

Essays on individual-level wage stickiness and forward guidance

Magnus Åhl



Essays on individual-level wage stickiness and forward guidance

Magnus Åhl

Academic dissertation for the Degree of Doctor of Philosophy in Economics at Stockholm University to be publicly defended on Monday 14 December 2020 at 14.00 in Nordenskiöldsalen, Geovetenskapens hus, Svante Arrhenius väg 12.

Abstract

Wage stickiness and household heterogeneity

Since long, it has been recognized that wage frictions provide realism to, and improve the properties of, macroeconomic models. Recently, another element of realism has generated a large macroeconomic literature: the addition of household heterogeneity, especially through uninsurable idiosyncratic risk on the level of worker productivity. In this chapter, I examine these two elements jointly. I incorporate wage stickiness on the household level into a standard macroeconomic model with household heterogeneity. A standard assumption when wages are sticky is that the labor demand is forcing, i.e., that households commit to supplying the amount of labor that is demanded, even if this is against their will. I show that in this setting, such an assumption is particularly unrealistic, and hence I relax it. I find that in an environment where the households cannot be forced to supply labor, productivity shocks can give rise to spells of severe underemployment -- a proxy for unemployment -- when wages are high relative to the productivity. When wages are low relative to productivity, hours worked rise, but only moderately so.

Wage stickiness and household heterogeneity in general equilibrium

In this chapter, I analyze the general-equilibrium implications of the heterogeneous-agents model with sticky wages that I develop in chapter 1. I find that the underemployment risk caused by a household-level wage stickiness has a large impact on the worker's precautionary motive to save, and hence also on the equilibrium interest rate. Moreover, the wage friction causes a high dispersion of labor supply across households, in turn leading to a low aggregate labor supply, and hence also to a low production. I show that the main findings are robust to variations in the key model parameters.

How big is the toolbox of a central banker? Managing expectations with policy-rate forecasts: Evidence from Sweden

Some central banks have decided to publish forecasts of their policy rates. Can such forecasts be used to manage market expectations of future policy rates? In this chapter, I use an event study and regression analysis on Swedish high-frequency data to conclude in the affirmative. Surprises in an announced policy-rate forecast by the central bank affect expectations of the future policy rate up to a horizon of approximately a year and a half. However, the response is not one-to-one, but is estimated to be less than one half. It is also decreasing with the forecast horizon. Moreover, I find that the actual decisions of the bank on its current policy rate -- to the extent their choices are surprises -- influence the market expectations. However, this mechanism is only active for short horizons (less than two quarters). The longer-run market expectations on the policy rate are not affected by policy-rate surprises today.

Stockholm 2020

<http://urn.kb.se/resolve?urn=urn:nbn:se:su:diva-186075>

ISBN 978-91-7911-340-7
ISBN 978-91-7911-341-4
ISSN 0346-6892



Department of Economics

Stockholm University, 106 91 Stockholm

ESSAYS ON INDIVIDUAL-LEVEL WAGE STICKINESS AND
FORWARD GUIDANCE

Magnus Åhl

Essays on individual-level wage stickiness and forward guidance

Magnus Åhl

©Magnus Åhl, Stockholm University 2020

ISBN print 978-91-7911-340-7

ISBN PDF 978-91-7911-341-4

ISSN 0346-6892

Cover picture: *Blå rök* by Mattias Bolin

Back-cover photo by Malin Åhl

Printed in Sweden by Universitetsservice US-AB, Stockholm 2020

Doctoral dissertation
Department of Economics
Stockholm University

Abstracts

Wage stickiness and household heterogeneity

Since long, it has been recognized that wage frictions provide realism to, and improve the properties of, macroeconomic models. Recently, another element of realism has generated a large macroeconomic literature: the addition of household heterogeneity, especially through uninsurable idiosyncratic risk on the level of worker productivity. In this chapter, I examine these two elements jointly. I incorporate wage stickiness on the household level into a standard macroeconomic model with household heterogeneity. A standard assumption when wages are sticky is that the labor demand is forcing, i.e., that households commit to supplying the amount of labor that is demanded, even if this is against their will. I show that in this setting, such an assumption is particularly unrealistic, and hence I relax it. I find that in an environment where the households cannot be forced to supply labor, productivity shocks can give rise to spells of severe underemployment — a proxy for unemployment — when wages are high relative to the productivity. When wages are low relative to productivity, hours worked rise, but only moderately so.

Wage stickiness and household heterogeneity in general equilibrium

In this chapter, I analyze the general-equilibrium implications of the heterogeneous-agents model with sticky wages that I develop in chapter 1. I find that the underemployment risk caused by a household-level wage stickiness has a large impact on the worker's precautionary motive to save, and hence also on the equilibrium interest rate. Moreover, the wage friction causes a high dispersion of labor supply across households, in turn leading to a low aggregate labor supply, and hence also to a low

production. I show that the main findings are robust to variations in the key model parameters.

How big is the toolbox of a central banker? Managing expectations with policy-rate forecasts: Evidence from Sweden

Some central banks have decided to publish forecasts of their policy rates. Can such forecasts be used to manage market expectations of future policy rates? In this chapter, I use an event study and regression analysis on Swedish high-frequency data to conclude in the affirmative. Surprises in an announced policy-rate forecast by the central bank affect expectations of the future policy rate up to a horizon of approximately a year and a half. However, the response is not one-to-one, but is estimated to be less than one half. It is also decreasing with the forecast horizon. Moreover, I find that the actual decisions of the bank on its current policy rate — to the extent their choices are “surprises” — influence the market expectations. However, this mechanism is only active for short horizons (less than two quarters). The longer-run market expectations on the policy rate are not affected by policy-rate surprises today.

*To Malin,
for providing the love and inspiration I needed,
and to Oskar,
for trading my time for his motivational love
at a very beneficial exchange rate*

Acknowledgements

The decision to come back to the academic world after several years of working was not easy to make, and required some time of consideration. However, I took that step, and a thesis materialized several years later. This journey would not have been possible without the support and guidance from several people, whom I devote this short text of thankfulness to. But I first want to point to an environment, rather than specific people: the IIES, with its narrow corridors and ill-painted walls (at least in office A865), has provided the best possible setting for me to produce a thesis, and make good friends, over the last years.

One important reason behind the IIES atmosphere is definitely Per Krusell, my advisor. I started to appreciate his brilliance in the first-year PhD macro course he taught, where he caught the tired students' attention by pointing out the resemblance between the name of a common fruit, and his own. From there, I have come to know him as a great and generous advisor, regarding the thesis, but also more broadly. He has a genuine interest in the PhD program in Stockholm/Uppsala, and devotes much time and energy to improving it and making sure everyone within it prosper. Recently, I have accepted that I can save time by deciding to agree with him immediately, because when I do not, I anyway always end up realizing, after some pondering, that he had a very good point in the first place. I also share his enthusiasm to exercise and follow sports. In particular, I have enjoyed playing football, tennis, and yes, occasionally even floorball (innebandy) with him. He has also, on several occasions, brought along colleagues to provide live support to his favorite football team Falkenbergs FF, which gave me much joy.¹ But what I have perhaps liked most about Per is how he, despite his obvious seniority, does not distance himself, but rather opens up and shows vulnerability. He has contributed to my life beyond research and this thesis.

¹On these occasions, the handful of us, wearing yellow shirts and scarfs provided by Per, approximately doubled the number of people publicly displaying their support for the away team. When writing this, FFF is in severe risk of relegation from the top Swedish league. The team is placed second to last, with 19 points, 6 short of "safe ground", after 24 out of 30 games played. Hence, they need all the support they can get. I grant them that support, as long as it is not at the expense of AIK, my favorite team.

I also want to acknowledge Jesper Lindé, who not only hired me for my first position at the Riksbank, but also early on accepted my request to become my second advisor for this thesis. In the end, I did not take advantage of this possibility to the extent I could have, which is entirely on me, and not on Jesper. He is an excellent researcher, and a very funny and enthusiastic person.

At the IIES, the supervision of macro PhD students is done more or less collectively. I am therefore very grateful for all the good advice I have received, as well as excellent teaching, from the macro faculty at the IIES: John Hassler, Timo Boppart, Kurt Mitman, Tobias Broer, Kathrin Schlafman, Alexandre Kohlhas, Yimei Zou, and Kieran Larkin.

One advantage of the IIES is that it produces research in a good mix of subfields within economics. Some examples of people I have enjoyed the company of are Ingvild Almås, Jon de Quidt, Mitch Downey, and Jakob Svensson. Anyone whom I have not mentioned by name here, should still feel part of contributing to the IIES atmosphere that I mentioned above.

The first two years of the PhD program, I spent in the Department of Economics. There, I want to acknowledge Johan Söderberg. I enjoy teaching a lot, and being Johan's teaching assistant several consecutive years for his second-year Master's macro course was both fun and educating. Karolina Ekholm was not in the Department of Economics during my time there, but our paths have crossed before and after that. Besides the conversations being interesting, I always feel encouraged and motivated after talking to her.

I believe that one of the most difficult jobs in an institute involved in economic research must be to handle the administration — not least dealing with all the researchers. I cannot imagine what the IIES would be without Christina Lönnblad. For sure, the conversations about tapirs would decrease significantly, and unfortunately. Moreover, the quality of the language in this thesis has benefited from Christina's help. Ulrika Gålnander, besides providing completely necessary aid in the final stages of a thesis, and in other areas, is the social glue of the IIES.² She is an expert in making friends, and has made one in me. Other people, in administrative roles, that I have enjoyed much working with over the years

²I often found her sitting in my chair chatting with my office mates, disappointed that I am back so soon, when I returned from lunch.

include Tove Happonen, Hanna Weitz, Annika Andreasson, Karl Eriksson, Anne Jenssen, and Anita Karlsson.

In the first paragraph of these acknowledgments, I mentioned that the decision to apply to the PhD program was a difficult one. Back then, when I had been working a few years at the Riksbank, it was my manager at the time, Ulf Söderström, who convinced me to apply. For this I am obviously very grateful. At the Riksbank, I also want to acknowledge the former board member Lars EO Svensson, who has always been encouraging and humble, as well as a good example of how great researchers sometimes give something back to the society in a very direct way. Over the years, staying in touch with the Riksbank, in particular my colleagues at the Modelling Division in the Monetary Policy Department, has been very rewarding. I have also received generous support and terms by my managers: Jens Iversen, and Vesna Corbo. On many occasions, I have discussed my research with my former manager David Vestin, who always has creative and insightful perspectives.

Throughout almost my entire time as a PhD student, I have shared an office with Jonna Olsson. It must be said that she has provided invaluable help with my research. But that aside, I have enjoyed her company, and learned so much from her. I sincerely admire her integrity and unobtrusive confidence. Our discussions, covering feminism, graph color schemes, grammar, politics, and much more, have truly had an impact on me in a good way.

