods on the discounted rate across individuals in model 1 first-order stochastically dominates that in version 2, because, as we recount above, the estimated net benefits of refinancing are slightly lower in version 2 than in version 1.

While our primary focus is on refinancing behavior, we should also point out that the heterogeneity of borrower refinancing propagates into substantial heterogeneity in the elasticities of their initial loan size with respect to discounted and reset rates. The mean borrower elasticity with respect to the discounted rate r_d equals -1.566 and its standard deviation equals 0.722. Borrowers with a lower k_i are more elastic to the discounted rate (and less elastic to the reset rate) than borrowers with a higher k_i because they are more likely to refinance regularly and thus pay the discounted rate—the elasticity to the discounted rate of lowest-k_i borrowers equals -2.603. The mean borrower elasticity with respect to the reset rate R among borrowers equals -0.466 and its standard deviation equals 0.813. Interestingly, some borrowers display a positive elasticity with respect to the reset rate, because if the reset rate increases (while keeping the discounted rate fixed), the benefits of refinancing increase, and thus some borrowers are more likely to refinance. This leads to a lower average expected interest rate, meaning that such borrowers will increase their initial loan size in response to a higher reset rate.

3.5 Counterfactual Analyses: Constant Interest Rate

We compare the outcomes for households in our estimated models under the dual rate structure, with a counterfactual in which all households simply pay a constant interest rate and have no need to refinance. We perform these comparisons for four different values of the constant interest rate, namely:

1. The average discounted rate, i.e., $r_c = 650$ bps.

2. The weighted average of the discounted and the reset rates, i.e.,

$$r_c = \frac{r_d(Q_0(r_d) + Q_1(r_d)) + RQ_2(R)}{Q_0(r_d) + Q_1(r_d) + Q_2(R)}. \tag{3.17}$$

   We calculate this weighted average using the aggregate balances in the data and obtain $r_c = 683$ bps.

3. The rate that yields the same revenue as the composite of the populations on the discounted rate and the reset rate.

   More precisely, in the baseline case aggregate lender revenues from all mortgages (on both discounted and reset rates) equal:

$$r_d(Q_0(r_d) + Q_1(r_d)) + RQ_2(R). \tag{3.18}$$

Under the assumption of aggregate lender revenues remaining constant across the two scenarios, the interest rate r_c must satisfy:

$$r_c Q(r_c) = r_d(Q_0(r_d) + Q_1(r_d)) + RQ_2(R). \tag{3.19}$$

In practice, this equality yields $r_c = 700$ bps.[24]

4. The average reset rate, i.e., $r_c = 759$ bps.

The values of the constant interest rate in cases 1 and 4 likely represent lower and upper bounds to interest rates in a counterfactual market with constant rates, respectively, whereas cases 2 and 3 use intermediate values. Because these intermediate values seem more plausible to us than the other values, we focus on these cases more extensively below.

We note here that our model focuses on cross-household differences in borrowers' inaction—i.e., the demand side of the mortgage market. Our counterfactual scenarios attempt to capture a range of differences in the magnitudes of borrower cross-subsidies. Clearly, changes in the profile of interest rates affect lender profits and revenues as well, and their supply-side responses could constitute an important ingredient for further analysis.[25]

Table 3.5.1 reports the results of the counterfactual mortgage market outcomes for the different combinations of interest rates (in different panels) and estimated models (in different columns) as ratios of their respective baseline values (i.e., in the dual-rate economy). Odd-numbered columns correspond to version 1 of the model, and even-numbered columns correspond to version 2 of the model.

Perhaps not surprisingly, Panel A shows that all reported statistics increase in a market with a constant interest rate equal to the average discounted rate. Interest rates decline for all borrowers, except for those who always refinance, thereby boosting mort-

[24]When working with multiple groups, we perform the cross-subsidy calculation using the interest rate that satisfies:

$$r_c \sum_{j=1}^{J} Q_j(r_c) = \sum_{j=1}^{J} \left(r(Q_{0,j}(r_d) + Q_{1,j}(r_d)) + RQ_{2,j}(R) \right), \tag{3.20}$$

where $Q_j(r_c)$ is the aggregate mortgage debt of group j when the interest rate is fixed at r_c. The difference between equations (3.19) and (3.20) is that aggregate revenues are calculated using the heterogeneous parameters across groups. In practice, the difference between the interest rates that satisfy equations (3.19) and (3.20) is only a few bps, with minimal effects on the counterfactuals reported in Table 3.5.1.

[25]Among others, Gurun et al. (2016), Guiso et al. (2022), Benetton et al. (2021), Allen and Li (2020), and Thiel (2021) study supply-side incentives in mortgage markets.

gage debt both at the extensive margin (increasing the total number of mortgages) and at the intensive margin (increases in average initial loan sizes and average loan balances). These outcomes are remarkably similar across the different estimated models, i.e., single versus multiple groups, with the slight difference that the multiple-group models display a higher sensitivity of household participation decisions to interest rates relative to the UK-wide case, and thus a larger increase in the number of mortgages as interest rates decline.

Similarly, Panel D shows that all reported statistics decrease in a market with a constant interest rate equal to the average reset rate, because interest rates increase for all borrowers. We note that the model estimated separately for each income group reported in columns (3)-(4) features a similar aggregate number of mortgages to the UK-wide case of columns (1)-(2), which in combination with the results in Panel A suggests that allowing for heterogeneity between groups pushes toward slightly greater sensitivity to downward rate movements than upward rate movements.

Panels B and C report several interesting outcomes which appear robust across different estimated models. Panels B and C are overall quite similar, we therefore focus our discussion on Panel B. We first describe the changes for the UK-wide case, and then for the multiple-group cases.

Panel B, UK-wide case. Panel B reports that the change in the profile of interest rates to a single-rate structure yields two main aggregate adjustments in opposite directions: the number of mortgages increases, but the average loan size decreases. Again, our discussion focuses on version 1 of the model in column (1), noting the key differences with version 2 in column (2).

More precisely, the first row of Panel B in Table 3.5.1 reports that the number of mortgages increases by 6.49 percent relative to the number of mortgages in the baseline economy. The reason for this increase is that there are many households with valuation v_i and with moderate or high costs $k_i^o = k_i$ just below the entry threshold $v_i^*(k_i^o)$ in the baseline economy who switch from renting a property in the dual-rate economy to taking a mortgage to buy a house in the counterfactual single-rate economy. These households would rarely refinance, and thus expect to pay an average rate close to the reset rate in the baseline dual-rate economy, which raises the costs of taking on a mort-

Figure 3.5.1: Change in Loan Size at Origination

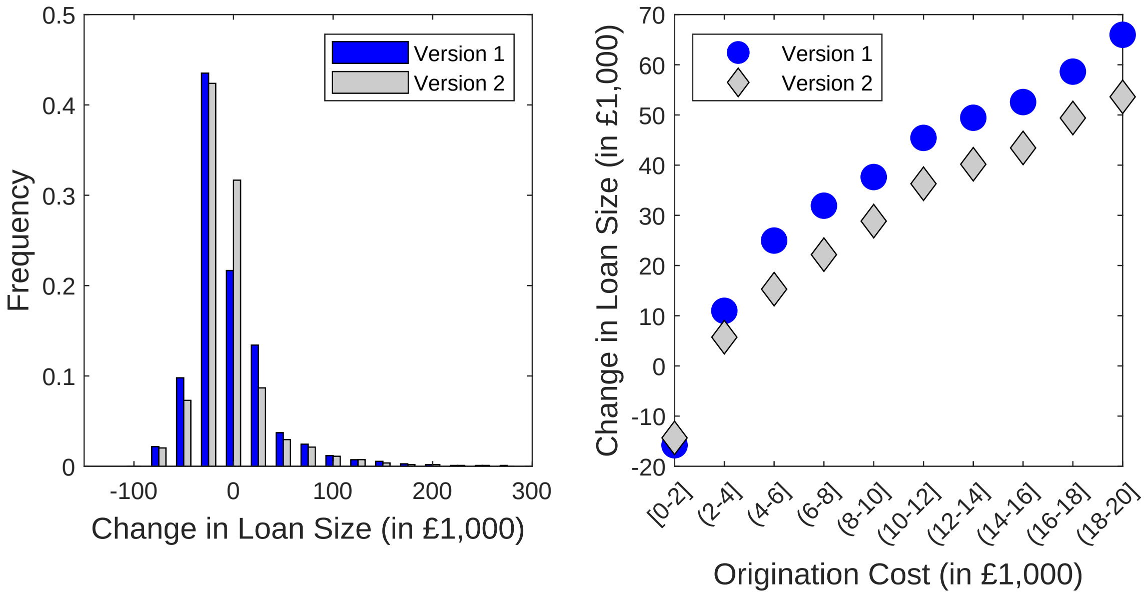

Notes: The left panel displays the distribution of the changes in loan sizes at origination between the counterfactual economy with constant interest rates and the baseline economy with discounted and reset rates. The right panel displays the average change in loan sizes for households with different values of their origination costs (in bins of £2,000). All statistics displayed are computed including only households who either participate in the mortgage market in the baseline dual-rate economy, or in the counterfactual single-rate economy, or in both.

gage. These households, therefore, choose to rent in the dual-rate world, but since they pay a lower rate in the counterfactual single-rate economy, they choose to buy a house by taking on a mortgage in the counterfactual. The mass of these households, on net, is greater than the mass of households with low $k_i^o = k_i$ who pay an average rate close to the discounted rate in the baseline economy but pay a higher rate in the single-rate economy. Such low k_i^o households switch from owning with a mortgage in the dual-rate economy to renting a property in the counterfactual single-rate economy, but their exit from the mortgage market is more than offset by new entrants into the single-rate mortgage market.

The second row of Panel B shows that the average initial loan size decreases by 2.37 percent of the average loan size in the baseline case, corresponding to a mortgage size reduction of £$4,653$. The main reason for this decline is the change in the composition of borrowers: marginal households who enter the mortgage market in the single-rate economy have smaller loan sizes than inframarginal households whose participation does not change.

More generally, the change in the average loan size combines borrowers who increase their mortgage amounts with borrowers who decrease them. The left panel of Figure 3.5.1 shows the full distribution of the changes in mortgage amounts of those households who participate in the mortgage market in the baseline dual-rate economy, in the counterfactual single-rate economy, or both. The heterogeneity of the changes in mortgage amounts at origination is apparent, with decreases in mortgage amounts more concentrated than increases.

The right panel of Figure 3.5.1 helps to rationalize the asymmetric adjustment in loan sizes. It displays how the average change in mortgage size varies with the origination cost k_i^o of refinancing costs. Borrowers with the lowest $k_i^o = k_i$ pay an interest rate close to 650 bps in the estimated mortgage market, because they almost always refinance, but they pay 683 bps in the counterfactual market with a constant interest rate. This higher rate induces them to reduce their loan sizes. In contrast, borrowers with the highest k_i^o pay an interest rate close to 759 bps in the baseline market, because they never refinance, but pay 683 bps in the counterfactual market. As a result, these borrowers increase their loan sizes. The increases in loan sizes are more dispersed than the decreases in loan sizes, because there is a bigger difference between the rates in the dual- and single-rate worlds paid by households with high k_i^o than that between the interest rates in the economies paid by those with low k_i^o.

The right panel of Figure 3.5.1 shows that version 2 displays similar qualitative patterns to those of version 1. The key difference is that the magnitudes of all adjustments in version 2 are smaller than those in version 1. The reason is that households forecast their future interest rates when they choose their loan size at origination. The uncertainty over the level of refinancing costs is larger in version 2 than in version 1, and thus households' interest rate forecasts converge to an intermediate value between the discount and reset rate. These forecasts resemble the counterfactual interest rate that we consider, and thus households respond less in their intensive margin of adjustment in version 2 than in version 1.

The third row of Panel B in Table 3.5.1 reports that the standard deviation of initial loan sizes declines quite substantially, by 4.04 percent of the standard deviation of the initial loan size (corresponding to £$4,978$) in the estimated baseline model. The reason is that one dimension of household heterogeneity, namely k_i, contributes to the determination of the loan size in the baseline model with refinancing. However, this dimension of heterogeneity becomes irrelevant when interest rates are constant. More specifically, the previous arguments suggest—and Figure 3.5.1 shows—that borrowers with larger loans in the baseline economy decrease their loan sizes in the counterfactual, whereas borrowers with smaller loans in the baseline economy increase their loan sizes in the counterfactual with constant interest rates and no refinancing. A constant, common interest rate thus pushes loan sizes to be more homogeneous.

The decline in initial loan size and the increase in the number of mortgages together combine to increase aggregate mortgage debt by 4.13 percent relative to the model with dual rates. Cross-subsidies are eliminated in the counterfactual, and one consequence of this change is that the mortgage market increases in size, although the effect is tempered by the opposing effects on the extensive and intensive margins.

The fourth and fifth rows of Panel B report that the patterns in the initial loan size distribution described above transfer to the aggregate loan balance distribution (i.e., including different cohorts of mortgages), with one additional subtle effect. In the baseline economy, on average, borrowers who originate large loans pay lower rates than borrowers with small loans. Hence, as loans amortize over time, the loan balances of borrowers with large loans tend to decline at a faster rate than the loan balances of borrowers with small loans, which compresses the distribution of loan balances over time. This force is absent in the counterfactual single-rate economy as all borrowers pay the same rate. Hence, the standard deviation of loan balances (normalized by that observed in the baseline economy) reported in the last row is slightly larger than the standard deviation of initial loan balances (also normalized with respect to the baseline economy) reported in the third row.

Finally, the last row of Panel B summarizes all the changes in a single money-metric ex-ante measure of consumer surplus, calculated for each household as $\max\left(W_0(v_i, k_i^o), \frac{\bar{u}}{1-\beta}\right)$. Consumer surplus increases by 3.94 percent in the single-rate economy relative to the dual-rate economy.

Panel B, Income Groups. The cases with multiple groups allow us to explain some of the observable heterogeneity in refinancing rates across income groups and geographies of the UK with heterogeneity in preferences v_i and costs k_i. These richer cases help us to evaluate whether and how the shift to a single mortgage rate structure leads to different outcomes for households in these groups.

Columns (3) and (4) of Table 3.5.1 report aggregated counterfactual estimates when the model is estimated using moments for different income groups. When we compare these aggregate statistics with those of the UK-wide model in columns (1) and (2), the differences appear small. The main difference is that column (3) exhibits a slightly larger adjustment in the extensive margin (i.e., the number of mortgages) than column (1), whereas the differences between columns (4) and (2) are minor. We now analyze how the results differ across income groups.

Figure 3.5.2 plots selected changes to mortgage market outcomes for each income group. The top-left panel shows that interest rates (in bps) are lower in the counterfactual economy for income groups up to roughly the 80th percentile of the income distribution in the sample, and are higher for the very highest income groups. This pattern is consistent with the regressive nature of the cross-subsidies in the dual-rate economy. The highest income group pays higher interest rates in the single-rate economy than the average rates they pay in the dual-rate economy. This is primarily because high-income households have larger loans, which gives them greater incentives to refinance promptly in the dual-rate economy.

The top-right panel adds refinancing costs to interest rates to calculate an all-inclusive mortgage cost. The pattern across groups is broadly similar to that seen in interest rates, with lower-income households paying lower mortgage costs under the counterfactual single-rate economy, whereas higher-income households pay mortgage costs similar to the costs that they pay in the baseline dual-rate economy.

The bottom-left panel shows that these changes in interest rates translate into an aggregate increase in the number of mortgages. Critically, the percentage increase is larger for lower-income groups, and minimal for the highest-income groups. The magnitude of the percentage increase is considerable for low-income households, because their homeownership rate is low in the baseline dual-rate economy (Table 3.2.2 reports that it equals 50% among the lowest-income households), and our model suggests that the design of the mortgage market is a contributing factor. In the single-rate market, there is a substantially greater entry of these low-income households into the housing and mortgage markets. As expected, version 2 displays smaller adjustments than version 1 for all income groups, because, when originating their mortgages in the baseline dual-rate economy, households can tailor their loan amounts to their future costs less precisely in version 2 than in version 1, because of their noisier information about their future costs. However, the qualitative pattern of the adjustments across income groups appear robust to the differential household information in the two versions.

The bottom-right panel plots the average percent differences between initial loan sizes in the single-rate economy and those in the dual-rate economy. While there are also important changes within groups, the across-group comparison highlights that higher-income groups adjust their average initial loan size downward more than lower-income groups. The adjustment in the average initial loan size of the highest-income group is sizable, because many of these borrowers—i.e., a larger fraction than among lower-income groups—almost always refinance and thus suffer a substantial increase in the interest rate that they pay, from $r_d = 650$ to $r_c = 683$. These loan size adjustments across income groups are very similar in the two versions of the model.

Overall, these panels suggest that the richer model with greater household heterogeneity across the income distribution implies that higher-income households pay lower rates and lower all-in mortgage costs than lower-income households in the current dual-rate structure. These patterns are consistent with the idea that the dual-rate structure fosters regressive cross-subsidies. The bottom panels of Figure 3.5.2 suggest that different income groups would respond to a single-rate structure with different types and levels of adjustments on both the intensive and extensive margins. In particular, raising participation in the mortgage market is the main adjustment for lower-income groups, whereas lowering initial loan sizes is the main adjustment for higher-income groups.

Finally, the changes in consumer surplus confirm that all income groups would benefit in the single-rate economy relative to the dual-rate economy, because they either pay lower interest rates or they save the refinancing costs $k_{i,t}$. Lower-income groups attain larger percentage-change increases in consumer surplus than higher-income groups.

Panel B, Regions. Columns (5) and (6) of Table 3.5.1 report aggregated counterfactual estimates when the model is estimated with parameters and moments for different UK regions. Once again, as with the model which incorporates greater heterogeneity across income groups, the aggregate statistics reported in columns (5) and (6) are remarkably similar to those of the UK-wide estimation reported in columns (1) and (2).

Figures 3.5.3 and 3.5.4 present maps that display some of the changes to mortgage market outcomes across different UK regions, for the models with more- and less-precise information about the persistent component of future refinancing costs at origination (i.e., versions 1 and 2), respectively. In each panel, darker colors indicate larger (positive) changes in the counterfactual market with constant interest rates when compared with the baseline dual-rate economy with discounted and reset rates.

The top-left map displays the change in average interest rates paid on mortgages, reported in bps. Households in the more prosperous regions of Greater London, the South East of England, and the East of England experience the largest increases in mortgage rates, whereas households in relatively less well-off regions and devolved administrations such as Northern Ireland, Wales, and the North East of England would experience the largest decreases in rates when moving to a single rate. These regional patterns are consistent with the dual-rate structure featuring regressive cross-subsidies across UK regions.

The top-right plot displays all-inclusive mortgage costs that sum (paid) refinancing costs and interest rates. The geographic patterns once again point to the regressive patterns found for interest rates in the top-left panel, with Scotland, Northern Ireland, and North East of England paying lower mortgage costs, whereas southern regions of the UK pay higher mortgage costs, in the counterfactual with a single rate.

In the counterfactual equilibrium, as seen earlier in the case of income groups, households endogenously adjust their mortgage market participation as well as their mortgage amounts. The bottom-left plot shows the change in the number of mortgages, which broadly increases the most in regions and devolved administrations that experience the largest decrease in mortgage rates and costs, such as Scotland and Northern Ireland. In contrast, southern regions of the UK experience smaller adjustments to mortgage market participation.

The bottom-right map displays the changes in the average initial loan size rate across regions. The differences in these averages when moving to the single-rate world mask larger within-region changes. That said, once again, southern regions' average initial loan sizes do shrink considerably more than less well-off regions' average initial loan sizes in the counterfactual single-rate world. Overall, the bottom maps confirm the pattern that the change in the profile of interest rates affects mostly the extensive margin in lower-income regions, and mostly the intensive margin in higher-income regions.

Finally, all regions would enjoy higher consumer surplus in the single-rate economy relative to the dual-rate economy, because they either pay lower interest rates or they save the refinancing costs $k_{i,t}$, with lower-income regions experiencing higher percentage-change increases in consumer surplus than higher-income regions.

3.6 Conclusion

We develop a model of mortgage refinancing and structurally estimate it on rich and granular data from the UK mortgage market. Our model matches broad features of the data, and the parameters reveal considerable heterogeneity in mortgage refinancing costs across households, echoing findings in prior literature. We use the estimated parameters to uncover regressive cross-subsidies in this market by conducting a counterfactual comparison with an alternative mortgage contract that features a constant interest rate and no need for refinancing.

This approach allows us to quantify cross-subsidies in this market setting. Using 2015 data, we set annual interest rates in our main counterfactual single-rate equilibrium to lie approximately 25bps above the average discounted rate and 30bps below the average reset rate that borrowers are routinely rolled on to at the expiration of the discounted rate fixation period. These are material changes given the importance of mortgages to household budgets.

The counterfactual scenario features different adjustments by low- and high-income groups. Low-income households enter the mortgage market in greater numbers, and raise their loan balances in response to the lower interest rates that they pay on average and the elimination of refinancing costs. Essentially, low-income groups are penalized by the dual-rate structure because they have smaller loan balances. Hence, they are more likely to pay the high reset rate in the dual-rate economy compared with high-income households. In contrast, high-income households mainly take on smaller loans in the single-rate counterfactual economy in response to their inability to take advantage of the discounted rate.

These findings highlight an important and novel dimension of inequality that would be invisible without our structural approach: we find that changes in mortgage rates increase entry into the housing and mortgage markets for low-income households; they also tend to push loan sizes to be more uniform across high- and low-income households. The economic size of these responses is substantial, even when we conservatively assume that households only observe noisy signals about their ongoing refinancing costs. Our results suggest that simplifying the design of mortgage refinancing and eliminating the costs associated with refinancing can cause forward-looking (even if imperfectly informed) households to participate more extensively in the mortgage market.

Our work has both methodological and economic contributions beyond the specific context that we study. First, we believe that our structural approach to estimating financial cross-subsidies by comparing the current and counterfactual market structures is a useful way to provide a money-metric assessment of the impacts of heterogeneity in household inaction. This has potentially wider implications for the field of household finance, where such heterogeneity is widely prevalent in many markets including credit and insurance. Our findings on the regressive nature of these cross-subsidies highlight that other household finance settings where high-income households benefit more due to their larger stakes and their greater propensity to take action may also contribute to inequality. In a broader sense, our results on the distribution of financial cross-subsidies in this important market show that studying household finance is helpful for the agenda of identifying the sources and consequences of wealth inequality, a continuing concern for society.

Table 3.5.1: Market Outcomes with Constant Interest Rates

	UK-WIDE		INCOME GROUPS		REGIONS	
	(1)	(2)	(3)	(4)	(5)	(6)
PANEL A: CONSTANT INTEREST RATE=650 BPS						
NUMBER OF MORTGAGES	1.10	1.11	1.15	1.10	1.12	1.10
MEAN INITIAL LOAN AMOUNT	1.06	1.05	1.05	1.06	1.05	1.05
STANDARD DEVIATION INITIAL LOAN AMOUNT	1.05	1.05	1.01	1.02	1.04	1.04
MEAN LOAN BALANCE	1.05	1.04	1.05	1.06	1.05	1.05
STANDARD DEVIATION LOAN BALANCE	1.06	1.05	1.02	1.03	1.04	1.04
CONSUMER SURPLUS	1.12	1.13	1.12	1.12	1.11	1.11
PANEL B: CONSTANT INTEREST RATE=683 BPS						
NUMBER OF MORTGAGES	1.06	1.07	1.10	1.07	1.07	1.06
MEAN INITIAL LOAN AMOUNT	0.98	0.97	0.96	0.97	0.97	0.97
STANDARD DEVIATION INITIAL LOAN AMOUNT	0.96	0.96	0.92	0.92	0.95	0.95
MEAN LOAN BALANCE	0.98	0.97	0.97	0.97	0.97	0.97
STANDARD DEVIATION LOAN BALANCE	0.97	0.96	0.93	0.94	0.96	0.96
CONSUMER SURPLUS	1.04	1.05	1.03	1.03	1.03	1.03
PANEL C: CONSTANT INTEREST RATE=700 BPS						
NUMBER OF MORTGAGES	1.04	1.04	1.07	1.05	1.04	1.04
MEAN INITIAL LOAN AMOUNT	0.94	0.94	0.92	0.92	0.93	0.93
STANDARD DEVIATION INITIAL LOAN AMOUNT	0.91	0.91	0.87	0.88	0.90	0.90
MEAN LOAN BALANCE	0.94	0.94	0.93	0.93	0.94	0.93
STANDARD DEVIATION LOAN BALANCE	0.92	0.92	0.89	0.89	0.92	0.91
CONSUMER SURPLUS	1.00	1.01	0.99	0.99	0.99	0.99
PANEL D: CONSTANT INTEREST RATE=759 BPS						
NUMBER OF MORTGAGES	0.95	0.95	0.95	0.97	0.94	0.95
MEAN INITIAL LOAN AMOUNT	0.83	0.82	0.80	0.79	0.82	0.81
STANDARD DEVIATION INITIAL LOAN AMOUNT	0.77	0.78	0.75	0.74	0.77	0.77
MEAN LOAN BALANCE	0.84	0.84	0.81	0.80	0.83	0.82
STANDARD DEVIATION LOAN BALANCE	0.80	0.80	0.77	0.76	0.79	0.79
CONSUMER SURPLUS	0.88	0.89	0.87	0.86	0.88	0.87

Notes: This table reports the statistics on the mortgage market in counterfactual markets with constant interest rates, as ratios of those of the estimated market with dual interest rates. The statistics in Panel A are calculated using a constant interest rate equal to the average discounted rate. The statistics in Panel B are calculated using a constant interest rate equal to the average interest rate equal to (3.17). The statistics in Panel C are calculated using a constant interest rate equal to the interest rate that satisfies the equal-revenue equation (3.19). The statistics in Panel D are calculated using a constant interest rate equal to the average reset rate. Odd-numbered columns correspond to version 1 of the model, and even-numbered columns correspond to version 2 of the model.

Figure 3.5.2: Changes in Market Outcomes by Income Groups

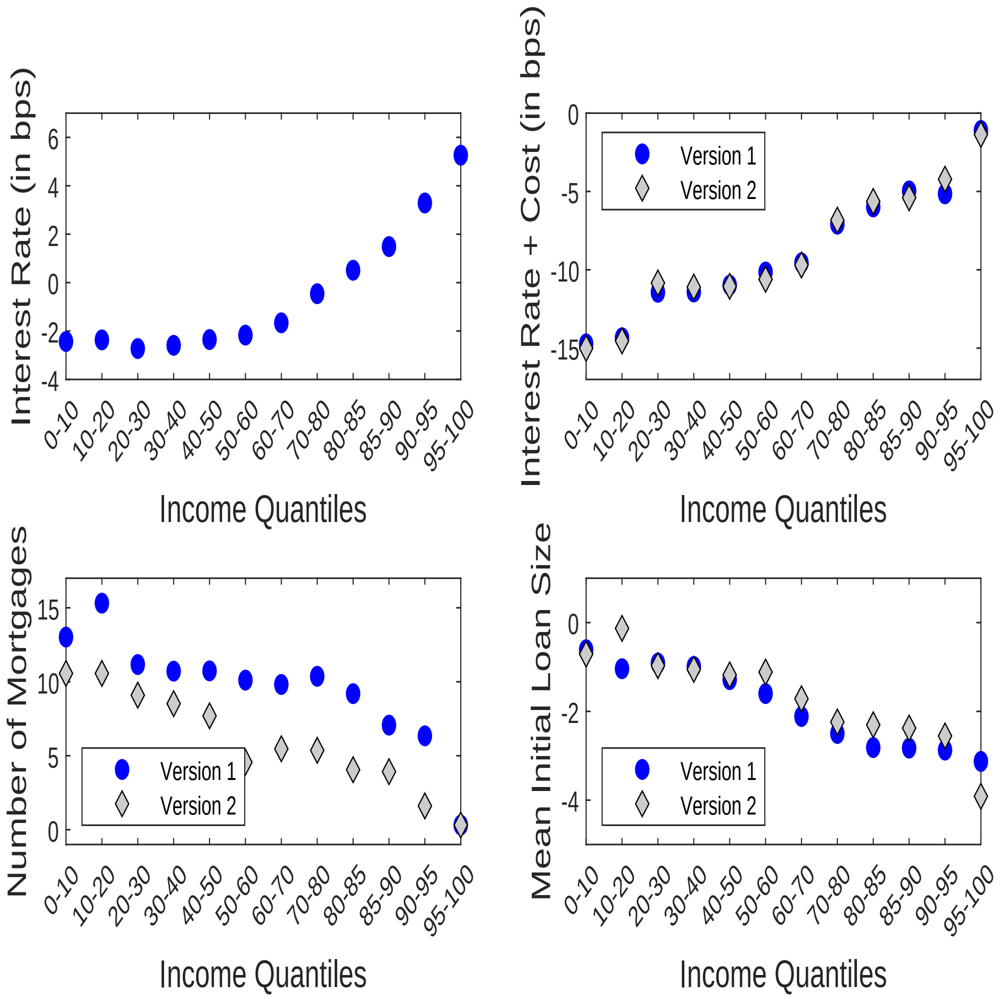

Notes: The top-left panel displays the change in interest rates (in bps); the top-right panel displays the change in mortgage costs, calculated as interest rates net of refinancing costs k; the bottom-left panel displays the percentage change in the number of mortgages; and the bottom-right panel displays the percentage change in the average initial loan size for each income group in the counterfactual case with a constant interest rate equal to $r_c = 683$ bps relative to the baseline case. Dark dots correspond to the model of column (3) and light diamonds correspond to the model of column (4) in Panel B of Table 3.5.1