Most of the time, Jonna and I have shared the office with someone else. Of all our former office mates, some deserve extra attention. Karin Kinnerud has helped, inspired, and encouraged me in so many ways. Given her background, I am also proud to have fought with her to the 18th hole for the unofficial IIES golf championship one year. Divya Dev, whom I have fought with many times on the squash court, has also always provide the most enjoyable company possible. Miriam Hurtado Bodell has an energy I always envied and tried to absorb, mostly without success. All of you will remain sources of inspiration to me.

Somehow, this thesis ended up consisting of only single-authored chapters. This is a bit unfortunate, since I have enjoyed several joint projects over the years. I will not dwell on the possible reasons that none of them found its way into this thesis, and I hope that they will eventually

lead somewhere. But I want to thank the people I have had the fortune to collaborate closely with: Jakob Almerud, Mattias Almgren, John Kramer, Ricardo Lima, and Andrea Papetti.

Due to a parental leave, I made a transition from one cohort of PhD students to another. I was fortunate to start out as part of my first cohort, consisting of Jonna, Serena Cocciolo, Selene Ghisolfi, Dany Kessel (whom I always viewed as belonging to our cohort, despite a head start), Matilda Kilström, Erik Lindgren, Jaakko Merilainen, Matti Mitrunen, Elisabet Olme, and Josef Sigurdsson. I believe that we peaked while singing at each others defense parties. But I have also been warmly welcomed into my latest cohort, especially by the macro students: Karin, Kasper Kragh-Sørensen, Markus Karlman, Fredrik Paues, and Has van Vlokhoven.

Somewhat orthogonal to the ordinary business, there are several sports activities taking place involving people from the IIES and elsewhere. I want to direct a special thank to everyone who have had a roll in this over the years, although I do not list them all here. These sport activities provide a welcome, and sometimes absolutely necessary, break. Besides some already covered, I want to take the opportunity to mention squash in the early years with Thomas Seiler, and the never-ending battles with Kasper on Wednesday football. I regard Kasper as a dear friend, so I hope he does not mind me mentioning his inspiring competitiveness and temper during several of these battles.

There are many other PhD students, not belonging to any of the above described categories, whom I have learned from and enjoyed spending time with: Karl Harmenberg, Richard Foltyn, Saman Darougeheh, Erik Öberg, Niels-Jakob Harbo Hansen, Hannes Malmberg, Sirius Håfström Dehdari, Benedetta Lerva, Gualtiero Azzalini, Agneta Berge, Tillmann von Carnap, José-Elías Gallegos Dago, Stefan Hinkelmann, Philipp Hochmuth, Markus Kondziella, Carolina Lindholm, Francesco Loiacono, Markus Peters, Sreyashi Sen, Fabian Sinn, Xueping Sun, and Claire Thürwächter, to mention some.

Also life outside of work has been an important source of energy during these years. I am grateful to all my friends who have supported me and beard with me when I have tried, often using far too many graphs, to explain what it is I actually do. Not the least I am grateful to Mattias

Bolin, who has created the beautiful cover picture on the front of this thesis.

My parents, Marika Åhl and Lars Jansson, although perhaps not having a direct impact on me pursuing a PhD in economics, have laid the foundation for this.

One thing that motivated me to choose this journey was a desire to be thoroughly challenged. The PhD program, beginning with two years of coursework, followed by intense thesis work, did this. I have been thoroughly challenged. One who knows this is my wonderful wife, Malin Åhl. She has stood by my side throughout this challenge, and I would never have finished this thesis without her. Since 2016 we also have Oskar in our lives, as an eternal source of inspiration. Thank you both, for everything.

Magnus Åhl
Solna, Sweden
October 2020

Contents

Introduction	i
References	v
1 Wage stickiness and household heterogeneity	1
1.1 Introduction	2
1.2 Model	7
1.3 Choice of parameter values	23
1.4 Results	25
1.5 Conclusions	48
References	50
1.A Appendix	54
2 Wage stickiness and household heterogeneity in general equilibrium	67
2.1 Introduction	68
2.2 Model	70
2.3 Choice of parameter values	88
2.4 Micro-level data (PSID 2017)	91
2.5 Results	95
2.6 Conclusions	128
References	129
2.A Appendix	132
3 How big is the toolbox of a central banker?	147
3.1 Introduction	148
3.2 Method	150
3.3 Results	161

3.4	Conclusions	179
	References	181
3.A	Appendix	184
	Sammanfattning (Swedish summary)	196

Introduction

The topics covered in this thesis are motivated by my work background in economic policy, specifically monetary policy. Working, first as a statistics assistant, then as an economist, in the Modelling Division within the Monetary Policy Department at Sveriges Riksbank has showed me the importance of conducting research that has relevance for decisions that directly affect many people. My ambition when I started writing this thesis, some time ago, was to return to policy work once I finished. During the process of working with the thesis, many thoughts around this have naturally flown through my mind. But in the end, I am back where I first intended.

The thesis consists of three self-contained chapters, of which the first two are more connected than the last. Chapter 3, which was written first, started as a spin-off from a joint project with Jakob Almerud and Andrea Papetti during a second-year course. The original project asked questions related to the *forward-guidance puzzle*, a term coined by Giannoni et al. (2015), i.e., the counter-intuitive result that a *credible* communication of a policy-rate adjustment far into the future has a much larger impact on the economy today, than the same immediate adjustment of the policy rate has. We figured that the key to solving the puzzle was in the word *credible*. A central bank communicating intentions far into the future could never be fully credible in doing so, since everyone is aware that the future is highly uncertain. Unanticipated events are bound to occur, creating incentives for the central bank to deviate from its communicated plan. I decided to take an empirical approach to this, exploiting that the Riksbank systematically communicates its intended future actions with the policy rate. How credible, or incredible, is this communication to the market? And could the conclusions resolve the forward guidance puzzle?

Chapters 1 and 2 are tightly connected and make use of the same

macroeconomic model. They started as one single project motivated by an insufficient understanding of the transmission mechanisms in more standard macroeconomic models. The core microfounded model used at the Riksbank has in common with most models used at central banks that it assumes all households to be alike. The homogenous households can thus be seen as one representative household. This has technical advantages, but modern macroeconomic research suggests that important aspects of the transmission mechanisms might be missing. In the last years, a large literature with the ambition to overcome these problems by introducing household heterogeneity in standard models with price frictions has emerged. Within this literature, I noticed a gap in the modeling of wage stickiness, a feature that is both realistic and important for correctly capturing the transmission mechanisms. The purpose of these two chapters is to help fill this gap in the literature. I now introduce the chapters, one at a time, in the order they appear in this thesis.

In the first chapter, **Wage stickiness and household heterogeneity**, I consider a standard model with household heterogeneity, that cannot be neutralized with insurance, creating a precautionary motive to save. I assume that idiosyncratic productivity is hit by stochastic shocks, and impose a household-level wage stickiness.³ The type of wage friction I assume is widely used in the new-Keynesian literature, at the aggregate level, but not at the individual level. The main question I address is if such wage stickiness can yield microeconomic mechanisms that are perceived as realistic and intuitive. How exactly should the wage friction be modeled to achieve this? What drives the households' decisions in this environment, and what is the effect on labor-market outcomes? The analysis is carried out in partial equilibrium, where the aggregate wage and the interest rate are exogenous.

Imposing a household-level wage friction produces two conceptual challenges. First, idiosyncratic productivity shocks implies that the amount of output that a household produces in a given period of time varies between households, and over time also varies within each household. In consequence, there are two different wages to consider: the compensation per time worked (*hourly wage*), and the compensation per output pro-

³I consider a household as consisting of one single individual, rather than several. Thus, I interchangeably use *individual* and *household* when referring to the same unit. I arbitrarily also choose a feminine pronoun, so *she* may refer to a household.

duced (*effective wage*). Either of these can be subject to the wage friction, with very different implications. Consider a situation where a household's wage was set in the past, and now her productivity drops, but she cannot update her wage. She produces less output in every hour worked, and if the hourly wage is fixed, the cost per hour for her employer remains the same. Trivially, her labor input is less valued by her employer, and the demand for it drops, which causes her underemployment. If, on the contrary, the effective wage is fixed, she indeed produces less per hour worked, but the cost she induces for her employer is also smaller. The effect on both the labor demand and her willingness to supply labor is ambiguous in this case. A main result is that it matters a lot which of these two wage concepts I assume to be sticky.

The other conceptual challenge refers to dramatic overemployment. Consider a scenario where a household's idiosyncratic productivity increases, and the wage cannot be adjusted. She now produces more output in every hour worked, and if the hourly wage is sticky, the cost per hour for her employer remains unchanged. Hence, her value to her employer increases, as does the demand for her labor services. If the labor demand is forcing to households, i.e., households commit to obeying the labor demand, which is the common assumption in the literature, she becomes overemployed. I.e., she supplies more labor than she would if she could update her wage. Idiosyncratic productivity shocks tend to be much larger than aggregate productivity shocks, so the overemployment may be very dramatic. This might be conceived as problematic. It can, however, be remedied by allowing households to supply less labor than is demanded from them. I refer to this as a *non-forcing* labor demand. It creates an asymmetry that allows a severe underemployment due to "too high" a wage, but prevents a severe overemployment due to "too low" a wage.

The two considerations described above — sticky hourly or sticky effective wages, and forcing or non-forcing labor demand — give rise to four potential versions of the model. I conclude that only one of these versions yields sound microeconomic mechanisms and plausible correlations at the individual level: a sticky hourly wage, and a non-forcing labor demand. Only in that version, we see a negative correlation between the idiosyncratic productivity and underemployment, in combination with an unbroken positive relationship between the productivity and individual

welfare. A decreasing idiosyncratic productivity when the wage is stuck creates a wage-productivity mismatch. The employer is better off by substituting to other workers, and the employee thus faces a lower labor demand. As a consequence of the falling demand, the worker suffers a spell of involuntary underemployment. The spell ends by either an opportunity to reset a lower wage, or a positive idiosyncratic productivity shock — both making the worker more valuable to the employer.

In the opposite situation, when a worker becomes more productive, and the hourly wage is stuck, the wage-productivity mismatch makes her more valuable to her employer. Labor demand increases, but she is not forced to meet the demand. Typically, her labor supply remains quite flat. However, were she to get a chance to update her wage, she would set it higher to erase the wage-productivity mismatch, which increases her earnings.

In the second chapter, **Wage stickiness and household heterogeneity in general equilibrium**, I close the same model of household-level wage stickiness by explicitly modeling the firm side of the economy, and endogenizing the aggregate wage and the interest rate to clear the markets for labor and government bonds. The resulting general equilibrium allows me to study how the aggregate outcomes are affected by micro-level wage stickiness and thus the mechanisms discussed in chapter 1. How is the precautionary motive to save affected by an underemployment risk, and what effect does that have on the equilibrium interest rate? Does rationing in the labor market, caused by idiosyncratic wage-productivity mismatches, affect aggregate production?

More precisely, I take the standard HANK model and add two things: a friction to the individual-household wages, and a relaxation of the assumption that the demand is forcing in the labor market.⁴ In addition to answering questions about aggregate outcomes, this framework is suitable for addressing distributional questions about, e.g., underemployment and wealth.