Figure 3.5.3: Regional Changes, Version 1

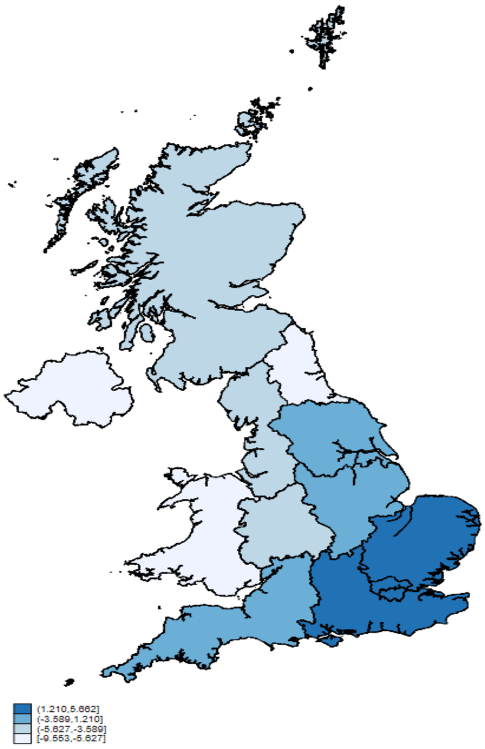

(a) Interest Rate

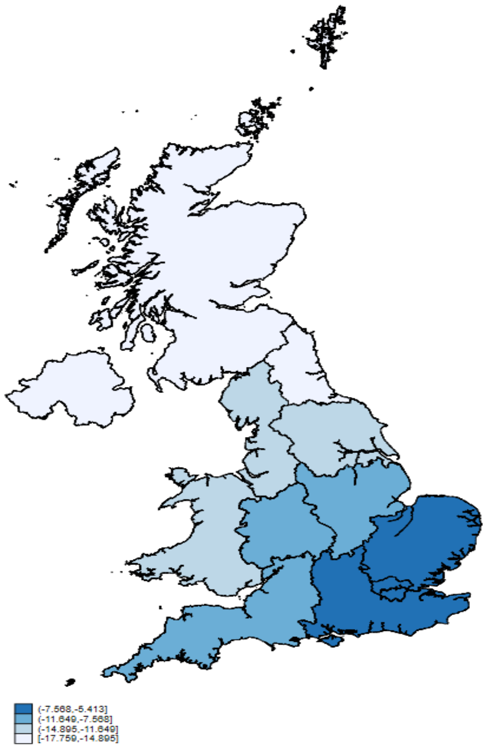

(b) Interest Rate + Refinancing Costs

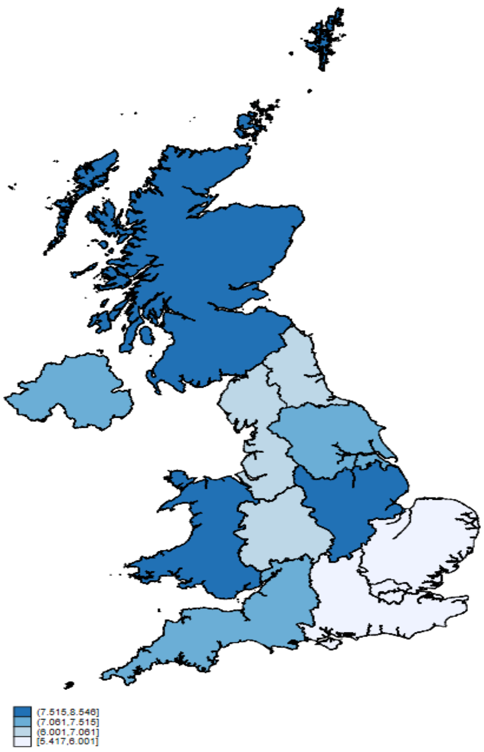

(c) Number of Mortgages

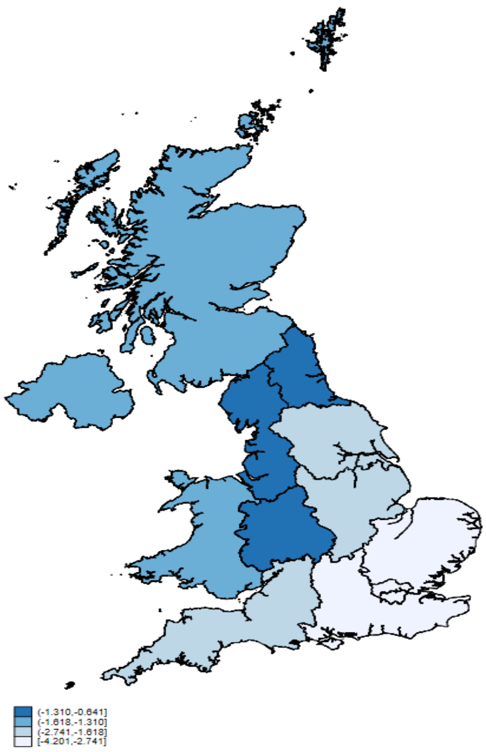

(d) Initial Loan Size

Figure 3.5.4: Regional Changes, Version 2

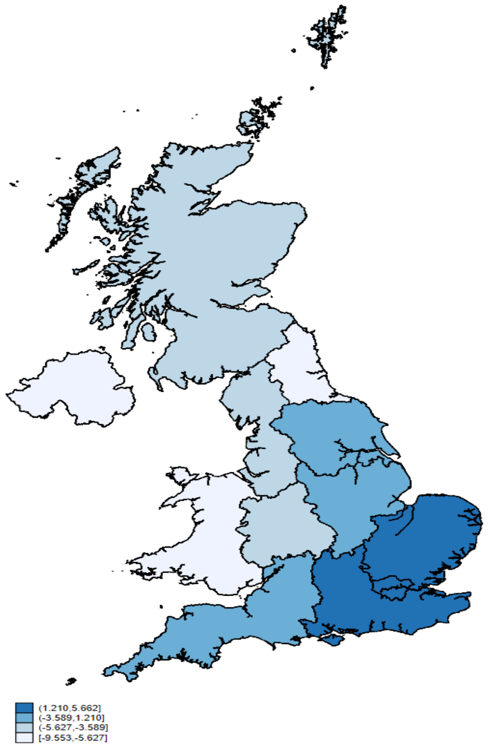

(a) Interest Rate

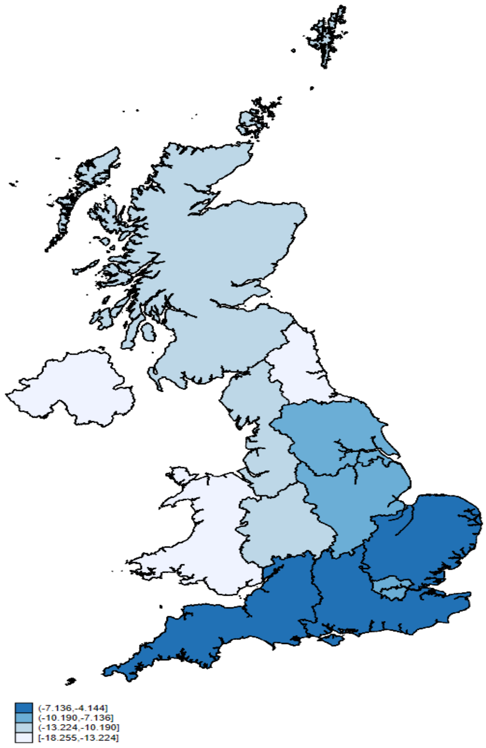

(b) Interest Rate + Refinancing Costs

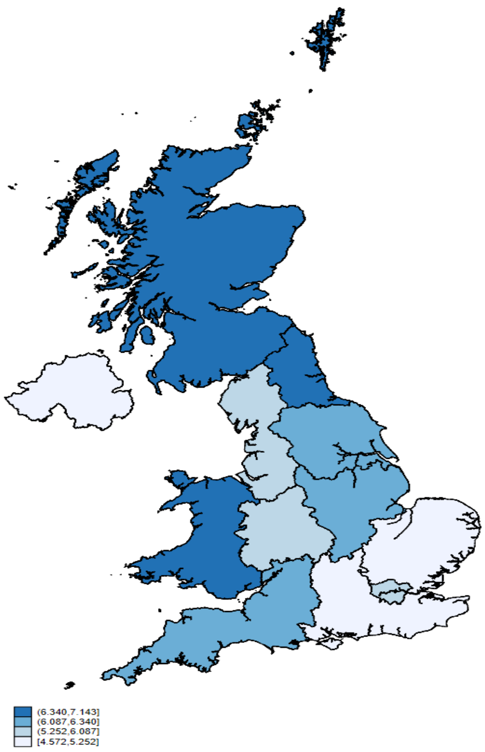

(c) Number of Mortgages

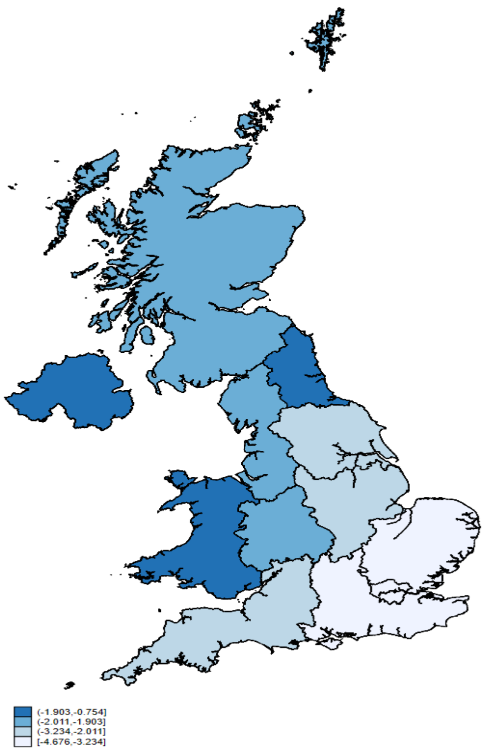

(d) Initial Loan Size

Appendices

3.A Dataset Construction

3.A.1 Primary Data Source: Stock of Owner-Occupier Mortgages

PSD007 includes loan-level information on the universe of mortgages in the owner-occupier or residential segment of the mortgage market. The owner-occupier segment includes first-time-buyers, home-movers, and refinancers who obtain mortgages from regulated financial institutions, such as deposit taking lenders and building societies. All regulated financial institutions are mandated by law to share this data with the FCA at a semi-annual frequency.

We have data on 6 PSD007 snapshots, reported half-yearly from mid-2015 to end-2017. Table 3.A.1 provides a description of the loan-level variables reported in PSD007 relevant to our study. In each snapshot, we observe the loan balance, original size of the loan, remaining term to maturity, original maturity, and interest rate for each mortgage recorded on the reporting date. The database also includes information on the type of interest rate and whether the mortgage is incentivized (i.e., on a discounted rate), and if so, the remaining period under the incentivized or discounted rate. The types of interest rates reported in the dataset are teaser, discounted, capped, standard variable rate, tracker and an unclassified other category. In some of our summary statistics, we use the reported interest rate to calculate a spread of the discounted rate over the yield on a nominal zero coupon UK Treasury maturing over the horizon over which the interest rate is fixed.[26]

[26] Discounted mortgages are fixed-term mortgages with a specified period under the discounted rate. In our model and in the data, mortgages automatically switch to the reset rate rate at the end of the discounted period. For example, in the case of a mortgage with a year remaining on the discounted rate, the spread is calculated over the yield on a nominal zero coupon UK Treasury bill maturing in a year. Reset rate mortgages are variable rate mortgages; the spread for reset rate mortgages is calculated based on the yield on short-term (6 month) UK Treasury bill.

Table 3.A.1: Description of Variables

Variable	Description
Current Loan Balance	Balance as on the date of reporting
Current Interest Rate	Interest rate charged on the mortgage
Spread	Spread over the yield on a nominal zero coupon bond maturing over a horizon comparable to the fixation period for interest rates (0 for mortgages under reset rate).
Original Loan Balance	Original size at the time of mortgage account opening date.
Original Term	Original term to maturity at the time of mort. account opening date.
Remaining Term	Remaining term to maturity.
Remaining Discounted Period	Remaining period under discounted rates.
Borrower Age	Borrower age as on the date of reporting.

Notes: The table above provides a brief description of mortgage level variables reported in PSD007 data relevant to our study.

Table 3.A.2 shows the overall balance of mortgages in 2015H1 by interest rate type and incentivized status. The table shows that a vast majority of the mortgages reported as being incentivized are also reported to be under teaser rates. Most mortgages under discounted and capped interest rates are also reported as being incentivized. However, there are a few discounted and capped mortgages which are reported as being non-incentivized and appear to have anomalous interest rates (we explain further below). We exclude such mortgages from our sample, and pool all incentivized mortgages reported as teaser, discounted, and capped interest rates into our discounted category in the paper.

Table 3.A.3 shows the average interest rate by interest rate type and incentivized status in the 2015H1 snapshot. Mortgages that we classify as discounted, i.e., incentivized mortgages on teaser, discounted and capped interest rates have lower average interest rates. Mortgages on reset rates (or Standard Variable Rates, SVRs) have higher average interest rates than these categories. There is a small group of mortgages on reset rates which are also reported as being incentivized, which bear interest rates comparable to that of the non-incentivized reset rate mortgages.[27] We treat all instances of mortgages on reset rates as non-incentivized.

Tracker mortgages are the remaining large category of mortgages. Their interest rates are benchmarked to the contemporaneous Bank of England base rate or LIBOR. Table 3.A.3 shows that the average interest rate of mortgages in this category are lower than other mortgage types. However, this category is distinct from the discounted rate mortgages, as these are not teaser rates fixed for a duration; they are subject to rate

[27] This is a data issue only in the 2015H1 snapshot.

Table 3.A.2: Mortgages in 2015H1: Total Balance by Interest Rate Type and Incentivized Status (in £ billions)

	Incentivized		
	No	Yes	Total
Teaser	11.4	442.6	454.0
Discount	1.3	7.3	8.7
Capped	0.0	0.5	0.5
SVR	208.5	6.1	214.6
Tracker	121.7	90.6	212.3
Other	39.0	0.1	39.1
Total	381.9	547.3	929.2

Notes: The table above shows the total balance in £ billions by type of interest rate, and whether the mortgage is reported as being incentivized in the mortgage snapshot for 2015H1.

Table 3.A.3: Mortgages in 2015H1: Average Interest Rate by Interest Rate Type and Incentivized Status

	Incentivized		
	No	Yes	Total
Teaser	5.83	3.35	3.48
Discount	3.04	3.31	3.26
Capped	4.02	2.91	2.99
SVR	3.79	3.63	3.79
Tracker	2.22	2.16	2.19
Other	2.88	2.80	2.88
Total	3.39	3.15	3.26

Notes: The table above shows the average interest rate by type of interest rate, and whether the mortgage is reported as being incentivized in the mortgage snapshot for 2015H1.

fluctuations, and there are rarely transitions from the reset and discounted rate category into this category. As the tracker category is relatively isolated from the other two categories and outside of our model, we restrict our study on cross-subsidies to mortgages under the discounted and reset rate categories.[28]

3.A.2 Borrower Incomes and House Prices

The PSD007 dataset does not include information on current borrower incomes, which are typically reported at mortgage origination. We obtain information at origination from the PSD001 dataset (a dataset similar to PSD007, but from earlier years). We use the same variable used to merge information across stock snapshots (since it uniquely identifies a mortgage) to merge the stock data with the loan origination data. We obtain the latest income reported to the lender at the time of origination (usually the first instance of the mortgage being issued, occasionally captured in a subsequent refinanc-

[28]The total number of mortgages, average interest rates, and outstanding balances reported in Tables 3.A.2-3.A.3 are before the data filtering and cleaning steps that we describe in this appendix and in the paper.

ing round), and scale it using local-area level income indices obtained from the Office of National Statistics to an estimate for 2015H1.

Importantly, the distribution of the year of the reported income recorded at the time of origination does not vary across regions or across income bins. This helps to validate the quality of our loan-level income data, and provides reassurance that our cross-subsidy estimate by income-quantiles in section 3.5 is not affected by any differential quality of reported income across UK regions.

Similar to income, house prices are typically appraised at origination. In order to compute current house prices, and in particular to obtain current loan-to-value ratios (LTV), we scale house prices observed at origination using local-area-level house price indices, reported by HM Land Registry. These house price indices are available at monthly frequency. In order to match the reporting frequency of the mortgage stock (PSD007), we use house price indices reported in June (for H1), and December (for H2 data). This approach is standard in the literature, and is consistent with lenders' own adjustments of loan-to-value ratios when households refinance.

3.A.3 Data Cleaning

In the preceding section, we discussed filtering out mortgages with anomalous interest rate types, tracker mortgages and mortgages under an unspecified "other" category. We implement further data cleaning steps to filter out observations with anomalous or inconsistent data on remaining discounted period, balance, interest rate, remaining term and borrower age, as well as borrowers who may not be able to refinance because they are underwater, in arrears, highly leveraged or have a loan balance that is too small.

Reported Remaining Discounted Period. Table 3.A.4 shows summary statistics for the remaining period on discounted rates (in years) for discounted mortgages across the six snapshots. The mean and standard deviation of remaining discounted period is consistent across the snapshots, except 2015H1. This is driven primarily by misclassification of reset rate mortgages as discounted mortgages in 2015H1, resulting in a mass of mortgages with remaining discounted periods greater than 10 years. These are cases where the remaining term is reported as the remaining discounted period, something not seen in other snapshots. We reclassify the reset rate (or SVR) mortgages with misre-

Table 3.A.4: Remaining Discounted Period in Years

(a) Raw database

	mean	sd	p10	p25	p50	p75	p90
Rem. discounted period (2015H1)	2.89	5.37	0.42	1.00	1.83	3.33	4.58
Rem. discounted period (2015H2)	2.21	2.27	0.33	0.92	1.75	3.00	4.42
Rem. discounted period (2016H1)	2.15	2.23	0.33	1.00	1.75	2.92	4.33
Rem. discounted period (2016H2)	2.09	2.20	0.42	0.92	1.67	2.83	4.25
Rem. discounted period (2017H1)	2.04	2.19	0.33	0.83	1.50	2.75	4.33
Rem. discounted period (2017H2)	2.11	2.21	0.33	0.83	1.58	3.00	4.50

(b) Filtered database

	mean	sd	p10	p25	p50	p75	p90
Rem. discounted period (2015H1)	2.11	1.52	0.42	1.00	1.83	3.08	4.25
Rem. discounted period (2015H2)	2.09	1.53	0.33	0.92	1.75	3.00	4.33
Rem. discounted period (2016H1)	2.06	1.51	0.42	1.00	1.75	2.92	4.25
Rem. discounted period (2016H2)	2.01	1.52	0.42	0.92	1.67	2.75	4.25
Rem. discounted period (2017H1)	1.96	1.57	0.33	0.83	1.50	2.75	4.25
Rem. discounted period (2017H2)	2.03	1.65	0.33	0.83	1.58	2.92	4.50

Notes: The above tables shows summary statistics for the remaining discounted period in months for discounted mortgages across the PSD007 snapshots. Panel (a) shows the summary statistics for the raw database; panel (b) shows the summary statistics after the filtering steps described in Section 3.A.3.

ported remaining discounted period as being not incentivized.

In addition, across all snapshots, there are few mortgages with remaining discounted periods less than -1 years, and greater than 11 years. We drop all such observations from the sample. Table 3.A.4 (b) shows that the distribution of remaining discounted period is similar across the different snapshots after implementing the filtering steps described above.

Reported Balance. Table 3.A.5 reports summary statistics for loan balances across snapshots. It shows that all the moments (including mean and standard deviation) for loan balances in 2017H2 are higher than that for other snapshots. This difference is driven by discounted rate mortgages; the loan balance moments for reset rate mortgages are stable across the snapshots.

We find that the high mean and standard deviation for discounted mortgages in 2017H2 is driven by misreported loan balances for two lenders. Hence, for the discounted mortgages issued by these two lenders in 2017H2, we replace the reported loan balance in 2017H2 with the estimated amortized loan balance based on the reported loan balance, remaining term, and discounted interest rate of 2017H1.[29]

[29]We estimate an amortized loan balance for 2017H2 only for discounted mortgages with at least 6 months on discounted periods in 2017H1. Further, we do this estimation only for mortgages which are on a capital and interest payment plan; i.e. we do not restate the 2017H2 loan balance for the small balance of interest only discounted mortgages in the stock of the two aberrant lenders. Further, we do not restate the loan balance of the newly issued discounted rate mortgages issued by these two lenders which leads

Table 3.A.5: Loan Balances for Raw and Filtered Databases

(a) Raw database

	mean	sd	p10	p25	p50	p75	p90
Balance (2015H1)	118,143	108,109	29,300	59,534	98,043	149,398	219,929
Balance (2015H2)	119,800	115,850	25,000	57,743	98,198	151,763	227,112
Balance (2016H1)	124,175	121,525	28,302	59,246	99,952	155,683	235,932
Balance (2016H2)	128,213	126,975	29,279	60,250	101,966	160,238	244,876
Balance (2017H1)	130,608	127,003	30,000	60,775	103,092	162,999	250,191
Balance (2017H2)	143,369	148,222	29,357	61,902	108,069	178,562	286,897

(b) Filtered database

	mean	sd	p10	p25	p50	p75	p90
Balance (2015H1)	123,325	98,092	38,770	64,821	101,620	152,765	223,988
Balance (2015H2)	127,332	105,483	38,758	65,237	103,424	157,122	233,834
Balance (2016H1)	130,092	111,117	38,309	65,061	104,278	160,214	241,596
Balance (2016H2)	133,558	116,289	38,336	65,680	106,060	164,432	250,009
Balance (2017H1)	134,998	117,715	37,984	65,622	106,807	166,905	254,782
Balance (2017H2)	140,451	125,369	37,953	66,386	109,479	173,774	269,100

Notes: The above tables shows summary statistics for the outstanding balance for mortgages across the PSD007 snapshots. Panel (a) shows the summary statistics for the raw database; panel (b) shows the summary statistics after the filtering steps described in Section 3.A.3.

Reported Interest Rate, Remaining Term, and Age. We drop all instances of negative interest rates, and winsorize interest rates at the 99.9 percentile point for each snapshot to address outliers which clearly arise from misreporting (for instance, interest rates of >1000%). We drop all instances of negative remaining terms, and winsorize the distribution at the 99.9 percentile point for each snapshot to address outliers which clearly arise from misreporting (for instance, remaining term of 9999 months). Finally, we drop all instances of reported negative age of borrowers.

3.A.4 Borrowers Potentially Ineligible to Refinance

Table 3.A.6 shows the proportion of mortgages/loans that are potentially ineligible for refinancing using specific criteria based on loan characteristics such as high LTV, low loan-balances or default status discussed in section 3.2 in the paper. Together, these filters account for 14% of mortgages in 2015H1, a figure that ranges from 13-14% across the snapshots.

Table 3.A.7 presents an alternate view of the effect of these criteria by showing the proportions of mortgages remaining after filtering by these criteria. The first row titled 'All' corresponds to the sample of loans in different snapshots for which we have non-missing data for all variables, and row 4 corresponds to the fraction of this sample

to slightly higher average loan balance and LTV for the 2017H2 snapshot. The higher loan balance for 2017H2 has no bearing on the cross-subsidy estimates in our paper based on data moments from the 2015H1 snapshot.

Table 3.A.6: Proportion of Loans Ineligible for Refinancing

	2015H1	2015H2	2016H1	2016H2	2017H1	2017H2
All	100%	100%	100%	100%	100%	100%
(1) LTV>=95	2.3%	1.9%	2.2%	2.4%	2.4%	3.6%
(2) Balance<=30000	6.5%	6.5%	6.7%	6.7%	6.9%	6.9%
(3) Non-performing	5.5%	5.0%	3.9%	3.9%	3.8%	3.6%
All excl. (1),(2),(3)	86.4%	87.2%	87.7%	87.4%	87.4%	86.3%

Notes: The table above shows the share of borrowers who fall under (1) to (3). "LTV" refers to the current loan-to-value ratio; "Balance" to the current loan balance; and "Non-performing" includes short-term arrears, loans in forbearance, and loans with a possession order.

Table 3.A.7: Sample With and Without Mortgages Unlikely to Be Able to Refinance

(a) Percentage of mortgages

	2015H1	2015H2	2016H1	2016H2	2017H1	2017H2
All	100%	100%	100%	100%	100%	100%
All excl. (1) LTV>=95	97.7%	98.1%	97.8%	97.6%	97.6%	96.4%
All excl. (2) Balance<=30000	93.5%	93.5%	93.3%	93.3%	93.1%	93.1%
All excl. (3) Non-performing	94.5%	95.0%	96.1%	96.1%	96.2%	96.4%
All excl. (1), (2), (3)	86.4%	87.2%	87.7%	87.4%	87.4%	86.3%

(b) R-r (spread)

	2015H1	2015H2	2016H1	2016H2	2017H1	2017H2
All	0.50	0.61	0.72	0.61	0.77	0.98
All excl. (1) LTV>=95 (3)	0.49	0.62	0.75	0.66	0.82	1.00
All excl. (2) Balance<=30000 (5)	0.51	0.62	0.73	0.62	0.79	0.99
All excl. (3) Non-performing (7)	0.46	0.56	0.68	0.57	0.73	0.94
All excl. (1), (2), (3)	0.45	0.58	0.72	0.62	0.80	0.96

Notes: The tables above show the share of borrowers (panel (a)) and the average interest rate spread between mortgages under reset (R) and teaser (r) rates (panel (b)) of mortgages excluding those under the filters (1) to (3). "LTV" refers to the current loan-to-value ratio; "Balance" to the current loan balance; and "Non-performing" includes short-term arrears, loans in forbearance, and loans with a possession order.

that remains after applying the filtering criteria to exclude borrowers that are potentially ineligible for refinancing. The 86.4% figure for 2015H1 corresponds to the sample on which we compute the cross-subsidies in our paper. Table 3.A.7 (b) shows that the spread between the average interest rate paid on reset rate and discount rate mortgages goes down with the application of these criteria.

3.A.5 Summary of Dataset Construction Steps

Table 3.A.8 shows a reconciliation from the owner-occupier mortgage stock of £ 989 mn in 2015H1 in the UK, to the sample of £ 470 mn mortgages relevant to our study. Column 2 in rows 2-8 corresponds to the effect in £ mn of each step discussed in Sections 3.A.1 - 3.A.4, while columns 3 and 4 show the size of the remaining data in £ mn and mn observations, respectively.

Table 3.A.8: Steps in Dataset Construction

	Drop (£M)	£Mn	Mn
Mortgage stock in 2015H1		**989**	**8.7**
Drop if non-unique postcode + d.o.b.	60	929	7.8
Drop if not reset/discounted	251	678	5.7
Drop discounted if reported as non-incentivised	13	665	5.6
Drop if not reset/discounted in future snapshots	68	597	5.0
Drop if inconsistent/anomalous	32	565	4.6
Drop if unreported income/region	53	512	4.1
Drop if ineligible to refinance	42	470	3.6
Sample for cross-subsidies		**470**	**3.6**

Notes: The above shows the total size of the mortgage stock in 2015H1 (row 1), and the effect of each step in data construction leading up to the sample used for cross-subsidies (row 9).

Figure B.1: Mortgage Flows across Discounted and Reset Rate Categories

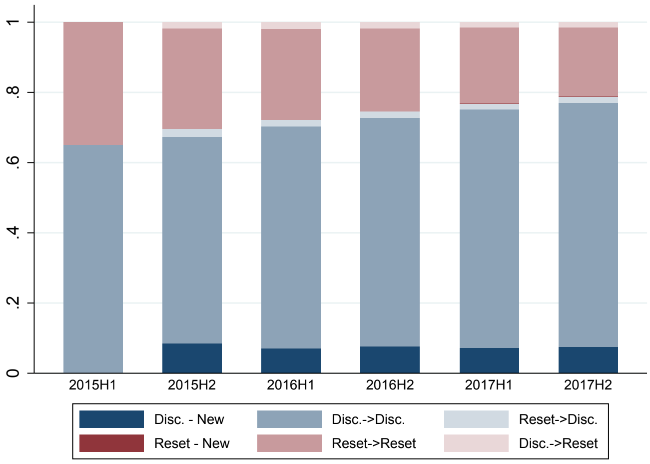

The figure above shows the proportion of mortgages under discounted rates (blue) and the reset rate (red) from the mortgage stock as reported at a half-yearly period from 2015H1 to 2017H2. The proportion of mortgages that are new to a snapshot are shown using a darker shade; and the proportion of mortgages that cross categories across snapshots are shown in a lighter shade.

3.B Merging Across Stock Snapshots

The high-quality disaggregated information in our database allows us to track mortgages across snapshots. In particular, we use the loan-level information on borrower date of birth and the 6-digit postcode to track mortgages across snapshots. These variables, when combined, provide a unique identifier for each mortgage.

We start with the 2015H1 snapshot as the base, and merge data from subsequent snapshots using this unique identifier. For each mortgage we can track whether it is discontinued between specific snapshots and whether it originated in any of the snapshots. Exploiting our ability to observe mortgages across snapshots, we also track whether a mortgage transitions across categories (discounted-to-reset rate or reset rate-to-discounted) between snapshots, or whether it continues in the same interest rate category. Across all 6 snapshots, the data track 6.00 million unique mortgages.