My main findings are that this model gives rise to micro distributions that qualitatively match the US data. The risk of underemployment, which is economically significant, amplifies the precautionary motive to

⁴HANK — Heterogeneous Agents New-Keynesian — was coined by Kaplan et al. (2018).

save. Compared to the model with flexible wages, this results in a lower equilibrium interest rate. The wage friction also distorts labor supply across households, resulting in a lower aggregate labor, and hence a lower output. Another consequence of labor supply being heavily dispersed across households, mainly caused by underemployment, is a larger earnings inequality. However, the precautionary savings behavior actually results in a more concentrated distribution of wealth. In terms of mechanisms, the underemployed tend to be wealth-poor and low-productive, in line with micro data on unemployment. I also show that the findings are robust to variations in the key model parameters.

In the third chapter, **How big is the toolbox of a central banker?**, I perform an event study of the impact on Swedish financial market expectations of surprises in the Riksbank's announcements following policy-rate decisions. An announcement consists of at least three parts: a policy-rate decision, a policy-rate forecast for the future, and a report analyzing the current economic situation. I distinguish the effect of surprises in the policy-rate forecast on expectations of the future policy rate. The identification relies in part on high-frequency changes of financial data around announcements, and I use movements in forward rates on different horizons to measure the impact on expectations.

My main finding is that a surprise in the policy-rate forecast published by the Riksbank does move market expectations of the future policy rate. However, the effect is not one-to-one, and only significant up to around a year and a half, which is shorter than the forecast horizon. The results are robust to a number of reasonable variations in my measures and variables.

References

- Giannoni, M., Patterson, C., and Negro, M. D. (2015). The Forward Guidance Puzzle. Technical report.
- Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary Policy According to HANK. *American Economic Review*, 108(3):697–743.

Chapter 1

Wage stickiness and household heterogeneity*

*I would like to thank Susanto Basu, Timo Boppart, Tobias Broer, Richard Foltyn, John Hassler, Karin Kinnerud, Per Krusell, Kurt Mitman, Jonna Olsson, Christina Patterson, Kathrin Schlafmann, David Vestin, and participants in the IIES Macro Group for helpful discussions.

1.1 Introduction

Macroeconomic policy, and stabilization policy in particular, affects the welfare of many millions of people. For guidance, policymakers have relied on macroeconomic theory, but the nature of this theory has changed significantly over at least the last century. From a theory essentially described in terms of macroeconomic aggregates, the developments in the 1980's, with the introduction of the real business-cycle theory (RBC), refocused on building theory from the bottom and up. I.e., the idea was to describe individual households' optimization problem, and aggregate their optimal behavior to macroeconomic variables. The next step in the historical development was to include price frictions, linking the nominal and real sides of the economy in a realistic way, as captured by the new-Keynesian (NK) dynamic stochastic general-equilibrium model (DSGE) — which is still the workhorse model in most central banks in developed countries. These models often include wage stickiness, since it improves the properties of the models. However, markets are often assumed to be complete, and hence these models lack meaningful heterogeneity on the household level. In response, models of incomplete markets, where the distribution of households matters for aggregate outcomes, were developed.¹ Recently, this literature has merged with the literature on sticky prices, reflected by a large and growing number of heterogeneous-agent new-Keynesian models (HANK). Some of these models include imperfectly substitutable workers, and wage stickiness at the aggregate level, but taking household heterogeneity seriously requires decisions about the wage to be made at the individual level, which is also confirmed by the data. I contribute to this literature by introducing wage stickiness on the household level, and analyzing how it affects the microeconomic mechanisms governing the labor market. The ultimate aim is to better capture the behavior of the millions of people affected by macroeconomic policy, and thereby improve the policy itself.

More precisely, I consider a Bewley-Huggett-Aiyagari model with a precautionary motive to save, and idiosyncratic shocks to household productivity. I impose household-level wage stickiness of the Calvo (1983)

¹Bewley (1986), Imrohoroglu (1989), Huggett (1993), and Aiyagari (1994) are famous early contributions. Krusell and Smith (1998) introduce aggregate uncertainty in this class of models.

type, i.e., opportunities to update the wage appear randomly, and are hence state independent. This type of wage friction is widely used in the new-Keynesian literature, at the aggregate level, since it was introduced by Erceg et al. (2000). This formulation is consistent with the practice of wage setting, where wages are fixed for long periods of time and regularly adjusted significantly.² Can such household-level wage stickiness yield microeconomic mechanisms that are perceived as realistic and intuitive, and if so, how should it be modeled? What drives the households' decisions in this environment, and what is the effect on labor-market outcomes? These are the questions I address in this chapter.³ The analysis is carried out in partial equilibrium. In a separate chapter, Åhl (2020), I take the same model to general equilibrium.

The approach I choose produces two conceptual challenges. First, stochastic idiosyncratic productivity implies that the amount of output a household produces in a given period of time varies between households, but also over time within each household. On top of this, empirical research shows that idiosyncratic productivity tends to vary an order of magnitude more than aggregate productivity, thus making the variation very large.⁴ In consequence, there are two different wages to consider: the compensation per unit of time worked, and the compensation per output produced. I refer to the former as *hourly wage*, and the latter as *effective wage*. Either of these can be subject to the wage friction, with very different implications. Consider a situation where a household's wage was set in the past, and now her productivity drops, but she cannot update her wage.⁵ She produces less output every hour worked, and if the hourly wage is fixed, the cost per hour for her employer is still the same. Trivially, her labor input is less valued by her employer, and the demand

²This is in contrast with smooth and recurrent wage setting à la Rotemberg (1982).

³Ideally, there would be household-level panel data on wages, productivity, and hours worked to which the model could be compared, in pursuit of answers to these questions. Such data on wages and hours worked, of course, exists, but not for idiosyncratic productivity. In empirical work, idiosyncratic productivity is often assumed to move one-for-one with the wage. One of the main points of this paper is to disentangle the wage from productivity at the micro level, to see how the labor market is impacted.

⁴To be clear, I assume aggregate productivity to be fixed, so all the variation stems from idiosyncratic movements.

⁵I consider a household as consisting of one single individual, rather than several. Thus, I interchangeably use *individual* and *household* when referring to the same unit. I arbitrarily also choose a feminine pronoun, so *she* may refer to a household.

for it drops, thereby causing her underemployment. If, on the contrary, the effective wage is fixed, she will indeed produce less per hour, but the cost she induces her employer is also smaller. The effect on both labor demand, and her willingness to supply labor, is ambiguous in this case. Which of these two possible assumptions, sticky hourly or effective wage, is more reasonable? I will argue that it is the sticky hourly wage.

The other conceptual challenge refers to dramatic overemployment. Consider the opposite situation as above: a household's idiosyncratic productivity increases, and the wage cannot be adjusted. She now produces more output in every hour worked, and if the hourly wage is sticky, the cost per hour for her employer is unchanged. She is the perfect employee — effective and cheap — and the demand for her labor services increases. If the household commits to obeying the labor demand, which is the common assumption in the literature, and which I refer to as a *forcing* labor demand, she becomes overemployed. I.e., she supplies more labor than she would if she could update her wage. Since idiosyncratic productivity shocks tend to be much larger than aggregate productivity shocks, the overemployment may be very dramatic, so she may end up working many times more than she would like. This might be conceived as problematic, and has been criticized before, but mainly in models without idiosyncratic shocks.⁶ It can, however, be remedied by relaxing the assumption that a household commits to obeying the labor demand, and allowing her to supply less (but never more) labor than what is demanded from her. I refer to this as a *non-forcing* labor demand. It creates an asymmetry that allows a severe underemployment due to “too high” a wage, but prevents severe overemployment due to “too low” a wage. I will argue that a non-forcing labor demand is the better choice.

The two considerations described above — sticky hourly or effective wages, and forcing or non-forcing labor demand — give rise to four potential versions of the model. I conclude that only one of these versions yields sound microeconomic mechanisms, and plausible correlations at

⁶Huo and Ríos-Rull (2020) show that overemployment, or violation of the labor participation constraint, is problematic in a model with aggregate shocks only. They also show that the problem is nonlinear, appearing only when shocks are large. Given that idiosyncratic productivity shocks tend to be large compared to aggregate shocks, this hints at how large the problem is in models with idiosyncratic productivity, like the one I study here.

the individual level: sticky hourly wages and a non-forcing labor demand. Only in this version do we see a negative correlation between idiosyncratic productivity and underemployment, in combination with an unbroken positive relationship between productivity and individual welfare. Decreasing idiosyncratic productivity when the wage is stuck creates a wage-productivity mismatch. The employer is better off by substituting to other workers, and the employee faces a lower labor demand. The elasticity of substitution between skill types is crucial in determining how much the demand drops. As a consequence of the falling demand, the worker suffers a spell of involuntary underemployment. The spell ends by either an opportunity to reset a lower wage, or a positive idiosyncratic productivity shock — both making the worker more valuable to the employer.

In the opposite situation, when a worker becomes more productive with the hourly wage being stuck, the wage-productivity mismatch makes her more valuable to her employer. Labor demand increases, but she is not forced to meet the demand. Typically, her labor supply remains quite flat, creating a gap between labor demand and supply, at the individual-worker level. However, were she to get a chance to update her wage, she would set it higher to increase her earnings. Both the wage-productivity mismatch and the demand-supply gap would disappear or shrink substantially.

Assuming the effective wage to be sticky is inferior because it makes the underemployment stem from an increased desire to supply labor in response to a positive productivity shock, rather than suppressed labor demand in response to a negative shock. Aside from being unintuitive in itself, a consequence is that the underemployed are high-productive households, which is not in line with the intuition. A forcing labor demand, on the other hand, is unrealistic because, with idiosyncratic productivity and a sticky hourly wage, I show that a household is forced to supply ridiculous amounts of labor in response to a positive productivity shock if the wage cannot be adjusted immediately.

I hope that these findings prove helpful in the future pursuit of a better understanding of how microeconomic decisions transmit to macroeconomic outcomes. The structure of the rest of the paper is the following: Immediately below are brief comments on the related literature, and section 1.2 introduces the model in detail, and distinguishes the four versions. Section 1.3 briefly discusses the parameter values I choose, section 1.4

presents and discusses some simulation results of the different versions, and finally section 1.5 concludes the chapter.

Related literature This paper is related to the literature on macroeconomic models with sticky wages, analyzing how the wage stickiness affects and improves the underlying mechanisms of the model. Early contributions include, but are not limited to Erceg et al. (2000), Christiano et al. (2005), and Schmitt-Grohé and Uribe (2005). The most obvious deviation from those papers is that they assume complete markets, or a representative agent.

I also contribute to the large and growing literature on macroeconomic models with incomplete markets and heterogeneous agents, stemming from early work by Bewley (1986), Imrohoroglu (1989), Huggett (1993), and Aiyagari (1994), and where Krusell and Smith (1998) introduced aggregate uncertainty. Although my paper does not explicitly include nominal frictions, specifically not frictions on goods prices, it closely relates to the recent HANK literature, e.g., Kaplan et al. (2018), McKay et al. (2016), Gornemann et al. (2016), Auclert (2017), Kaplan and Violante (2018), and Luetticke (2020).⁷ I mainly deviate in two ways: I apply wage stickiness, and my analysis is in partial equilibrium, while most of the mentioned papers have a general-equilibrium analysis.