Figure B.1 provides a breakdown of discounted and reset rate mortgages in a given snapshot to show the cross-flows between these two mortgage groups across consecutive snapshots. For each snapshot starting 2015H2, we show the discounted (in blue) and reset (in red) mortgages by whether they are new (darkest shade), in the same category in the previous snapshot (lighter shade, e.g. discounted-to-discounted), or cross-flow from a different category in the previous snapshot (lightest shade, e.g. discounted-to-reset).

Figure B.2: Transition Probabilities Over Time

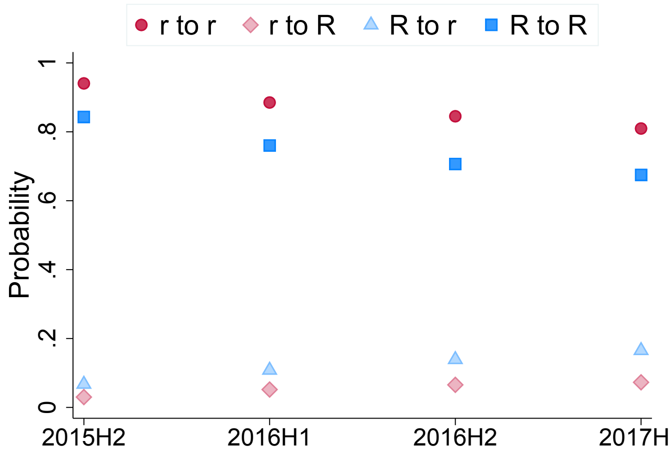

Notes: This figure displays transition probabilities between rate types, namely from discounted to discounted, discounted to reset, reset to discounted, and discounted to discounted, over time. The transition probabilities are measured between the share of mortgages in the 2015H1 stock, and 6 to 24 months later (2015H2, 2016H1, 2016H2, 2017H1). Omitted transition probabilities are from discounted or reset rate to account closure.

Observing these cross-flows by tracking mortgages across snapshots underlies one of the key moments we use to estimate our model—the transition probability of reset rate mortgages to be on discounted rates over a 2-year window. Figure B.2 shows the transition probabilities across different categories over 6-month to 24-month horizons for mortgages in the cleaned data for the 2015H1 snapshot (i.e., from 2015H1 up to 2015H2 and 2017H1, respectively). Transition from reset to discounted rate mortgages (R to r in the figure, light blue) increases from 6.73% of the mortgages after 6-months to 16.52% after 24-months.

In our sample, on average, around 345,000 mortgages are originated, and around 178,000 mortgage accounts are discontinued every 6 months. Given that new mortgages in each snapshot are predominantly discounted rate mortgages, this leads to an increase in the share of discounted mortgages from 65.0% in 2015H1 to 78.7% in 2017H2. We discuss potential arguments explaining this shift in Section 3.B.1. In comparison, the proportion of mortgages that flow across the two groups between snapshots is relatively stable.

Figure B.3: Average Interest Rate of Mortgages under Discounted and Reset Rates

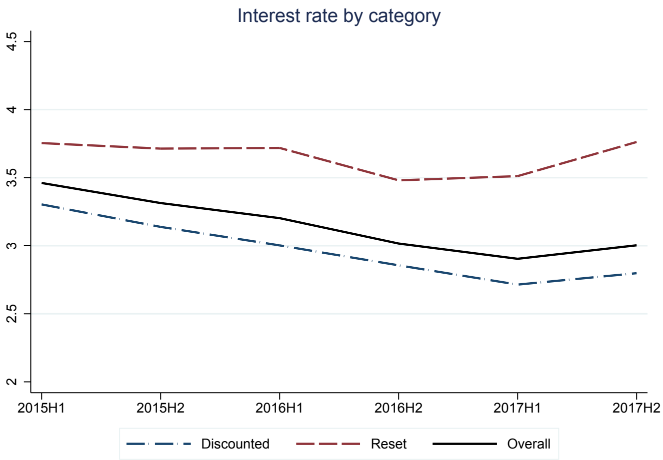

Notes: The figure above shows the equal-weighted interest rate for mortgages under discounted/reset rates in the snapshots of mortgage stock reported at a half-yearly period from 2015H1 to 2017H2.

3.B.1 Time-Series Evolution of Mortgage Stock

The UK mortgage market has experienced a number of changes from 2015H1 to 2017H2. In particular, Figure B.1 shows a decline in the share of mortgages paying the reset rate. We note that a number of factors could explain the drop in the number of mortgages on reset rates since 2015H1. First, Figure B.3 displays a decline in the average discounted rate, and an increase in the spread between average reset rates and discounted rates after 2015H1. Second, as reported in Financial Conduct Authority (2019b), there has been an increase in lenders' focus on retaining existing customers through internal switching, and an increased role of intermediaries in prompting borrowers to undertake beneficial switches.[B.1]

Table B.1 shows the averages of selected variables in the dataset across snapshots from 2015H1 to 2017H2. The average size of discounted rate loans has risen steadily over time, while the average size of loans on the reset rate has decreased. This is consistent with the change in refinancing incentives over time highlighted above. We also observe an increase in cash-out refinancing as evident in the average loan balance of discounted-to-discounted rate refinanced mortgages. The average remaining term

[B.1] We direct interested readers to more recent changes in the UK mortgage market aimed at facilitating switching at the time of refinancing. For instance, Financial Conduct Authority (2020a) reflects on increased use of technology and other remedies to facilitate switching; and recent policies have made it easier for for financial groups to switch customers from a group's closed book or lender to an active one (Financial Conduct Authority, 2020b), with the objective to make intra-group switching easier), and modified affordability assessments while refinancing for borrowers with up-to-date payments (Financial Conduct Authority, 2019a).

Table B.1: Summary Statistics over Mortgage Snapshots

Snapshots	2015H1	2015H2	2016H1	2016H2	2017H1	2017H2
Average loan size in £						
Discounted	140,647	143,611	145,431	147,815	149,792	152,278
Reset	112,692	111,176	109,285	108,468	107,038	105,799
Average remaining term in years						
Discounted	20.57	20.68	20.71	20.82	20.87	20.90
Reset	16.84	16.53	16.15	15.93	15.62	15.23
Average remaining term (value-weighted) in years						
Discounted	21.67	21.85	21.95	22.11	22.22	22.30
Reset	17.00	16.68	16.30	16.07	15.72	15.33
Average remaining discounted period in years						
Discounted	2.11	2.09	2.05	2.00	1.97	2.03
Average remaining discounted (value-weighted) period in years						
Discounted	2.10	2.10	2.05	1.99	1.95	2.02
Average interest rate						
Discounted	3.30	3.14	3.00	2.86	2.71	2.80
Reset	3.75	3.71	3.72	3.48	3.51	3.76
Average interest rate (value-weighted)						
Discounted	3.20	3.03	2.89	2.75	2.60	2.68
Reset	3.72	3.69	3.69	3.45	3.48	3.74
Average borrower age						
Discounted	41	41	41	41	41	41
Reset	44	45	45	46	46	46

Notes: The table above share summary statistics of mortgages for the stock snapshots from 2015H1 to 2017H2. The sample includes mortgages under two categories - those under discounted rates, and under the reset rate. Please see Appendix Table 3.A.1 for a description of the underlying variables.

on discounted rate mortgages rises from around 20.6 to 20.9 years (21.7 to 22.3 value-weighted), while the average remaining term on reset rate loans decreases through the sample period, from around 16.8 years to 15.2 years (17.0 to 15.3 value-weighted). The average remaining discounted period on discounted loans is 24 to 25 months in all snapshots of the data, reflecting the modal discounted period of 2 years observed in the data.[B.2] Finally, we observe an increase in the average interest rate gap between loans on reset rate and discounted rates over time, from 45bp (52bp value-weighted) in 2015H1, to 96bp (106bp value-weighted) in 2017H2. In all sample periods, the loan-balance weighted rate spread is higher than the equal-weighted rate spread. This effect stems mainly from larger discounted rate mortgages having lower rates on average, which is consistent with wealth-based heterogeneity in mortgage refinancing efficiency.

Legacy reset rate mortgages

As mentioned in Section 3.2.3, our data includes mortgages by two large lenders who offered to cap the reset rate for mortgages issued up to and during the 2007-09 financial crisis at 250 bp. In 2015H1, roughly 90% of all reset rate mortgages for these lenders were on these historically low rates, while the rest of the lenders had raised the interest rate paid when moving to reset rate mortgages to more than 400 bp. Consequently, excluding mortgages by these lenders more than doubles the spread between the average interest paid on reset and discounted mortgages—around 110 bp in 2015H1 (52 bp in our sample including the two large lenders) to 144 bp two years later in 2017H1.

[B.2] Mortgages under the discounted period are essentially fixed-rate loans. At origination, the most common discounted period is 2 years, followed by 5-year fixed-rate loans.

3.C Mortgage Moments by Region

Table C.1 shows summary statistics across UK regions and devolved administration; we evaluate cross-subsidies by UK regions in Section 3.5.

The regions are ordered by the average borrower income in the third column. The second column reports the population size, which shows that higher-income regions, such as London and the South East, are also the most populous. Higher-income regions tend to have higher balances as well a larger share of borrowers on the discounted rate, consistent with the patterns that we document across income quantiles in Table 3.2.2. One noticeable difference with the patterns in Table 3.2.2 is that the highest-income region, London, has the lowest homeownership rate, suggesting rich heterogeneity within- and across-regions. Finally, the average spread does not display an obvious pattern across regions.

Table C.1: Summary Statistics for the Mortgage Stock in 2015H1, By Region

REGIONS	POPULATION (1,000)	INCOME (£)	HOMEOWNERS (%)	BALANCE (£)	DISCOUNTED (%)	SPREAD (RESET-DISC.)
NORTHERN IRELAND	1,852	46,236	0.69	88,790	0.59	0.58
WALES	3,099	46,443	0.67	100,026	0.62	0.36
NORTH EAST (ENGLAND)	2,625	46,465	0.61	93,488	0.60	0.29
YORKSHIRE AND THE HUMBER	5,390	47,138	0.63	100,650	0.64	0.41
EAST MIDLANDS (ENGLAND)	4,677	49,331	0.67	106,786	0.64	0.29
NORTH WEST (ENGLAND)	7,175	49,439	0.64	103,406	0.63	0.38
WEST MIDLANDS (ENGLAND)	5,755	50,270	0.65	110,089	0.61	0.28
SCOTLAND	5,373	51,463	0.60	102,084	0.61	0.42
SOUTH WEST (ENGLAND)	5,472	55,248	0.67	128,260	0.67	0.30
EAST OF ENGLAND	6,076	62,041	0.67	146,888	0.69	0.47
SOUTH EAST (ENGLAND)	8,949	68,143	0.67	165,072	0.69	0.47
LONDON	8,667	85,598	0.49	207,592	0.69	0.84

Notes: The table above shows summary statistics of mortgages from the stock data reported in 2015H1, split by UK regions. Regional population for June 2015 is in 1,000s. Appendix Table 3.A.1 contains a description of the underlying variables.

3.D Prompt and Sluggish Refinancers

The goal of this Appendix is to check whether the dual-rate structure serves the function of screening borrowers on the basis of risk. To do so, we build on the analysis of Cloyne et al. (2019) by evaluating observable differences in the characteristics of the borrowers on different rates.

Table D.1 compares households who are on the reset and discounted rates, respectively. The table further splits the data on whether the current loan balance is above median (columns 1 and 2), or below median (columns 3 and 4).

We are particularly interested in any differences between borrowers on the discounted and reset rates evident in mortgage contract characteristics and borrower attributes that capture riskiness. Table D.1 shows that borrowers on the reset rate have smaller original and current loan balances and shorter remaining terms, consistent with rationally lower incentives to refinance.[D.1] However, the two sets of borrowers have comparable loan-to-value ratios and if anything, reset rate borrowers have slightly higher incomes and slightly lower loan-to-income ratios. These patterns are consistent for borrowers with both below and above median loan balances, suggesting that they are not just driven by smaller loan balances that reflect older contracts and/or lower leverage. These summary statistics are consistent with lenders earning greater risk-adjusted returns from borrowers on the reset rate, since they pay higher average rates, but have similar LTVs and LTIs, which are standard indicators of borrower riskiness.

We further study households on discounted and reset rates by assessing the subsample of households whose fixed-rate period expires in a subsequent snapshot of the mortgage stock. We then compare borrowers who refinance within 6 months of the fixed-rate expiration window ("prompt"), and "sluggish" borrowers who delay refinancing past this point.[D.2] We conduct this comparison using borrowers whose fixed-rate contract expired between 2015H1 and 2015H2. Table D.2 shows that households who refinance promptly have typically larger (origination and current) loan balances (cross-sectional means are displayed with cross-sectional standard deviations below in paren-

[D.1] By definition, as a result of being on the reset rate, the current interest rate paid by these borrowers in the table is higher than that of borrowers on the discounted rate.

[D.2] This definition of prompt and sluggish will not necessarily correspond to optimal and sub-optimal refinancing, as optimality depends on borrowers' specific circumstances, as we describe in more detail in the model section.

Table D.1: Characteristics by Rate Type and Loan Balance

	Disc. (high)	Reset (high)	Disc. (low)	Reset (low)
Original Loan Balance	205,634	187,929	82,467	86,351
	(117,247)	(97,522)	(27,204)	(32,456)
Original Term	24.84	23.97	21.35	23.04
	(6.52)	(6.42)	(7.59)	(6.94)
Current Loan Balance	196,766	176,047	72,114	69,137
	(112,791)	(90,371)	(20,954)	(21,058)
Current Interest Rate	3.22	3.72	3.41	3.78
	(0.93)	(0.97)	(0.96)	(0.98)
Remaining Term	22.56	17.53	18.15	16.37
	(7.02)	(6.80)	(7.86)	(7.01)
Borrower Age	39.47	44.35	42.07	44.49
	(8.78)	(9.32)	(10.63)	(10.56)
Current Income	74,855	78,754	38,665	41,894
	(261,183)	(109,242)	(37,069)	(48,369)
Current LTV	62.76	62.64	51.46	53.32
	(17.01)	(17.07)	(21.61)	(22.20)
Current LTI	2.98	2.65	2.21	1.97
	(1.61)	(2.37)	(1.16)	(1.06)
N	1,283,633	511,486	1,051,121	743,988

Notes: The table above compares means and standard deviations (in parentheses) for households across rate types, split by current loan balance. Columns 1 and 2 report values for households who have above median loan balances, while columns 3 and 4 report values for households who have below median loan balances.

theses), slightly longer remaining terms, higher income levels, but comparable LTI and LTV ratios.