But also the HANK literature contains examples of wage stickiness, of which Hagedorn et al. (2019a), Hagedorn et al. (2019b), Broer et al. (2019), Auclert et al. (2020), and Bayer et al. (2020) are some notable examples. However, all of these examples model the wage stickiness at an aggregate level and/or use smooth stickiness à la Rotemberg (1982). Moreover, they do not systematically evaluate the implications of the sticky wage for individual households' labor supply, which is the focus of this paper.

One consequence of individual-level wage stickiness; in combination with idiosyncratic productivity shocks, and monopolistic competition in the labor market; is that the distinction between forcing and non-forcing labor demand becomes crucial. Although exactly this setting does not, to my best knowledge, exist in the literature, the deficiency of forcing labor

⁷I choose to cite a working-paper version of Auclert (2019), because the published version does not contain the HANK model section.

demand is discussed in a representative-agent setting in Huo and Ríos-Rull (2020).⁸ They conclude that forcing labor demand is problematic even without idiosyncratic fluctuations in productivity. I show that this problem becomes much more severe with idiosyncratic productivity shocks.

Finally, the solution method I employ to solve the households' problem relates to the literature on generalizations of the endogenous gridpoint method introduced by Carroll (2006). Barillas and Fernández-Villaverde (2007) is an early example, which is not applicable because it only handles multiple control variables which are not endogenous state variables. White (2015), Druedahl and Jørgensen (2017), and Ludwig and Schön (2018) are all more recent examples, allowing for multiple endogenous state variables. However, none of their methods are directly applicable in my case, since they require a closed-form solution for the first-order conditions. That is not the case in this paper, where the optimal wage choice lacks a closed-form solution.

1.2 Model

This paper presents a model of idiosyncratic shocks to productivity, and a household-level wage friction. The wage friction creates a potential gap between the labor demand and the desired labor supply, at the individual level. The idiosyncratic productivity shocks make the gap occasionally large. Altogether, this makes the assumption of whether labor demand is forcing — as in the seminal paper by Erceg et al. (2000), and many papers since — or not very important.

Moreover, with heterogeneous productivity, i.e., how much output a worker produces in a certain time input, there are two different wages to consider; the compensation per time input, and the compensation per effective output. For simplicity, I henceforth refer to these as the *hourly wage* and the *effective wage*, respectively. Which of these is assumed to be subject to the wage friction has a substantial impact on the mechanism at the micro level in the economy. These two crucial binary assumptions give rise to four versions of the model, illustrated in table 1.1.

To aid the understanding of the difference between an hourly and an

⁸See the references therein for an overview of the scarce literature discussing this topic.

Table 1.1: Model versions

	Forcing labor demand	
	<i>Yes</i>	<i>No</i>
Wage friction		
<i>Hourly</i>	Version 1	Version 2
<i>Effective</i>	Version 4	Version 3

effective wage, I provide a very hands-on example. Imagine two workers, A and B , whose job is to pack boxes. Both pack exactly 10 boxes in every hour worked, and are hence equally productive. A has an *hourly* wage of x , and B has an *effective* wage of x per every 10 boxes. Presently, they are compensated alike, for identical work. Both their wages are sticky. Now imagine that they are both struck by identical positive productivity shocks, so they start packing 12 boxes in every hour worked. However, they are not able to update their wages. A is still paid x per hour, but her compensation per box has gone down from $x/10$ to $x/12$, i.e., her effective wage has decreased. B , however, is still compensated by $x/10$ per box she packs, and hence her hourly wage has now increased from x to $1.2x$. It is intuitive that the shock made A more valuable to the firm they work for. Whether B is more valuable to the firm or not is unclear, but certainly her value relative to A has decreased, since they perform the exact same work at different costs to the employer. A describes the situation in versions 1 and 2 of the model, while B belongs to versions 3 and 4.

The focus of this paper is on household behavior, and common to all versions is that households face a dynamic consumption-savings choice. Wages are subject to a Calvo (1983)-type friction, meaning that with some probability each period, a household also gets to choose her wage, and can hence influence the demand for her labor services. A household that is

unable to update her wage is stuck with the wage, hourly or effective, she had in the last period. I begin by presenting the firm sector, since that is where the labor demand is decided. Then I continue with a detailed discussion of the households' problem in the different model versions.

1.2.1 The representative firm

The main interest of this paper is household behavior, where labor demand matters substantially. However, since the demand comes from firms, some comments on that sector are necessary. There is a representative firm with the static objective of minimizing the labor costs subject to a target level for effective labor input.⁹

Households are heterogeneous in skill types, which are symmetric and substitutable with a constant elasticity of δ . More precisely, the production technology to aggregate labor skill types is, as in Dixit and Stiglitz (1977):

$$L = \left(\int_0^1 [e(i)g^l(i)]^{\frac{\delta-1}{\delta}} di \right)^{\frac{\delta}{\delta-1}}, \quad (1.1)$$

where households are indexed by $i \in [0, 1]$, $e(i)$ is the idiosyncratic productivity of household i , and $g^l(i)$ is the endogenous labor supply, in time units, from household i . The product $e(i)g^l(i)$ is hence the supply of effective labor by household i , which is what enters the firm's aggregation technology (1.1). The imperfectly substitutable skill types give the households monopolistic power in the labor market

As mentioned above, labor demand can be forcing or not to households. This makes a difference to the firm since if demand is forcing, all households are responsive to shifts in labor demand. However, if labor demand is not forcing, there is a subset of households that are unresponsive to a marginal shift in labor demand. The labor supply from unresponsive households is hence exogenous to the firm, and cannot be part of the cost-minimizing labor-demand decision. The two cases are handled in detail below.

⁹One can think of this labor being utilized to produce the consumption good. However, since this paper does not aim at a general-equilibrium analysis, that step is not necessary.

Forcing labor demand: versions 1 and 4

With forcing labor demand, all households commit to supplying the amount of labor demanded by the firm. This is the case in the seminal work by Erceg et al. (2000), introducing wage stickiness in the new-Keynesian model, and it has been standard since then; see, e.g., Christiano et al. (2005), Smets and Wouters (2007), and Hagedorn et al. (2019a).

The problem is to minimize wage-bill costs, subject to an exogenous target for aggregate effective labor, N . Formally:

$$\min_{n(i)} \left\{ \int_0^1 g^w(i) n(i) di \right\} \quad \text{subject to} \quad \left(\int_0^1 [e(i) n(i)]^{\frac{\delta-1}{\delta}} di \right)^{\frac{\delta}{\delta-1}} \geq N \quad (M.2)$$

where $n(i)$ is labor demanded from household i , and $g^w(i)$ is the hourly wage of household i . The resulting labor-demand scheme is given by:¹⁰

$$n(i) = \left(\frac{g^w(i)}{W} \right)^{-\delta} N e(i)^{\delta-1}, \quad (1.3)$$

where

$$W \equiv \frac{1}{N} \int_0^1 g^w(i) n(i) di \quad (1.4)$$

is an aggregate wage index in the economy.¹¹

Assuming that $\delta > 1$, which is well in line with the literature, we can keep the following inequality in mind:

$$\frac{\partial n(i)}{\partial e(i)} = (\delta - 1) \left(\frac{g^w(i)}{W} \right)^{-\delta} N e(i)^{\delta-2} > 0. \quad (1.6)$$

This means that with a fixed hourly wage, labor demand is increasing

¹⁰See appendix 1.A.1 for derivation details.

¹¹In versions 3 and 4, $g^w(\cdot)$ refers to the *effective*, not the *hourly*, wage. This changes the labor-demand scheme (1.3) to

$$n(i) = \left(\frac{g^w(i)}{W} \right)^{-\delta} N e(i)^{-1}, \quad (1.5)$$

see equation (1.11) below, linking the hourly and the effective wages.

with idiosyncratic productivity, which is intuitive.¹²

Non-forcing labor demand: versions 2 and 3

Under the assumption of non-forcing labor demand, a household is free to supply any amount of labor less than or equal to the demand she faces. This is not standard in the literature, but arguably more realistic; see, e.g., Huo and Ríos-Rull (2020).

This assumption complicates the firm's problem in the sense that some of the households are unresponsive to a marginal shift in the demand for their labor services (i.e., they are unconstrained by labor demand), while others follow shifts in labor demand. This must be accounted for when the firm solves its problem. Formally, let $\mathbf{1}_d(i) \in \{0, 1\}$ be an indicator of whether labor demand is binding or not for household i . The firm's problem is:

$$\min_{\mathbf{1}_d(i)n(i)} \left\{ \int_0^1 g^w(i)n(i)di \right\} \quad \text{subject to} \quad \left(\int_0^1 [e(i)n(i)]^{\frac{\delta-1}{\delta}} di \right)^{\frac{\delta}{\delta-1}} \geq N.$$

The constraint can be rewritten as

¹²The counterpart with sticky effective wage, versions 3 and 4, is

$$\frac{\partial n(i)}{\partial e(i)} = - \left(\frac{g^w(i)}{W} \right)^{-\delta} N e(i)^{-2} < 0.$$

Note the opposite sign compared to the case of a sticky hourly wage, meaning that with a fixed *effective* wage, labor demand is *decreasing* with idiosyncratic productivity. In this case, the intuition is not as clear. One way to think about this result is that demand for effective labor, i.e., $e(i)n(i)$, depends on the effective wage and aggregate outcomes. So with an unchanged effective wage, increased idiosyncratic productivity means that it takes fewer hours worked to meet the demand for effective labor.

$$N \leq \left(\int_0^1 [e(i)n(i)]^{\frac{\delta-1}{\delta}} di \right)^{\frac{\delta}{\delta-1}} \iff$$

$$N \leq \left(\int_0^1 [\mathbf{1}_d(i)e(i)n(i) + (1 - \mathbf{1}_d(i))e(i)n(i)]^{\frac{\delta-1}{\delta}} di \right)^{\frac{\delta}{\delta-1}} \stackrel{(*)}{\iff}$$

$$N \leq \left(\int_0^1 [\mathbf{1}_d(i)e(i)n(i)]^{\frac{\delta-1}{\delta}} + [(1 - \mathbf{1}_d(i))e(i)n(i)]^{\frac{\delta-1}{\delta}} di \right)^{\frac{\delta}{\delta-1}} \iff$$

$$N \leq \left(\int_0^1 [\mathbf{1}_d(i)e(i)n(i)]^{\frac{\delta-1}{\delta}} di + \underbrace{\int_0^1 [(1 - \mathbf{1}_d(i))e(i)n(i)]^{\frac{\delta-1}{\delta}} di}_{\equiv n_s} \right)^{\frac{\delta}{\delta-1}} \iff$$

$$\underbrace{\left(N^{\frac{\delta-1}{\delta}} - n_s \right)^{\frac{\delta}{\delta-1}}}_{\equiv N_d} \leq \left(\int_0^1 [\mathbf{1}_d(i)e(i)n(i)]^{\frac{\delta-1}{\delta}} di \right)^{\frac{\delta}{\delta-1}}.$$

Step (*) is possible because of the boolean nature of $\mathbf{1}_d(i)$. n_s is an auxiliary variable connected to the labor supplied by all unresponsive households (sub-index s for *supply-driven* labor), and N_d is an auxiliary variable representing aggregate labor demanded only from households constrained by labor demand; hence sub-index d for *demand-driven*. Importantly, N_d is exogenous to the firm. The above shows that the firm problem with non-forcing labor demand is isomorphic to the firm problem with forcing labor demand, as long as $\int_0^1 \mathbf{1}_d(i)di > 0$, i.e., some households are constrained by labor demand. The corresponding labor-demand scheme is given by:

$$n^*(i) = \left(\frac{g^w(i)}{W_d} \right)^{-\delta} N_d e(i)^{\delta-1}, \quad (1.7)$$

where

$$W_d \equiv \frac{1}{N_d} \int_0^1 \mathbf{1}_d(i) g^w(i) n(i) di$$

is an auxiliary wage index for responsive households only. I use the asterisk (*) in the notation to distinguish the two cases of forcing and non-forcing labor demand, and to show that in this case, labor demand is notional in the sense that it does not necessarily bind, and thus should not be confused with the actual labor employed.^{13,14}

It may appear as if the wage fails to clear the labor market when labor demand is not forcing. As long as some households are responsive to labor demand, this is not the case. However, since I consider partial equilibrium here, I refer to a more detailed discussion in Åhl (2020).

To summarize the representative firm in the model, it demands labor services from the households, as a function of their idiosyncratic productivity and wage. The labor demand may be forcing or not, depending on which version is considered, and it differs between the two cases of forcing and non-forcing labor demand. If the hourly wage is sticky, labor demand is upward-sloping in productivity as long as the wage is stuck. However, if the effective wage is sticky, this relation is inverted, and labor demand is downward-sloping in productivity as long as the wage is stuck.

1.2.2 Households

The focus of this paper is to investigate the mechanisms driving household behavior, in particular the labor-market outcomes, when wages are sticky at the household level. The focus is also on how household behavior

¹³In versions 3 and 4, g^w refers to the *effective*, not the *hourly*, wage. This changes the labor-demand scheme (1.7) to

$$n^*(i) = \left(\frac{g^w(i)}{W_d} \right)^{-\delta} N_d e(i)^{-1}, \quad (1.8)$$

see equation (1.11) below.

¹⁴The counterpart of equation (1.6) with non-forcing labor demand is:

$$\frac{\partial n^*(i)}{\partial e(i)} = (\delta - 1) \left(\frac{g^w(i)}{W_d} \right)^{-\delta} N_d e(i)^{\delta-2} > 0. \quad (1.9)$$

I.e., labor demand is increasing in idiosyncratic productivity, if the hourly wage is fixed.

is affected by some crucial assumptions introduced above. This section first presents elements that are independent of these assumptions — idiosyncratic productivity, market structure, and the wage friction — and then looks into the households’ problem for each version to which the assumptions give rise.

Idiosyncratic productivity

As in many other papers investigating heterogeneity in household behavior with incomplete markets, I assume that a household’s productivity is subject to idiosyncratic shocks. This is interpreted as a shock to how much output (as valued by the firm) the household produces in a given time unit of work. In this model, labor is supplied on the intensive margin, and it is important to bear in mind that it is the *time* spent supplying labor that renders disutility. However, the disutility is not affected by the productivity component.

In the literature, it is common to distinguish between permanent and transitory shocks to productivity.¹⁵ For simplicity, however, I follow McKay et al. (2016) and let (the logarithm of) the idiosyncratic productivity follow an auto-regressive process of order one, $AR(1)$. More specifically:

$$\log(e_t) = \rho_e \log(e_{t-1}) + \varepsilon_t^e, \quad \varepsilon_t^e \sim N(0, \sigma_e), \quad (1.10)$$

where e_t is the idiosyncratic productivity of a household at time t , ρ_e is the persistence of the process, and ε_t^e is an idiosyncratic innovation to the process at time t , which is normally distributed with mean 0 and standard deviation σ_e . Note that although the innovations to the process are idiosyncratic, the two parameters of the process, ρ_e and σ_e , are common to all households. In that sense, all households are ex-ante identical.

Equation (1.10) suggests a continuous support for e . In the numerical exercises, however, the stochastic process for e has discrete support and follows a Markov chain, using the method of Rouwenhorst (1995).¹⁶

¹⁵E.g., section 4.2.2 in Krueger et al. (2016).

¹⁶An alternative method, also common in the literature, is Tauchen (1986). Choosing that method instead does not change the overall results.

Market structure

Along with a large literature, I assume that markets are incomplete. To be more precise, there is no market for Arrow and Debreu (1954)-type securities, and thus, assuming a *representative household* would yield meaningful differences for the aggregate variables. Hence, individual household behavior is crucial in the analysis.

I also assume that there is an exogenous limit to borrowing. As a result, households save in government bonds to smooth consumption and to insure against being borrowing constrained in the future, as a consequence of bad idiosyncratic shocks. A household that is borrowing constrained cannot smooth consumption. The latter motive to save is often referred to as “precautionary savings”.

Wage friction

Households choose their own wage, subject to a Calvo (1983)-type friction. This means that in each period, there is a probability $1 - \theta_w$ that a household may update her wage, and consequently a probability θ_w that she cannot update her wage. In the latter case, she keeps the same wage as in the previous period, which I label w_{-1} , and sometimes refer to as the *beginning-of-period wage*. Note that the opportunities to update the wage occur stochastically and are independent of the past, e.g., the time since the last wage update.¹⁷

It is important to realize that with idiosyncratic productivity, there are two different wages to be considered: the compensation per time unit of labor supplied, the *hourly wage*, denoted by w^h ; and the compensation per effective labor output, the *effective wage*, denoted by w^e . Labor income is given by w^hl , or equivalently w^eel , where l is “hours worked”, and el is

¹⁷The literature offers alternative ways to model wage frictions, of which the updating cost in Rotemberg (1982) is perhaps the most common besides Calvo (1983). There are two main reasons that I choose Calvo (1983) over Rotemberg (1982). First, if one views the underemployment caused by a mismatching wage and productivity as a proxy for unemployment, the idea of paying to end an unemployment spell is not very attractive. Second, one technical advantage to choosing Rotemberg (1982) in a setting with complete markets is that it does not give rise to dispersion between households, as does the Calvo (1983) assumption. However, this advantage is no longer valid with incomplete markets and idiosyncratic shocks, where the households are already dispersed for other reasons.

effective labor output as it enters the labor aggregator (1.1).¹⁸ From this we conclude that the hourly and effective wages relate as

$$w^h = w^e e. \quad (1.11)$$

Either of the two wages is assumed to be subject to the updating friction, with quite different results, as will be clear in section 1.4.

The household's problem

A household is dynastic and lives forever with certainty. She receives utility from consumption, and disutility from the time spent supplying labor. Her income consists of earnings, dividends from (non-traded) firm ownership, and interest payments from (traded) bond holdings. She also either pays taxes to, or receives transfers from, the government. She decides how much to spend on consumption in each period, and how much to save in bonds maturing in the next period. In random periods, she may also reset her wage, which affects the labor demand she faces for as long as it is valid. If labor demand is assumed to be non-forcing, she also chooses her labor supply, capped by demand.

Although they are quite similar, I choose to state the households' problem separately for the four versions. The small differences are key to the insights of this paper. The problem is recursive, and thus I do not use a sub-index to denote the time period. Instead, a prime ($'$) indicates the next period, while a sub-index of minus one ($_{-1}$) indicates the previous period (equivalent to beginning-of-period).

Version 1: sticky hourly wage and a forcing labor demand A household solves

$$V(b, w_{-1}^h, e; f) = \max_{c, w^h, l, b'} \left\{ U(c, l) + \beta E \left[V(b', w^h, e'; f') \right] \right\} \quad (1.12)$$

subject to a budget constraint

$$c + \frac{b'}{R} = w^h l + b + D - T,$$

¹⁸To be more precise, l is the time a household spends supplying labor, not necessarily measured in the unit of hours.

a forcing labor-demand function

$$l = n(w^h, e), \quad (1.13)$$

and the wage-updating friction

$$w^h = w_{-1}^h \quad \text{if } f = 0.$$

$V(\cdot)$ is the value function (defined as the infinite sum of discounted expected future stream of instant utility), $U(\cdot)$ is the instant-utility function, b is real bond holdings, w^h is the real hourly wage, e is idiosyncratic productivity, $f \in \{0, 1\}$ is an indicator for wage updating (f for the Calvo “fairy”), c is consumption, l is labor supply (in time units), R is the gross real interest rate, D is real dividend income from firm profits, and T is real taxes to (or transfers from) the government. The parameter β is the time-discount factor. In general, lower-case Latin letters refer to household-level variables, while upper-case letters refer to aggregate variables (or functions, which should then be clear). (b, w_{-1}^h, e) is the state vector for the household, of which e is exogenous. $n(w^h, e)$ in equation (1.13) is given by the labor-demand function (1.3) from the firm’s problem.

The solution to the households’ problem can be summarized in the following decision rules: $c = g^c(b, w_{-1}^h, e; f)$ for consumption, $w^h = g^w(b, w_{-1}^h, e; f)$ for the (hourly) wage, $l = g^l(b, w_{-1}^h, e; f)$ for labor supply, and $b' = g^b(b, w_{-1}^h, e; f)$ for bond savings. Note that for non-resetting households ($f = 0$), trivially $g^w(b, w_{-1}^h, e; 0) \equiv w_{-1}^h$ for all (b, w_{-1}^h, e) , due to the wage friction. Also note that for resetting households ($f = 1$), the state variable w_{-1}^h is irrelevant, and hence superfluous as an argument in the decision rules and the value function. The reason is that a resetting household chooses the wage before her labor-market outcome is formed, so the beginning-of-period wage has no effect on anything. Nevertheless, for completeness, I write out w_{-1}^h as an argument.

One more thing to note is that, due to the labor demand being forcing, the decision rule for labor supply is completely determined by the wage, productivity and exogenous aggregate factors:

$$g^l(b, w_{-1}^h, e; f) = n(w^h, e) \stackrel{\text{by (1.3)}}{=} \left(\frac{g^w(b, w_{-1}^h, e; f)}{W} \right)^{-\delta} N e^{\delta-1}. \quad (1.14)$$

Labor demand, $n(\cdot)$, is a function of the wage. The wage, in turn, is a function of the full state vector and the outcome of the wage friction, f . Hence, labor demand also varies with the state vector and f . It can be useful to think of labor demand as a composite function of the full state vector and the wage-updating status: $n(b, w_{-1}^h, e; f)$.