Table D.2: Refinancing Decisions Due Between 2015H1 and 2015H2

	Prompt	Sluggish	Repaid/Moved
Original Loan Balance	155,221	120,994	158,886
	(107,043)	(81,381)	(113,046)
Original Term	23.82	23.33	23.15
	(6.97)	(7.79)	(7.78)
Current Loan Balance	143,670	111,216	143,441
	(102,432)	(77,468)	(107,417)
Current Interest Rate	3.33	3.75	3.38
	(0.98)	(1.10)	(1.03)
Remaining Term	20.79	20.14	19.77
	(7.47)	(8.36)	(8.26)
Borrower Age	40.66	41.13	41.13
	(9.42)	(11.05)	(10.73)
Current Income	59,104	51,332	65,300
	(64,198)	(120,794)	(69,529)
Current LTV	57.35	57.10	54.16
	(17.95)	(20.18)	(19.43)
Current LTI	2.66	2.50	2.48
	(1.44)	(1.40)	(1.36)
N	207,320	49,586	19,482

Notes: The table above compares means and standard deviations (in parentheses) for households whose fixed-rate contract was due to expire between 2015H1 and 2015H2. Column 1 reports households who refinanced within 6 months after their contract expiration date ("prompt"), column 2 reports households who did not refinance in that window ("sluggish"), and column 3 reports households who prepaid the loan and leave the sample.

3.E Option Value of Staying on Reset Rate

In this Appendix, we evaluate to what extent households not refinancing reflects an optimal option exercise, rather than their refinancing cost. As we explain below, reset rates do not come with prepayment penalties and so allow households to flexibly refinance in case interest rates go down. If there is an option value associated with waiting on the reset rate, it is difficult to interpret k as a behavioral parameter, the need for which is eliminated in the counterfactual. In order to evaluate the option value, we apply a standard approach in the literature and customize it to see if it can deliver an option value that justifies paying the reset rate for short periods, and find that it cannot, for reasonable parameter values.

3.E.1 Reset Rates and the Option To Refinance

UK mortgages have fixed rates and substantial prepayment penalties over the discounted fixation period, such that households typically lock in the rate and refinance next at the end of the fixation period. The reset rate, also called the "Standard Variable Rate" (SVR), in contrast, allows households to retain the option to refinance flexibly when interest rates go down. Figure E.1 shows the average reset rate and variation in 2-year government bond yields since 2011.

While reset rates do appear linked to underlying UK Treasury rates, they adjust at a slower rate and have less variability. Households incur a cost to retain the option to refinance, which is the difference between this slow-to-adjust and high reset rate R and the discounted rate r that they can currently access in the market. In what follows, we conduct a quantitative assessment of the net benefit of this refinancing option.

3.E.2 Optimal Refinancing Differential

In a standard refinancing framework with long-term fixed-rate mortgages, households rationally evaluate the present value of interest payments that they make under the new rate into which they refinance, and compare the payments they would make on this rate with those on the rate they would otherwise be in, accounting for any refinancing costs incurred, plus any difference between the value of the refinancing option that they give up, and the value of the new refinancing option that they acquire (Chen and Ling, 1989;

Figure E.1: Reset Rates Over Time

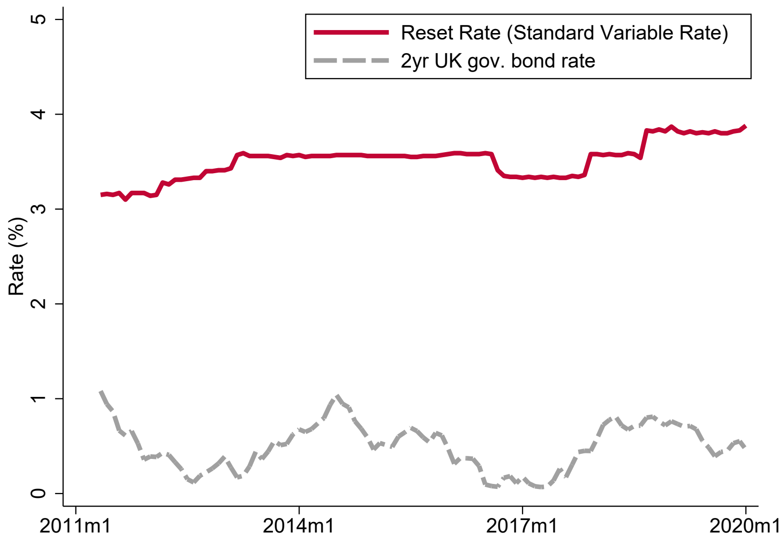

Notes: This figure shows the average Standard Variable Rate and 2-year UK Treasury bond yield at monthly frequency, as reported by the Bank of England Database.

Agarwal et al., 2013b). Households optimally exercise their option to refinance when the new rate is sufficiently lower than the rate they would otherwise bear (the "old rate"). This decision can be characterized using a "threshold," which is a specific value of the differential between the new and old rates beyond which it is rational to refinance.

Agarwal et al. (2013b) derive an analytical solution to this class of refinancing problems. They propose that households should refinance when the difference between the current mortgage interest rate (r_t) and the old rate (r_0), denoted by Δr, is greater than the optimal threshold Δr^*

$$\Delta r^* \equiv \frac{1}{\psi}\left(\phi + W\left(-\exp(-\phi)\right)\right),$$

where $W(.)$ is the principal branch of the Lambert W-function, $\psi = \frac{\sqrt{2(\rho+\lambda)}}{\sigma_r}$, and $\phi = 1 + \psi(\rho + \lambda)\frac{\kappa/M}{(1-\tau)}$. The optimal threshold depends on the real discount rate ρ, the expected real rate of exogenous mortgage repayment λ, the standard deviation of the mortgage rate σ_r, and κ/M, the ratio of refinancing cost and outstanding loan balance.

To make this framework applicable for UK borrowers considering whether or not to refinance into a discounted rate mortgage, we set κ to the median persistent component k_i of borrowers, £634, a conservative estimate of total refinancing cost, and M

as the average loan balance £130,871. σ_r is set to the historical standard deviation of 2-year UK real rates, at 0.0193. We further follow Agarwal et al. (2013b) to compute the rate of mortgage repayment as

$$\lambda = \mu + \frac{r_0}{\exp\left[r_0 T\right] - 1} + \pi,$$

where μ is the (annual) probability of prepayment, T is the remaining loan maturity, which we set to 30 years, and π is the rate of inflation, which we set to 2%. We set μ to 0.5 to capture the expected holding period of 2 years.[E.1]

For this representatively calibrated household, the optimal refinancing differential that we obtain is 111 basis points. This estimate provides a quantitative sense of by how much interest rates need to decrease for the refinancing option to be valuable, i.e. to be "in the money". For comparison, Figure E.2 shows how the optimal refinancing differential would vary by loan size. Intuitively, since the refinancing cost is fixed, the refinancing differential is decreasing in loan size, as the refinancing benefit is scaled by the loan balance. The majority of loans in the stock of mortgages have loan balances that are smaller than £200,000, and so require differentials between around 100 to 150 basis points.

Next, we simulate how likely it is that this refinancing threshold is hit, i.e., how likely the option is to be in the money, and the expected value of the option.

3.E.3 Simulation of Option Value of Refinancing

For a given old interest rate r_0, the optimal refinancing threshold calculated in the previous subsection (i.e., 111 basis points in our calibration) characterizes the level of the current mortgage rate r^* below which it is optimal to refinance. The expected value of the refinancing option is then

$$\begin{aligned} &\int_{-\infty}^{+\infty} (r_0 - r_s)\mathbb{I}_{(r_s \leq r*)} f(r_s) dr_s \\ &\quad = \int_{-\infty}^{r_s \leq r^*} (r_0 - r_s) f(r_s) dr_s. \end{aligned} \tag{E.1}$$

We can simulate the expected value of this expression by specifying a data-generating

[E.1] In this framework, the time over which the refinancing benefit accrues is primarily captured via the prepayment probability μ, rather than T.

Figure E.2: ADL Threshold Under Different Loan Sizes

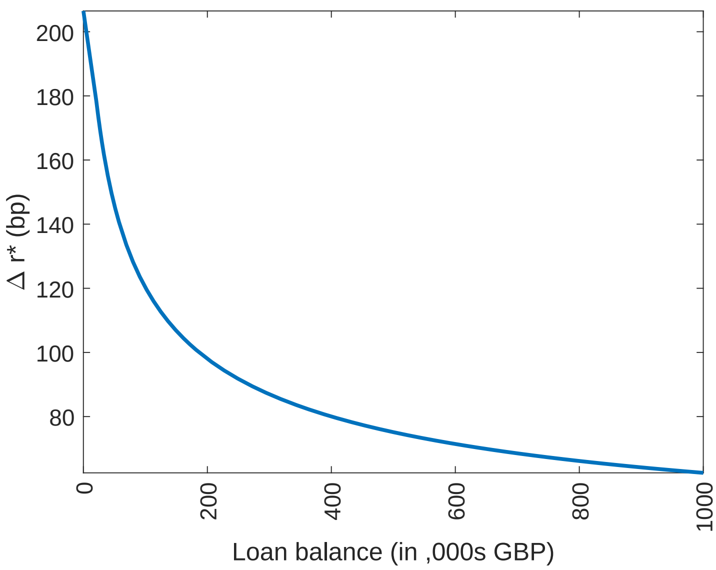

Notes: This figure plots the optimal refinancing threshold using the formula by Agarwal et al. (2013b) under the main calibration, when varying the loan balance.

process for interest rates. Suppose interest rates follow a standard AR(1) process:

$$r_t = (1 - \rho_r)\mu_r + \rho_r r_{t-1} + \epsilon_t, \tag{E.2}$$

where ϵ_t is a normally distributed white noise shock with mean zero and variance σ_ϵ^2, and ρ_r is the autocorrelation coefficient. The variance of the white noise shock is related to the variance of interest rates σ_r via $\sigma_\epsilon = \sqrt{\sigma_r^2 \cdot (1 - \rho_r^2)}$.

Households in our setting compare two options. The first is to take out a new 2-year fixed-rate contract right away (subsequently refinancing again after 2 years). The second is to stay on the reset rate and then refinance into a new 2-year fixed-rate contract whenever the optimal refinancing threshold is met. We simulate the expected NPV of such an option exercise, assuming that households can choose to wait up to 2 years to exercise the option to refinance. We then evaluate the NPV over a period of 4 years to ensure that we compare all households' NPV of option exercise over a similar period, regardless of when they fix again within the two-year window.

Figure E.3 shows simulations of the expected NPV of the option under different calibrations of the interest rate process and optimal refinancing thresholds. Panels (a), (b) and (c) show that the NPV is only positive if interest rates are substantially more volatile

Figure E.3: Expected NPV of Refinancing Option on Reset Rate

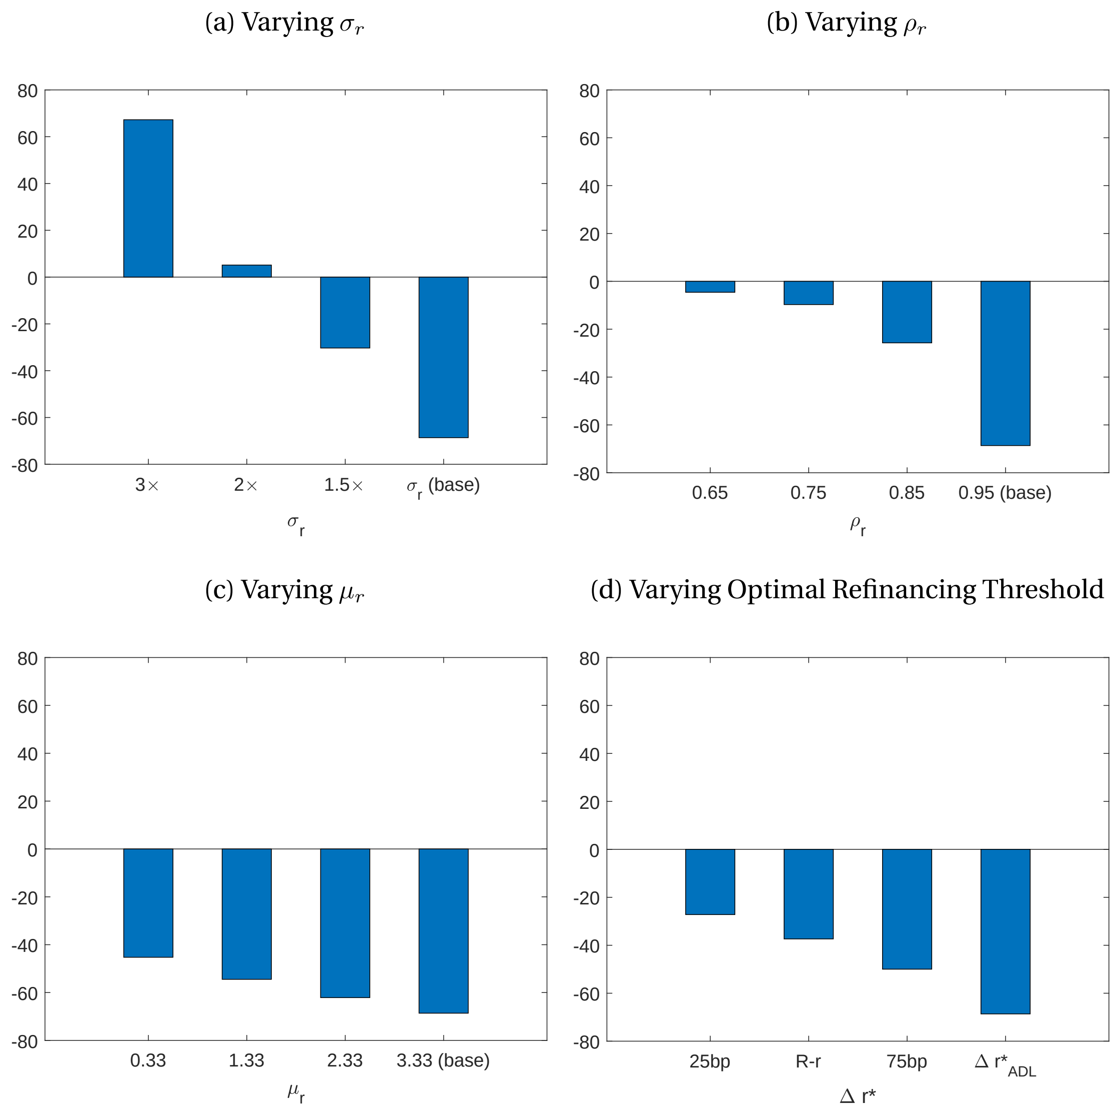

Notes: This figure displays the simulated net present value of the option to refinance when staying on the reset rate, under different calibrations of the interest rate process and the optimal refinancing threshold.

than they have been historically, by a factor of 2 to 3. Varying the persistence or long-run average of the interest rate process within reasonable ranges are not sufficient to yield a positive value. In addition, Figure E.3d shows the NPV for different assumptions about the refinancing threshold—lowering it to 25 and 75bp rather than calibrated ADL value of 111 bp. We also consider an alternative benchmark, which is a simple rule-of-thumb to refinance whenever the current interest rate makes up for the reset rate differential $R - r$. Under all these scenarios, we find that the expected value of the option remains negative.