Recall the result (1.6) that, all else equal, labor demand is increasing with idiosyncratic productivity. This means that with forcing labor demand, increased idiosyncratic productivity without an opportunity to reset the wage *mechanically* leads to higher labor supply. Naturally, this increases the earnings, but also renders a higher disutility from labor supply. Which effect that dominates for welfare depends on the individual state of the household, which is illustrated in section 1.4.1. But it is noteworthy that welfare does not necessarily increase with productivity, which is counter-intuitive.¹⁹

On the other hand, the opposite situation illustrates why Calvo (1983) wage stickiness provides a proxy for unemployment. If stuck with a high wage during a period of falling idiosyncratic productivity, labor demand falls, causing the hours worked to drop involuntarily. Although this occurs on the intensive margin in the model economy considered here, it has obvious similarities with the more binary states of employment or unemployment in the real world.

Version 2: sticky hourly wage and a non-forcing labor demand

Version 2 differs from version 1 in that labor demand is not forcing, but instead works as a maximum for labor supply. A household solves

$$V(b, w_{-1}^h, e; f) = \max_{c, w^h, l, b'} \left\{ U(c, l) + \beta E \left[V(b', w^h, e'; f') \right] \right\} \quad (1.15)$$

subject to a budget constraint

$$c + \frac{b'}{R} = w^h l + b + D - T, \quad (1.16)$$

¹⁹With forcing labor demand, the worker commits to supplying the demanded labor. In the real world, we never see workers supplying very high amounts of labor *against their will*, just because they have a low wage combined with high abilities.

a non-forcing labor-demand function

$$l \leq n^*(w^h, e), \quad (1.17)$$

and the wage-updating friction

$$w^h = w_{-1}^h \quad \text{if } f = 0.$$

Equation (1.17) is the only difference from version 1. Specifically, note the inequality. Labor demand $n^*(w^h, e)$ is given by equation (1.7) from the firm's problem.

An important difference from version 1 is that in version 2, labor supply, denoted by the decision rule $g^l(b, w_{-1}^h, e; f)$, is a “true” choice variable in the sense that it is not completely determined by the wage choice. I.e., equation (1.14) is not valid in version 2. Although equation (1.9) teaches us that labor demand is increasing with idiosyncratic productivity, labor supply is not necessarily so. And there is no reason to believe that individual welfare could decrease as idiosyncratic productivity increases, all else equal.

Version 3: sticky effective wage and a non-forcing labor demand

Version 3 differs from version 2 in the sense that the effective, not the hourly, wage is sticky. This means that a non-updating household keeps the effective wage from the previous period. If her idiosyncratic productivity changes, the hourly wage changes one-to-one, see equation (1.11). This is typically not how a wage contract is set up for an employee, but might be a better description of the labor income of someone self-employed. A household solves

$$V(b, w_{-1}^e, e; f) = \max_{c, w^e, l, b'} \{U(c, l) + \beta E[V(b', w^e, e'; f')]\}$$

subject to a budget constraint

$$c + \frac{b'}{R} = w^e e l + b + D - T,$$

a non-forcing labor-demand function

$$l \leq n^*(w^e, e), \quad (1.18)$$

and the wage-updating friction

$$w^e = w_{-1}^e \quad \text{if } f = 0.$$

$n^*(w^e, e)$ in equation (1.18) is given by equation (1.8). Note that in versions 3 and 4, a household chooses the effective wage, so $g^w(\cdot)$ refers to the effective wage, w^e , not the hourly wage, w^h .

Version 4: sticky effective wage and a forcing labor demand

Finally, version 4 differs from version 3 in that labor demand is forcing. It also differs from version 1 in that it is the effective, not the hourly, wage that is sticky. A household solves

$$V(b, w_{-1}^e, e; f) = \max_{c, w^e, l, b'} \{U(c, l) + \beta E[V(b', w^e, e'; f')]\}$$

subject to a budget constraint

$$c + \frac{b'}{R} = w^e e l + b + D - T,$$

a forcing labor-demand function

$$l = n(w^e, e),$$

and the wage-updating friction

$$w^e = w_{-1}^e \quad \text{if } f = 0.$$

$n(w^e, e)$ in equation (1.18) is given by equation (1.5).

Common to all four versions of the model is the optimality condition that applies to any household not constrained by the borrowing limit, also known as the *Euler equation*:

$$\frac{\partial U(c, l)}{\partial c} = \beta RE \left[\frac{\partial U(c', l')}{\partial c'} \right]. \quad (1.19)$$

Although not relevant to all households, it holds for households not constrained by the borrowing limit, and plays an important roll in solving the model, which is discussed in detail in appendix section 1.A.2.

Employment status

There are different ways of modeling employment status in a macroeconomic model. One of the most popular ways stems from work by Diamond-Mortensen-Pissarides, see, e.g., Pissarides (1985), and builds on vacancy postings, job searching, and matching of searchers with vacancies. In that class of models, the employment status changes on the extensive margin, and is either zero or one. I take a different approach and let any deviation from a full employment status of a household stem from a wage-productivity mismatch. This, in turn, is a consequence of the Calvo (1983)-type wage friction. With this approach, the employment status of a household changes on the intensive margin, and is a real positive number.²⁰ Since the source of fluctuations in employment is the inability to reset the wage, I find it natural to define the employment status as the actual labor supply relative to the hypothetical supply that would prevail if the household could reset her wage:

$$\text{Employment status} = \frac{g^l(b, w_{-1}, e; f)}{g^l(b, w_{-1}, e; 1)}, \quad (1.20)$$

where w represents w^h for versions 1 and 2, and w^e for versions 3 and 4. Note that, for any wage-updating household ($f = 1$), the employment status is trivially equal to one.

I refer to a value of the employment status strictly below one as *underemployment*. Further, it is meaningful to distinguish between two types of underemployment: *involuntary* and *voluntary*, where the former can be thought of as unemployment, while the latter reminds us of suppressed labor-force participation. In addition to an employment status < 1 , involuntary underemployment is characterized by $g^l(b, w_{-1}, e; f) = n^*(b, w_{-1}, e; f)$, i.e., labor demand is a binding constraint.²¹ Thus, involuntary underemployment stems from the wage being “too high”, given

²⁰The way I define employment status, it can be above 100%.

²¹Note that if labor demand is forcing, as in versions 1 and 4, labor demand binds by default, and hence all underemployment is involuntary.

the current idiosyncratic productivity.

Voluntary underemployment, on the other hand, is characterized by employment status < 1 and $g^l(b, w_{-1}, e; f) < n^*(b, w_{-1}, e; f)$, so it is a deliberate choice by the household to supply less labor than she would if she could reset her wage. I.e., voluntary underemployment stems from the wage being “too low”, given the current idiosyncratic productivity.

Distribution of households

The economy modeled in this paper features rich household heterogeneity, by which I mean that the state vector of a household is three-dimensional, and hence there is a three-dimensional joint distribution of households. The three dimensions are wealth, (beginning-of-period) wage and productivity, (b, w_{-1}, e) where w_{-1} denotes the hourly (w_{-1}^h) or effective (w_{-1}^e) wage depending on the version considered. Wealth and wage are endogenous choice variables and thus, the marginal distributions in these dimensions are model outcomes, and differ between the model versions. However, idiosyncratic productivity is exogenous, so the marginal distribution in this dimension is given by the ergodic properties of the productivity process (1.10), and is hence the same for all model versions.

It is necessary to know the distribution of households to calculate aggregate variables, such as the total demand for government bonds, the aggregate labor, and the aggregate demand for consumption goods. The analysis in this chapter is performed in partial equilibrium, without any clearing of the mentioned markets, and it is not necessary to keep track of the distribution of households to address the questions asked here. However, for an analysis in general equilibrium, the distribution of households is absolutely essential, and is hence in more focus in the accompanying chapter Åhl (2020).

Summary of the households

To summarize, the households in this model choose how much to consume, save, and work, and also what wage to use in the labor market. This is done subject to a stochastic process for idiosyncratic productivity, a wage friction, market incompleteness, a borrowing constraint, and a labor-demand function. Wage-productivity mismatches, caused by the wage

friction in combination with idiosyncratic variations in productivity, cause the employment status of a household to vary. The main focus is on involuntary underemployment.

Solving the household's problem described above is challenging, mainly because there are two endogenous state variables. The algorithm I develop to solve the household's problem contributes to the literature on generalizations of the endogenous-gridpoint method by Carroll (2006). The details are given in appendix section 1.A.2.

1.3 Choice of parameter values

This section presents, and very briefly discusses, the parameter values used in the analysis. The aim is to use parameter values that are non-controversial. The values used are found in table 1.2.

Table 1.2: Parameter values

<i>Parameter</i>	<i>Value</i>	<i>Interpretation</i>	<i>Comment</i>
β	0.99	Discount factor	Quarterly, standard
δ	6	Subst. elast., skill types	Galí (2008), chapter 6
θ_w	0.75	$1 - Pr(\text{update wage})$	Galí (2008), chapter 6
ρ_e	0.966	Persistence, idios. prod.	McKay et al. (2016)
σ_e	0.13	St. dev., idios. prod. innov.	McKay et al. (2016)
σ	2	Risk aversion	McKay et al. (2016)
φ	2	Inverse Frisch elasticity	McKay et al. (2016)
$R - 1$	-0.88%	Net quarterly real rate	
W	0.57	Real wage index	
D	0.41	Firm dividends	
T	-0.01	Government tax (transfer if < 0)	
N	0.82	Aggregate labor demand	
N_d	0.07	Labor demand from responsive households	
W_d	0.93	Real wage index for responsive households	

The time-discount factor $\beta = 0.99$ means that a household values utility one period into the future one percent less than utility now. One model period is to be interpreted as one quarter of a year. An elasticity of substitution between skill types, δ , of 6 is widely used in the literature.

$\theta_w = 3/4$ corresponds to, on average, one year between wage updates, which is in line with the literature.²²

I discretize the idiosyncratic productivity process (1.10), as in Rouwenhorst (1995), to 15 possible levels. The process is governed by persistence parameter ρ_e , and volatility parameter σ_e , which are estimated on micro data in Floden and Lindé (2001), and transformed to quarterly frequency in McKay et al. (2016). The resulting ergodic distribution of e is illustrated in figure 1.1. Median idiosyncratic productivity is 1, while the average is 1.13.

I use a standard MaCurdy (1981) instant-utility function:

$$U(c, l) = \frac{c^{1-\sigma} - 1}{1-\sigma} - \frac{l^{1+\varphi}}{1+\varphi}, \quad (1.21)$$

with preference parameters $\sigma = 2$ for risk aversion, and $1/\varphi = 1/2$ for the Frisch elasticity. Both are standard in the literature. This choice of utility function, separable in consumption and leisure, renders a diminishing marginal utility of consumption, and an increasing marginal disutility of labor:

$$\begin{aligned} \frac{\partial U}{\partial c}(c) &= c^{-\sigma} \\ \frac{\partial U}{\partial l}(l) &= -l^\varphi. \end{aligned} \quad (1.22)$$

As for the aggregate steady-state variables, I refer to the general-equilibrium study in Åhl (2020). The strategy is to use the steady-state values for version 2, the preferred choice. The steady state differs between the different model versions; especially the version with sticky hourly wages and forcing labor demand stands out; but to facilitate the comparison, I keep these values fixed for all model versions in this analysis.

Some things worth noting are that the real interest rate is well into negative territory, -3.5% annualized. This yields taxes that are negative, i.e., the households receives lump-sum transfers from the government in each period. However, these transfers are quite small compared to the dividends received lump-sum from firms. The dividends are large, contributing 46% of the average overall income.

²²See, e.g., Barattieri et al. (2014).




Figure 1.1: Ergodic distribution of the idiosyncratic productivity, e .

1.4 Results

This section presents and discusses the results of the paper. The emphasis is on the intuition behind the results, and a comparison of the different model versions. The versions are handled in turn.

1.4.1 Version 1: sticky hourly wage and a forcing labor demand

In version 1, the hourly wage is subject to the wage friction. This means that if the idiosyncratic productivity of a household changes from one period to another, without an opportunity to update the wage, the effective wage changes one-for-one in the opposite direction, see equation (1.11).²³ If the idiosyncratic productivity rises, the effective wage drops proportionally, and the other way around. This is important since it is the effective wage that determines the demand for labor, or more precisely the demand for *effective* labor, el ; see equation (1.3).²⁴ As a result, during a period with a fixed hourly wage, an increase in the idiosyncratic productivity results

²³Equation (1.11) again: $w^h = w^e e$.

²⁴Equation (1.3) again:

$$n(i) = \left(\frac{g^w(i)}{W} \right)^{-\delta} N e(i)^{\delta-1}.$$

in a lower effective wage, and thus a higher demand for that specific household's labor.

A helpful way of understanding the mechanisms of the model, and illustrating interesting situations, is to study a simulation of one single household for a large number of periods. The studied household is subject to exogenous idiosyncratic shocks to her productivity, and responds endogenously. As a result, she moves around in the distribution of households. If the simulation is long enough, the time she spends in different regions of the state space corresponds to the economy-wide density of households in that region. However, I do not focus on the distribution of households here. Instead, I dig deeper into a couple of extra interesting episodes during a long simulation. These episodes are shown in figures 1.2–1.4. All three figures cover the same two episodes, but focus on different things. Both episodes are characterized by a spell of fixed wage and a changing idiosyncratic productivity. During the first episode, contained in periods 374–384, the productivity decreases; while in the second episode, contained in periods 445–458, the productivity increases. The period numbers lack meaning, and can be viewed as completely arbitrary, but it might be helpful to keep in mind that a period represents a quarter of a year.

Figure 1.2 focuses on labor-market outcomes: wage, hours worked, and employment status. The simulated household enters the first episode with a high productivity (solid line), about double the economy-wide average, and also a high wage (dashed line). In period 376, she gets the opportunity to reset her wage, but has no reason to change it from the current level. Then follow five quarters with a fixed wage. The fixed-wage spell is, however, unknown to her when she sets her wage in period 376. She only knows the constant probability of resetting her wage in every future period. At this point, her high wage matches her high productivity, so the demand for her labor services is at a reasonable level. By definition,

Moving $e(i)^{-1}$ to the left-hand side, and $e(i)^\delta$ to the numerator inside the parenthesis, gives effective labor on the left-hand side, and effective wage in the numerator on the right-hand side: $e(i)n(i) = \left(\frac{g^w(i)/e(i)}{W} \right)^{-\delta} N$.

her employment status is 1 in the resetting period, see equation (1.20).²⁵

In the first quarter of the fixed-wage spell, period 377, her productivity drops quite drastically. In the new situation, she produces less output for each hour worked. With the hourly wage being fixed, this means that the firm's wage cost per output produced by her increases, i.e., her effective wage has increased. At this higher effective wage, the firm still wants to buy her labor services, since she is not perfectly substitutable for a cheaper employee, but not as much as before her productivity drop. The firm now demands less output from her, suggesting that she would work less hours. But she is also less productive, suggesting she will have to work more hours to meet any demand for effective labor. In the dotted line for labor supply, we see that the former effect dominates and, in fact, her hours worked drop to close to zero.²⁶ To be clear, her effective labor input, el , which is what is valued by the firm, drops even more.

Although enjoying more leisure, this situation is sub-optimal to the household in the sense that she would prefer a lower wage, which would render the demand for her labor services to increase, and she would work and earn more. This is illustrated by the dash-dotted line for employment status being far below 1. Period 377 marks the beginning of an underemployment spell, and since it is the result of a drop in labor demand, it is involuntary. In fact, since labor demand is forcing in version 1, all underemployment must be involuntary.

In period 378, productivity jumps up, but not all the way to the original level. This increases the labor demand somewhat, but not enough to get the employment status above 0.5. In the subsequent periods, productivity drops further, down to the economy-wide median level. Since the household is stuck with a far-above average wage, this brings down both the labor demand and the employment status to close to zero.

In period 382, an opportunity to reset the wage finally appears. The household's idiosyncratic productivity remains flat at the median level, and we can see that the new wage, which matches this productivity level,

²⁵Equation (1.20) again:

$$\text{Employment status} = \frac{g^l(b, w_{-1}, e; f)}{g^l(b, w_{-1}, e; 1)}.$$

²⁶The elasticity of substitution between skill types, δ , determines the relative size of the two effects.

is substantially lower than the previous wage, which was set with twice as high productivity.

When the opportunity to update the wage occurs, the employment status mechanically jumps to 1. However, note that the desired level for hours worked is less than half of what it was before the underemployment spell. This is not primarily explained by the instant disutility from labor supply, but rather by a fear of an extreme future labor supply — a mechanism explained in the second episode of interest.




Figure 1.2: Simulation of one household, selected illustrative episodes. The horizontal axis shows the simulation period. Vertical dotted lines indicate periods of wage updating, $f = 1$. The series shown are:

Idiosyncratic productivity (e), solid line with dots

Hourly wage (w^h), dashed line

Labor supply (l), dotted line

Employment status, see equation (1.20), dash-dotted line (peaks at 8.2 in period 453 in the right-hand side panel)

The next episode of interest, presented in the right-hand panel of figure 1.2, begins in period 447, with an opportunity to reset the wage with an idiosyncratic productivity slightly above average. Not much happens until period 451, when productivity exogenously increases. With the hourly wage fixed, the higher productivity means that the effective wage decreases, which in turn increases the demand for effective labor from the simulated household. The higher productivity is not enough to meet the increased demand for effective labor, so also the hours worked must increase. The household has no choice in this, since the labor demand is forcing and

hence must be met. However, the figure reveals that this is undesirable for the household. The employment status rises to around 200% of the desired level, meaning that the household would prefer a higher wage, which would cut the labor demand to half of the current level.

In periods 452–453, the productivity increases even more. This enhances the increased forcing labor demand much more. In fact, the household is forced to work 820% of what she would prefer, if she got the opportunity to reset the wage. The figure is cut below this, to keep other details visible. It is worth mentioning that this extreme labor supply stems from the forcing labor demand, and not from a desire for high earnings. This is a severe case of overemployment.

In period 454, productivity drops somewhat, bringing the hours worked back to more reasonable levels, although still more than twice the desired level. In period 456, the household resets her wage at a high enough level to bring labor demand down to the desired level.

With an increasing marginal disutility of labor, see equation (1.22), the extremely high number of hours worked caused by the wage friction and the positive shocks to idiosyncratic productivity, are very costly to a household in terms of welfare.²⁷ Since the stationarity of the autoregressive productivity process, see equation (1.10), makes the future productivity likely to increase when the level is low, the households are very reluctant to choosing a low wage in this situation.²⁸ A low wage would likely lead to very high labor supply before the next opportunity to reset. It is optimal to avoid this by choosing a higher wage, suppressing the labor demand, and hence earnings, to a very low level. This is what causes the low labor supply in periods 382–384, compared to periods 374–376, in the simulation. It is also illustrated, in comparison to versions 2 and 3, in appendix figures 1.12 and 1.13.

Figure 1.3 contains the same idiosyncratic outcome for productivity as

²⁷Equation (1.22) again:

$$\frac{\partial U}{\partial l}(l) = -l^\varphi.$$

²⁸Equation (1.10) again:

$$\log(e_t) = \rho_e \log(e_{t-1}) + \varepsilon_t^e, \quad \varepsilon_t^e \sim N(0, \sigma_e).$$

figure 1.2, but shifts the focus from the labor market to the consumption-savings decision of the simulated household during the same episodes. Beginning with the leftmost episode of underemployment, we see that the drop in hours worked causes the earnings to be very low during this period.²⁹ The household enters the period with some, but not much, wealth. In response to the lower income, to keep up consumption, she starts dis-saving, and the wealth starts decreasing from period 378.³⁰ The dis-saving continues and at the end of the underemployment spell, she is at the borrowing limit. Interestingly, she chooses to stay at the borrowing limit also after the underemployment spell has ended in period 382. The reason is the low earnings caused by the low labor supply discussed above.

The dis-saving is not enough to smooth consumption at the initial level when the negative productivity shock hits, so it drops in response to the lower income. During the rest of the underemployment spell, consumption is kept rather smooth, with a slight decrease caused by both further lowering of the income, and less room for dis-saving as the borrowing limit approaches. After the underemployment spell, the new lower earnings allows consumption to recover only slightly, but not near the pre-underemployment level.

In the second episode of interest, featuring radical overemployment, the household enters with zero wealth and rather low earnings, leaving no room for building wealth. However, when the productivity rises, in periods 451–452, the earnings rise drastically due to the very high labor supply. In response to the increased earnings, consumption rises to about twice the beginning-of-episode level. Even more dramatic is the increase in wealth, which in only two quarters goes from zero to more than the average annual income in the economy. In my view, these are involuntary savings, rather than driven by precaution, consumption-smoothing, or any other motive.

In period 456, when the overemployment is ended by an opportunity to increase the wage, the earnings come down. However, consumption is kept high by dis-saving of the now large wealth that was built up in a short period of time. A more comprehensive overview of the savings and

²⁹Recall that although the earnings are close to zero, the income is still substantially higher due to dividends from firm profits.

³⁰Wealth in period 377 is predetermined, and was decided in period 376, before the productivity dropped.




Figure 1.3: Simulation of one household, selected illustrative episodes. The horizontal axis shows the simulation period. Vertical dotted lines indicate periods of wage updating, $f = 1$. The series shown are:

Idiosyncratic productivity (e), solid line

Wealth (b), dashed line

Consumption (c), dotted line

Earnings (w^{hl}), dash-dotted line

consumption decisions of non-updating ($f = 0$) households is shown in appendix figures 1.14 and 1.15.

Before leaving version 1, let us have a look at the welfare of the simulated household, in terms of value; see the Bellman equation (1.12).³¹ It is quite intuitive to believe that the value is co-varying positively with the idiosyncratic productivity. The reason is that a higher productivity allows more output in a shorter period of time worked. Higher output renders higher earnings, which can be traded for higher consumption in the present, or in the future via savings. Higher consumption yields a higher utility, see the utility function (1.21).³² And the less time spent working yields less instant disutility, i.e., a higher utility.