Lastly, Figure E.4 shows that the expected value of the refinancing option is around

Figure E.4: Varying Optimal Refinancing Threshold and σ_r

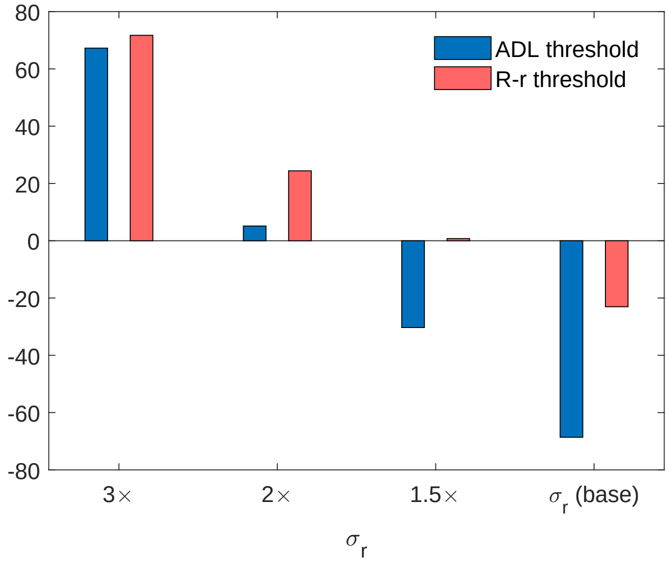

Notes: This figure displays the simulated net present value of the option to refinance when staying on the reset rate, under different calibrations of interest rate volatility and comparing the ADL optimal refinancing threshold with a threshold that corresponds to the difference between the reset and discounted rate.

20 basis points when households refinance based on the $R - r$ threshold, under a counterfactual interest rate volatility that is twice the historical average. Based on these simulations, we conclude that the option value of staying on the reset rate is not economically significant under a benchmark calibration of interest rates. This is because the cost of retaining the option is very high (this is the spread of the reset rate over the discounted rate) which penalizes waiting, and because the window over which the option can be exercised is relatively short, corresponding to the typical fixation window of 2 years.

Bibliography

Abraham, K. G., Haltiwanger, J., Sandusky, K., and Spletzer, J. (2019). The rise of the gig economy: fact or fiction? In *AEA Papers and Proceedings*, volume 109, pages 357–61.

Adams, A., Freedman, J., and Prassl, J. (2018). Rethinking legal taxonomies for the gig economy. *Oxford Review of Economic Policy*, 34(3):475–494.

Agarwal, S., Diao, M., Pan, J., and Sing, T. F. (2013a). Labor supply decisions of singaporean cab drivers.

Agarwal, S., Driscoll, J. C., and Laibson, D. I. (2013b). Optimal mortgage refinancing: a closed-form solution. *Journal of Money, Credit and Banking*, 45(4):591–622.

Agarwal, S., Rosen, R. J., and Yao, V. (2016). Why do borrowers make mortgage refinancing mistakes? *Management Science*, 62(12):3494–3509.

Akbarpour, M., Alimohammadi, Y., Li, S., and Saberi, A. (2021). The value of excess supply in spatial matching markets. *arXiv preprint arXiv:2104.03219*.

Allcott, H. and Taubinsky, D. (2015). Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market. *American Economic Review*, 105(8):2501–38.

Allen, J. and Li, S. (2020). Dynamic competition in negotiated price markets. Technical report, Bank of Canada Staff Working Paper.

Alvaredo, F., Chancel, L., Piketty, T., Saez, E., and Zucman, G. (2017). Global inequality dynamics: New findings from wid.world. *American Economic Review*, 107(5):404–409.

Andersen, S., Brandon, A., Gneezy, U., and List, J. A. (2014). Toward an understanding of reference-dependent labor supply: Theory and evidence from a field experiment. Working Paper 20695, National Bureau of Economic Research.

Andersen, S., Campbell, J. Y., Nielsen, K. M., and Ramadorai, T. (2020). Sources of inaction in household finance: Evidence from the danish mortgage market. *American Economic Review*, 110(10):3184–3230.

Angrist, J. D., Caldwell, S., and Hall, J. V. (2021). Uber versus taxi: A driver's eye view. *American Economic Journal: Applied Economics*, 13(3):272–308.

Armstrong, M. and Vickers, J. (2012). Consumer protection and contingent charges. *Journal of Economic Literature*, 50(2):477–493.

Ashenfelter, O., Doran, K., and Schaller, B. (2010). A Shred of Credible Evidence on the Long-run Elasticity of Labour Supply. *Economica*, 77(308):637–650.

Athey, S., Castillo, J. C., and Chandar, B. (2021). Service quality on online platforms: Empirical evidence about driving quality at uber. *Available at SSRN 3499781*.

Augenblick, N. and Rabin, M. (2019). An experiment on time preference and misprediction in unpleasant tasks. *Review of Economic Studies*, 86(3):941–975.

Ausubel, L. M. (1991). The failure of competition in the credit card market. *American Economic Review*, 81(1):50–81.

Badarinza, C., Campbell, J. Y., and Ramadorai, T. (2016). International comparative household finance. *Annual Review of Economics*, 8:111–144.

Badarinza, C., Campbell, J. Y., and Ramadorai, T. (2018). What calls to arms? international evidence on interest rates and the choice of adjustable-rate mortgages. *Management Science*, 64(5):2275–2288.

Belgibayeva, A., Bono, T., Bracke, P., Cocco, J., and Majer, T. (2020). When discounted rates end: The cost of a taking action in the mortgage market. Occasional Paper 54, Financial Conduct Authority.

Benetton, M. (2021). Leverage regulation and market structure: An empirical model of the uk mortgage market. *Journal of Finance*, 76(6):2297–3053.

Benetton, M., Gavazza, A., and Surico, P. (2021). Mortgage pricing and monetary policy. Technical Report 16456, Centre for Economic Policy Research.

Benhabib, J. and Bisin, A. (2018). Skewed wealth distributions: Theory and empirics. *Journal of Economic Literature*, 56(4):1261–1291.

Beraja, M., Fuster, A., Hurst, E., and Vavra, J. (2019). Regional heterogeneity and the refinancing channel of monetary policy. *The Quarterly Journal of Economics*, 134(1):109–183.

Berger, D., Milbradt, K., Tourre, F., and Vavra, J. (2021). Mortgage prepayment and path-dependent effects of monetary policy. *American Economic Review*, 111(9):2829–2878.

Berger, D., Milbradt, K., Tourre, F., and Vavra, J. (2023). Refinancing frictions, mortgage pricing and redistribution. working paper.

Berkouwer, S. B. and Dean, J. T. (2019). Credit and attention in the adoption of profitable energy efficient technologies in kenya.

Bernhardt, A., Campos, C., Prohofsky, A., Ramesh, A., and Rothstein, J. (2022). Independent contracting, self-employment, and gig work: Evidence from california tax data. Technical report, National Bureau of Economic Research.

Bernheim, B. D. and Rangel, A. (2004). Addiction and cue-triggered decision processes. *American economic review*, 94(5):1558–1590.

Bernheim, B. D. and Rangel, A. (2009). Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics. *The Quarterly Journal of Economics*, 124(1):51–104.

Bernheim, B. D. and Taubinsky, D. (2018). Behavioral public economics. In *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 1, pages 381–516. Elsevier.

Bertolini, A., Bogg, A., Colclough, C., Cole, M., Farrar, J., Ford, M., Gutierrez, C., Huws, U., Powdrill, T., Rix, M., Spencer, N., Sharp, T., Williamson, J., and Wishart, R. (2021). *Seven Ways Platform Workers are Fighting Back*.

Besley, T. and Case, A. (2000). Unnatural experiments? estimating the incidence of endogenous policies. *The Economic Journal*, 110(467):672–694.

Best, M. C., Cloyne, J. S., Ilzetzki, E., and Kleven, H. J. (2020). Estimating the elasticity of intertemporal substitution using mortgage notches. *The Review of Economic Studies*, 87(2):656–690.

Blundell, R. and MaCurdy, T. (1999). Labor supply: A review of alternative approaches. *Handbook of labor economics*, 3:1559–1695.

Boeri, T., Giupponi, G., Krueger, A. B., and Machin, S. (2020). Solo self-employment and alternative work arrangements: A cross-country perspective on the changing composition of jobs. *Journal of Economic Perspectives*, 34(1):170–95.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, 121(5):803–843.

Borjas, G. (1980). The relationship between wages and weekly hours of work: The role of division bias. *Journal of Human Resources*, 15(3):409–423.

Bresnahan, T. F. and Reiss, P. C. (1990). Entry in monopoly market. *The Review of Economic Studies*, 57(4):531–553.

Bresnahan, T. F. and Reiss, P. C. (1991). Entry and competition in concentrated markets. *Journal of political economy*, 99(5):977–1009.

Broughton, A., Gloster, R., Marvell, R., Green, M., Langley, J., and Martin, A. (2018). *The Experiences of Those in the Gig Economy*.

Brunnermeier, M. K. and Nagel, S. (2008). Do wealth fluctuations generate time-varying risk aversion? micro-evidence on individuals. *American Economic Review*, 98(3):713–736.

Byrne, S., Devine, K., King, M., McCarthy, Y., and Palmer, C. (2023). The last mile of monetary policy: Inattention, reminders, and the refinancing channel. Technical report, National Bureau of Economic Research.

Caldwell, S. and Danieli, O. (2020). *Outside options in the labor market*. Pinhas Sapir Center for Development, Tel Aviv University.

Caldwell, S. and Harmon, N. (2019). Outside options, bargaining, and wages: Evidence from coworker networks. *Unpublished manuscript, Univ. Copenhagen*, pages 203–207.

Caldwell, S. and Oehlsen, E. (2018). Monopsony and the gender wage gap: Experimental evidence from the gig economy. *Massachusetts Institute of Technology Working Paper*.

Caldwell, S. and Oehlsen, E. (2021). Gender differences in labor supply: Experimental evidence from the gig economy. *Unpublished*.

Camerer, C., Babcock, L., Loewenstein, G., and Thaler, R. (1997a). Labor supply of new york city cabdrivers: One day at a time. *The Quarterly Journal of Economics*, 112(2):407–441.

Camerer, C., Babcock, L., Loewenstein, G., and Thaler, R. (1997b). Labor Supply of New York City Cabdrivers: One Day at a Time*. *The Quarterly Journal of Economics*, 112(2):407–441.

Campbell, J. Y. (2006). Household finance. *The Journal of Finance*, 61(4):1553–1604.

Campbell, J. Y. (2013). Mortgage market design. *Review of Finance*, 17(1):1–33.

Campbell, J. Y., Ramadorai, T., and Ranish, B. (2019). Do the rich get richer in the stock market? evidence from india. *American Economic Review: Insights*, 1(2):225–240.

Capen, E. C., Clapp, R. V., and Campbell, W. M. (1971). Competitive bidding in high-risk situations. *Journal of petroleum technology*, 23(06):641–653.

Caplin, A. and Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203.

Castillo, J. C. (2020). Who benefits from surge pricing? *Available at SSRN 3245533*.

Chang, T. and Gross, T. (2014). How many pears would a pear packer pack if a pear packer could pack pears at quasi-exogenously varying piece rates? *Journal of Economic Behavior Organization*, 99(C):1–17.

Chen, A. H. and Ling, D. C. (1989). Optimal mortgage refinancing with stochastic interest rates. *Real Estate Economics*, 17(3):278–299.

Chen, K.-M., Ding, C., List, J. A., and Mogstad, M. (2020). Reservation wages and workers' valuation of job flexibility: Evidence from a natural field experiment. Technical report, National Bureau of Economic Research.

Chen, M. K., Chevalier, J. A., Rossi, P. E., and Oehlsen, E. (2017). The Value of Flexible Work: Evidence from Uber Drivers. NBER Working Papers 23296, National Bureau of Economic Research, Inc.

Chen, M. K., Currier, L., Rossi, P. E., and Chevalier, J. A. (2022). Suppliers and demanders of flexibility: The demographics of gig work. *Unpublished*.

Chen, M. K., Rossi, P. E., Chevalier, J. A., and Oehlsen, E. (2019). The value of flexible work: Evidence from uber drivers. *Journal of political economy*, 127(6):2735–2794.

Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: Theory and evidence. *American economic review*, 99(4):1145–77.

Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2002). Defined contribution pensions: Plan rules, participant choices, and the path of least resistance. *Tax Policy and the Economy*, 16:67–113.

Chou, Y. (2000). Testing Alternative Models of Labor Supply. Evidence from Taxi-Drivers in Singapore. Technical report.

Cloyne, J., Huber, K., Ilzetzki, E., and Kleven, H. (2019). The effect of house prices on household borrowing: A new approach. *American Economic Review*, 109(6):2104–2136.

Cohen, P., Hahn, R., Hall, J., Levitt, S., and Metcalfe, R. (2016). Using big data to estimate consumer surplus: The case of uber. Technical report, National Bureau of Economic Research.

Collins, B., Garin, A., Jackson, E., Koustas, D., and Payne, M. (2019). Is gig work replacing traditional employment? evidence from two decades of tax returns. *Unpublished paper, IRS SOI Joint Statistical Research Program*.

Cook, C., Diamond, R., Hall, J. V., List, J. A., and Oyer, P. (2021). The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. *The Review of Economic Studies*, 88(5):2210–2238.

Cornick, P., Lepanjuuri, K., and Wishart, R. (2018). *The Characteristics of Those in the Gig Economy.*

Crawford, V. P. and Meng, J. (2011). New york city cab drivers' labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income. *American Economic Review*, 101(5):1912–32.

Cullen, Z. and Farronato, C. (2021). Outsourcing tasks online: Matching supply and demand on peer-to-peer internet platforms. *Management Science*, 67(7):3985–4003.

Datta, N. (2019). Willing to pay for security: a discrete choice experiment to analyse labour supply preferences.

DeFusco, A. A. and Mondragon, J. (2020). No job, no money, no refi: Frictions to refinancing in a recession. *The Journal of Finance*, 75(5):2327–2376.

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic literature*, 47(2):315–72.

DellaVigna, S. and Malmendier, U. (2006). Paying not to go to the gym. *american economic Review*, 96(3):694–719.

DellaVigna, S. and Pollet, J. M. (2009). Investor inattention and friday earnings announcements. *The Journal of Finance*, 64(2):709–749.

Doran, K. (2014). Are long-term wage elasticities of labor supply more negative than short-term ones? *Economics Letters*, 122(2):208–210.

Dubal, V. (2019). An uber ambivalence: Employee status, worker perspectives, & regulation in the gig economy. *UC Hastings Research Paper*, (381).

Dubal, V. B. (2021). Economic security & the regulation of gig work in california: From ab5 to proposition 22. *European Labour Law Journal*, page 20319525211063111.

Dupas, P., Robinson, J., and Saavedra, S. (2015). The daily grind : Cash needs , labor supply and self-control.

Economides, N., Seim, K., and Viard, V. B. (2008). Quantifying the benefits of entry into local phone service. *the RAND Journal of Economics*, 39(3):699–730.

Einav, L. and Finkelstein, A. (2011). Selection in insurance markets: Theory and empirics in pictures. *Journal of Economic perspectives*, 25(1):115–38.

Ellison, G. (2005). A model of add-on pricing. *The Quarterly Journal of Economics*, 120(2):585–637.

Fagereng, A., Guiso, L., Malacrino, D., and Pistaferri, L. (2020). Heterogeneity and persistence in returns to wealth. *Econometrica*, 88(1):115–170.

Farber, H. S. (2005). Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers. *Journal of Political Economy*, 113(1):46–82.

Farber, H. S. (2008). Reference-dependent preferences and labor supply: The case of new york city taxi drivers. *American Economic Review*, 98(3):1069–82.

Farber, H. S. (2015). Why you Can't Find a Taxi in the Rain and Other Labor Supply Lessons from Cab Drivers*. *The Quarterly Journal of Economics*, 130(4):1975–2026.

Farrell, J. and Klemperer, P. (2007). Coordination and lock-in: Competition with switching costs and network effects. In Armstrong, M. and Porter, R., editors, *Handbook of Industrial Organization*, volume 3, pages 1967–2072. Elsevier.

Fehr, E. and Goette, L. (2007). Do workers work more if wages are high? evidence from a randomized field experiment. *American Economic Review*, 97(1):298–317.

Fields, G. S. (2013). Self-employment in the developing world. *Paper submitted to the High Level Panel on the Post-2015 development agenda. Cornell University and Iza.*

Financial Conduct Authority (2019a). Changes to mortgage responsible lending rules and guidance—feedback on CP19/14 and final rules. Policy Statement PS 19/27, Financial Conduct Authority. Accessed: 2021-10-19.

Financial Conduct Authority (2019b). Mortgage market study: Final report. Market Study MS 16/2.3, Financial Conduct Authority.

Financial Conduct Authority (2020a). FCA issues research on mortgage switching. https://www.fca.org.uk/news/news-stories/fca-issues-research-mortgage-switching. Accessed: 2021-10-19.

Financial Conduct Authority (2020b). Removing barriers to intra-group switching and helping borrowers with maturing interest-only and part-and-part mortgages. Policy Statement PS 20/11, Financial Conduct Authority. Accessed: 2021-10-19.

Financial Conduct Authority (2021). Mortgage prisoners review. Policy Paper CP 576, Financial Conduct Authority.

Fisher, J. (2022). The cost of labor supply biases. *Unpublished*.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, 129(4):1661–1710.

Gabaix, X. and Laibson, D. (2006a). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics*, 121(2):505–540.

Gabaix, X. and Laibson, D. (2006b). Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets. *The Quarterly Journal of Economics*, 121(2):505–540.

Gabaix, X., Laibson, D., Moloche, G., and Weinberg, S. (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, 96(4):1043–1068.

Ganserer, A., Gregory, T., and Zierahn, U. (2022). Minimum wages and the rise in solo self-employment.

Garin, A., Jackson, E., and Koustas, D. (2022). New gig work or changes in reporting? understanding self-employment trends in tax data. *Understanding Self-Employment Trends in Tax Data (May 19, 2022). University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2022-67).

Garin, A. and Koustas, D. (2021). The distribution of independent contractor activity in the united states: Evidence from tax filings. *Joint Statistical Research Program of the Statistics of Income Division of the IRS*.

Gavazza, A. and Lizzeri, A. (2021). Frictions in product markets. In Ho, K., Hortaçsu, A., and Lizzeri, A., editors, *Handbook of Industrial Organization*, volume 4, pages 433–484. Elsevier.

Goettler, R. L. and Clay, K. (2011). Tariff choice with consumer learning and switching costs. *Journal of Marketing research*, 48(4):633–652.

Goetzmann, W. N., Spaenjers, C., and Van Nieuwerburgh, S. (2021). Real and private-value assets. *The Review of Financial Studies*, 34(8):3497–3526.

Goldin, J. and Reck, D. (2017). Optimal defaults with normative ambiguity. *Available at SSRN 2893302*.

Goldman, T. and Weil, D. (2021). Who's responsible here? establishing legal responsibility in the fissured workplace. *Berkeley J. Emp. & Lab. L.*, 42:55.

Greenwald, D. L., Leombroni, M., Lustig, H., and Van Nieuwerburgh, S. (2021). Financial and total wealth inequality with declining interest rates. Working Paper 28613, National Bureau of Economic Research.

Grubb, M. D. and Osborne, M. (2015). Cellular service demand: Biased beliefs, learning, and bill shock. *American Economic Review*, 105(1):234–71.

Gruber, J. (1994). The incidence of mandated maternity benefits. *The American economic review*, pages 622–641.

Gruber, J. and Köszegi, B. (2001). Is addiction "rational"? theory and evidence. *The Quarterly Journal of Economics*, 116(4):1261–1303.

Gruber, J. and Kőszegi, B. (2004). Tax incidence when individuals are time-inconsistent: the case of cigarette excise taxes. *Journal of Public Economics*, 88(9-10):1959–1987.

Gruber, J. and Krueger, A. B. (1991). The incidence of mandated employer-provided insurance: Lessons from workers' compensation insurance. *Tax policy and the economy*, 5:111–143.

Guiso, L., Pozzi, A., Tsoy, A., Gambacorta, L., and Mistrulli, P. E. (2022). The cost of steering in financial markets: Evidence from the mortgage market. *Journal of Financial Economics*, 143(3):1209–1226.

Gurun, U. G., Matvos, G., and Seru, A. (2016). Advertising expensive mortgages. *The Journal of Finance*, 71(5):2371–2416.

Haggag, K. and Paci, G. (2014). Default tips. *American Economic Journal: Applied Economics*, 6(3):1–19.

Hall, J. V., Horton, J. J., and Knoepfle, D. T. (2021). Pricing in designed markets: The case of ride-sharing. Technical report, Tech. rep., Working paper, Massachusetts Institute of Technology.

Hall, J. V. and Krueger, A. B. (2018). An analysis of the labor market for uber's driver-partners in the united states. *Ilr Review*, 71(3):705–732.

Hamilton, B. H. (2000). Does entrepreneurship pay? an empirical analysis of the returns to self-employment. *Journal of Political economy*, 108(3):604–631.

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7):2643–82.

Handel, B. R. and Kolstad, J. T. (2015). Health insurance for" humans": Information frictions, plan choice, and consumer welfare. *American Economic Review*, 105(8):2449–2500.

Handel, B. R., Kolstad, J. T., Minten, T., and Spinnewijn, J. (2020). The social determinants of choice quality: evidence from health insurance in the netherlands. Technical report, National Bureau of Economic Research.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the econometric society*, pages 1029–1054.

Harris, S. D. and Krueger, A. B. (2015). *A Proposal for Modernizing Labor Laws for Twenty-First-Century Work: The" Independent Worker"*. Brookings Washington, DC.

Heidhues, P. and Kőszegi, B. (2018). Behavioral industrial organization. In Bernheim, B. D., Vigna, S. D., and Laibson, D., editors, *Handbook of Behavioral Economics - Foundations and Applications 1*, volume 1, pages 517–612. Elsevier.

Hoffman, M. and Burks, S. V. (2020). Worker overconfidence: Field evidence and implications for employee turnover and firm profits. *Quantitative Economics*, 11(1):315–348.

Honka, E. (2014). Quantifying search and switching costs in the us auto insurance industry. *The RAND Journal of Economics*, 45(4):847–884.

Hsiaw, A. (2013). Goal-setting and self-control. *Journal of Economic Theory*, 148(2):601–626.

Hubmer, J., Krusell, P., and Smith Jr, A. A. (2020). Sources of us wealth inequality: Past, present, and future. In *NBER Macroeconomics Annual 2020, volume 35*. University of Chicago Press.

Hurst, E., Keys, B. J., Seru, A., and Vavra, J. (2016). Regional redistribution through the us mortgage market. *American Economic Review*, 106(10):2982–3028.

Illanes, G. (2016). Switching costs in pension plan choice. Technical report, Northwestern University.

Jackson, E. (2019). Availability of the gig economy and long run labor supply effects for the unemployed. *Work. Pap., Stanford Univ., Stanford, CA Google Scholar Article Location*.

Jäger, S., Roth, C., Roussille, N., and Schoefer, B. (2022). Worker beliefs about outside options. Technical report, National Bureau of Economic Research.

Katsnelson, L. and Oberholzer-Gee, F. (2021). *Being the Boss: Gig Workers' Value of Flexible Work*. Harvard Business School.

Katz, L. F. and Krueger, A. B. (2017). The role of unemployment in the rise in alternative work arrangements. *American Economic Review*, 107(5):388–92.

Katz, L. F. and Krueger, A. B. (2019). The rise and nature of alternative work arrangements in the united states, 1995–2015. *ILR review*, 72(2):382–416.

Keane, M. P. (2011). Labor supply and taxes: A survey. *Journal of Economic Literature*, 49(4):961–1075.

Keys, B. J., Pope, D. G., and Pope, J. C. (2016). Failure to refinance. *Journal of Financial Economics*, 122(3):482–499.

Kolsrud, J. (2018). Sweden: Voluntary unemployment insurance.

Koustas, D. (2018). Consumption insurance and multiple jobs: Evidence from rideshare drivers. *Unpublished working paper*.

Koustas, D. K. (2019). What do big data tell us about why people take gig economy jobs? In *AEA Papers and Proceedings*, volume 109, pages 367–71.

Lachowska, M., Mas, A., Saggio, R., and Woodbury, S. A. (2023). Work hours mismatch. Technical report, National Bureau of Economic Research.

Lambrecht, A. and Skiera, B. (2006). Paying too much and being happy about it: Existence, causes, and consequences of tariff-choice biases. *Journal of marketing Research*, 43(2):212–223.

Lea, M. (2010). International comparison of mortgage product offerings. Technical report, Research Institute for Housing America.

Liu, L. (2022). The demand for long-term mortgage contracts and the role of collateral. Technical report, Imperial College London.

Lockwood, B. B. (2020). Optimal income taxation with present bias. *American Economic Journal: Economic Policy*, 12(4):298–327.

Luco, F. (2019). Switching costs and competition in retirement investment. *American Economic Journal: Microeconomics*, 11(2):26–54.

Mahoney, N. (2022). Principles for combining descriptive and model-based analysis in applied microeconomics research. *Journal of Economic Perspectives*, 36(3):211–22.

Mas, A. and Pallais, A. (2017). Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–59.

Mas, A. and Pallais, A. (2020). Alternative work arrangements. forthcoming. *Annual Review of Economics*.

Miles, D. (2004). The uk mortgage market: Taking a longer-term view. Technical report, Imperial College London.

Morgul, E. F. and Ozbay, K. (2015). Revisiting labor supply of new york city taxi drivers: Empirical evidence from large-scale taxi data. Technical report.

Mullainathan, S., Schwartzstein, J., and Congdon, W. J. (2011). A reduced-form approach to behavioral public finance.

Nelson, S. (2022). Private information and price regulation in the us credit card market. Mimeo, University of Chicago.

Nevo, A., Turner, J. L., and Williams, J. W. (2016). Usage-based pricing and demand for residential broadband. *Econometrica*, 84(2):411–443.

Nguyen, Q. and Leung, P. (2009). Revenue targeting in fisheries: The case of hawaii longline fishery.

Nielsen, K. and Rehbeck, J. (2022). When choices are mistakes. *American Economic Review*, 112(7):2237–68.

Parrott, J. A. and Reich, M. (2018). An earnings standard for new york city's app-based drivers. *New York: The New School: Center for New York City Affairs*.

Pires, P. (2022). How much can you make? misprediction and biased memory in gig jobs. *Unpublished*.

Piskorski, T. and Seru, A. (2018). Mortgage market design: Lessons from the great recession. *Brookings Papers on Economic Activity*, 2018(1):429–513.

Prassl, J. (2018). *Humans as a service: The promise and perils of work in the gig economy*. Oxford University Press.

Ravenelle, A. J. (2019). *Hustle and gig: Struggling and surviving in the sharing economy*. Univ of California Press.

Reck, D. and Seibold, A. (2020). The welfare economics of reference dependence. *Working Paper*.

Reichling, F. and Whalen, C. (2012). Review of estimates of the frisch elasticity of labor supply.

Robles-Garcia, C. (2022). Competition and incentives in mortgage markets: The role of brokers. Mimeo, Stanford University.

Rochet, J.-C. and Tirole, J. (2006). Two-sided markets: a progress report. *The RAND journal of economics*, 37(3):645–667.

Sallee, J. M. (2014). Rational inattention and energy efficiency. *The Journal of Law and Economics*, 57(3):781–820.

Scharfstein, D. and Sunderam, A. (2016). Market power in mortgage lending and the transmission of monetary policy. *Unpublished working paper. Harvard University*, 2.

Schmidt, M.-A. (2018). The daily labor supply response to worker-specific earnings shocks. Technical report, Working Paper.

Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics*, 50(3):665–690.

Smith, J. E. and Winkler, R. L. (2006). The optimizer's curse: Skepticism and postdecision surprise in decision analysis. *Management Science*, 52(3):311–322.

Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. *Journal of the European Economic Association*, 13(1):130–167.

Stango, V. and Zinman, J. (2016). Borrowing high versus borrowing higher: Price dispersion and shopping behavior in the us credit card market. *Review of Financial Studies*, 29(4):979–1006.

Stanton, C. T. and Thomas, C. (2021). Who benefits from online gig economy platforms? Technical report, National Bureau of Economic Research.

Stanton, R. (1995). Rational prepayment and the valuation of mortgage-backed securities. *The Review of Financial Studies*, 8(3):677–708.

Stanton, R. and Wallace, N. (1998). Mortgage choice: What's the point? *Real Estate Economics*, 26(2):173–205.

Taubinsky, D. and Rees-Jones, A. (2018). Attention variation and welfare: theory and evidence from a tax salience experiment. *The Review of Economic Studies*, 85(4):2462–2496.

Thakral, N. and Tô, L. T. (2017). Daily labor supply and adaptive reference points. *Working Paper*.

Thakral, N. and Tô, L. T. (2021). Daily labor supply and adaptive reference points. *American Economic Review*, 111(8):2417–43.

Thaler, R. H. (1988). Anomalies: The winner's curse. *Journal of economic perspectives*, 2(1):191–202.

Thaler, R. H. and Sunstein, C. R. (2009). *Nudge: Improving decisions about health, wealth, and happiness.* Penguin.

Thiel, J. (2021). Price discrimination, switching costs and welfare: Evidence from the dutch mortgage market. Mimeo, Copenaghen Business School.

Weyl, E. G. (2010). A price theory of multi-sided platforms. *American Economic Review*, 100(4):1642–72.

Wooldridge, J. M. (2015). Control function methods in applied econometrics. *Journal of Human Resources*, 50(2):420–445.

Zhang, D. (2022). Closing costs, refinancing, and inefficiencies in the mortgage market. Technical report, Harvard University.