³¹Equation (1.12) again:

$$V(b, w_{-1}^h, e; f) = \max_{c, w^h, l, b'} \left\{ U(c, l) + \beta E \left[V(b', w^h, e'; f') \right] \right\}.$$

³²Utility function (1.21) again:

$$U(c, l) = \frac{c^{1-\sigma} - 1}{1-\sigma} - \frac{l^{1+\varphi}}{1+\varphi}.$$




Figure 1.4: Simulation of one household, selected illustrative episodes. The horizontal axis shows the simulation period. Vertical dotted lines indicate periods of wage updating, $f = 1$. The series shown are:
Idiosyncratic productivity (e), solid line (left-hand scale)
Value, $V(\cdot)$, dashed line (right-hand scale)

Figure 1.4 yet again shows the same idiosyncratic outcome for productivity as figures 1.2 and 1.3, during the same episodes, but it also shows the evolution of the value function.³³ In the first episode, featuring underemployment, the value behaves in line with the intuition, co-varying very closely with productivity, for precisely the reasons described above. What is interesting is the second period, featuring the radical overemployment. We can clearly see that as the labor supply and employment status shoots through the roof, in period 452, the close co-variance of value and productivity is spectacularly broken. Instead, the value drops a great deal. The failure of the intuition in this case is that with a forcing labor demand, the household is not able to choose how to trade-off the higher productivity between instant leisure and higher earnings, and hence consumption. In this case, she is so valuable to the firm that she is forced to work more, although she would prefer to enjoy more leisure. This set-up has a flavor of “slavery”, in the sense that she is forced to work much more than desired. The big difference is that the work is paid, and she becomes very rich as a result. Nevertheless, she suffers from her increased productivity. This is also illustrated in appendix figure 1.17.

³³Note the different scales on the vertical axes: productivity on the left-hand scale, and value on the right-hand scale.

Figure 1.18 in the appendix shows contemporaneous and lagged correlations between individual-level variables. Although cumbersome to read, there are some insights to be gained from it. One is that the value is negatively correlated with both labor supply and savings, for the reasons discussed above. However, given that the average number of periods between wage updates is quite small, a household suffering from overemployment is likely to escape the spell soon, and then with some substantial wealth. This is reflected in the fact that although a high labor supply today tends to make a household miserable today, she is likely to have a drastically increased value just a few periods into the future.

To summarize version 1, sticky hourly wages and forcing labor demand capture involuntary underemployment caused by a suppressed labor demand. In particular, the labor demand is sensitive to a wage that is high compared to the idiosyncratic productivity. The way for a household to escape an underemployment spell, if the productivity does not rise exogenously, is to wait for an opportunity to reset a lower wage. This is in line with the intuition of underemployment being a proxy for real-world unemployment.

On the other hand, due to the symmetric nature of the labor market, a mismatch where the wage is low in comparison to the productivity causes a very high labor demand. The forcing labor demand results in severe suffering, despite very high earnings and consumption, due to involuntary overemployment. It also results in very volatile involuntary savings. Moreover, the fear of overemployment prevents the households from choosing low wages when being low-productive, which, in turn, suppresses aggregate labor supply and hence aggregate output. This is not in line with any intuition of the real-world economy.

1.4.2 Version 2: sticky hourly wage and a non-forcing labor demand

Version 2 is similar to version 1 in that the wage friction works in the same way, i.e., the hourly wage is subject to the friction. However, it deviates in that the labor demand is not forcing, but rather constitutes a ceiling to what labor each household can supply.

I study version 2 with the aid of the same type of simulation analysis as was used for version 1, i.e., simulating one household for many periods,

and focusing on a couple of interesting episodes in the simulation. To facilitate the comparison, I use the same sequence of (stochastic) exogenous outcomes for the idiosyncratic productivity and the wage-updating opportunities, and also focus on the same episodes as in version 1.

Figure 1.5 shows the labor-market outcomes from the two highlighted episodes. The first episode is entered with a high productivity, more than twice that of the median household. The wage is also high, on par with the productivity. Hence, as the simulated household gets an opportunity to reset the wage in period 376, she chooses to stay with the same wage. However, in the next period, her productivity drops and the wage is fixed. With her lower productivity, she produces less output in every hour worked, and is hence less valuable to the firm.³⁴ In response, the demand for her labor services comes down, and remains low (with some fluctuation due to the variation in idiosyncratic productivity) until the next opportunity she gets to update her wage, in period 382. During this spell, the low labor demand is binding, indicated by the grey shading in the figure. Moreover, the binding labor demand is substantially below the labor-supply level the household would prefer, if she could reset her wage (lower). This is indicated by the employment status far below 1. The low employment status indicates a spell of underemployment, and the binding labor demand makes it involuntary.

Another perspective of the situation is that the desired labor supply is an upward-sloping function of the wage (all else equal), while demand is downward-sloping in the wage, see equation (1.7).³⁵ Without any wage friction, the outcome of each household would always be at the intersection of the desired labor supply and demand. However, with the wage friction, this is typically not the case. The actual outcome of hours worked, when labor demand is not forcing, is instead at the minimum of supply and demand. Figure 1.6 shows demand curves and supply curves for households unable to update the wage, $f = 0$. It does so for three arbitrary example levels of idiosyncratic productivity; just below half the median, the median,

³⁴The situation is very similar to version 1, described above. Some more details are given in that description.

³⁵Equations (1.7) again:

$$n^*(i) = \left(\frac{g^w(i)}{W_d} \right)^{-\delta} N_d e(i)^{\delta-1}.$$

and just above twice the median. The figure also indicates the wage choice of an updating household, $f = 1$. The underemployment spell of the leftmost episode in figure 1.5 can be understood, qualitatively, as the wage being set at the yellow circle in the rightmost panel of figure 1.6, and is then fixed for a number of periods. In the coming periods, productivity drops and the curves shift towards those of the middle panel. Since the wage is stuck at a higher level than the intersection of the curves, the demand is lower than the supply, and hence binding and determines the actual hours worked. As can be seen, demand is close to zero at this wage-productivity combination, leaving the household severely underemployed; or *unemployed*, if you prefer that interpretation.

Returning to the simulation in figure 1.5, the involuntary underemployment spell is ended by the household resetting a much lower wage in period 382. Labor demand jumps back up, as does the labor supply. Although difficult to read from the figure, the new labor-supply level is slightly lower than prior to the underemployment spell. The reason is that the wage is lower, and at this lower wage, it is not as advantageous for the household to work extra hours. Indirectly, this is because the lower wage was set with a lower productivity. A pattern of more productive households choosing to work more is common in this class of models, and holds for updating households, $f = 1$. The case of non-updating households, $f = 0$, is handled in the second episode of interest, discussed below.

Another thing worth noting in the first episode of interest is that the labor demand binds *only* during the underemployment spell, not in periods of wage updating or the other periods in the left-hand panel of figure 1.5. Focusing on the wage-updating periods 376 and 382, this means that the household could choose a wage that is exactly so much higher that the labor demand perfectly matches the desired labor supply. This would mean the exact same amount of hours worked, but higher earnings due to the higher wage, in that period.³⁶ If the wage-setting problem was static, the household would push up the wage until the labor demand coincides with her desired supply. However, the dynamic nature of the problem, imposed by the wage friction, makes it more complicated. Choosing a

³⁶This can also be seen, just barely, in figure 1.6, where the chosen wage tends to be slightly below the intersection of supply and demand.

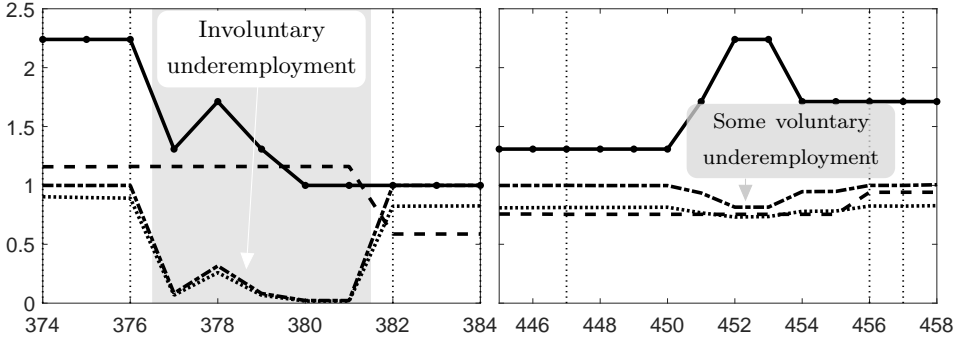


Figure 1.5: Simulation of one household, selected illustrative episodes. The horizontal axis shows the simulation period. Vertical dotted lines indicate periods of wage updating, $f = 1$. Shaded areas indicate periods where labor demand is binding. The series shown are:

Idiosyncratic productivity (e), solid —

Hourly wage (w^h), dashed - -

Labor supply (l), dotted ···

Employment status, see equation (1.20), dash-dotted - · -

wage just below the one making supply and demand coincide works as an insurance against involuntary underemployment in the event of negative future shocks to the idiosyncratic productivity. This can also be seen in appendix figure 1.16.

The second episode of interest in figure 1.5 covers a fixed-wage spell between periods 447 and 456. The household resets the wage in the beginning, without any need to change it. Some periods of unchanged productivity are followed by a couple of periods when the household's productivity rises substantially. This means that she produces more value to the firm in each hour she works; and since her hourly wage is fixed, it also means that the demand for her labor services rises.³⁷ However, labor demand is not forcing in version 2, and she chooses not to supply more labor in response to the increased demand. This is a big difference from version 1.³⁸ Instead of obeying the labor demand, her optimal response

³⁷Labor demand is not plotted in figure 1.5, but it behaves similarly to the actual hours worked in version 1, shown in figure 1.2. Figure 1.6 might also be helpful in understanding the relation between the labor demand, the wage, and the productivity.

³⁸A comparison of figures 1.2 for version 1 and 1.5 for version 2 illustrates the




Figure 1.6: Supply and demand for labor. The columns represent different idiosyncratic productivity levels. Blue is **labor demand** from the firm, $n^*(w^h, e)$, red is **labor supply** by a household with fixed wage, $g^l(b, w_{-1}, e; 0)$, and yellow circles indicate the **wage choice** by an updating household, $g^w(b, w_{-1}, e; 1)$. *Note:* both the labor supply and the wage choice are also affected by the state variable wealth, b , but only marginally. For simplicity, I have fixed wealth to $b = 1$ throughout the figure.

is to slightly lower her labor supply. The reason for this is that with a higher idiosyncratic productivity, her outlooks for the future are better due to the persistence of the productivity process. Hence, she can afford to enjoy some more leisure presently. This can be seen in a slight shift to the left of the red labor-supply curve as the idiosyncratic productivity increases in figure 1.6.

As mentioned above, a lower labor supply in response to a higher idiosyncratic productivity is an uncommon result in this class of models. Typically, the optimal response to a positive productivity shock is to take the opportunity to work more, while the wage is high, and enjoy more leisure in the future. The reason why this is not the case here is that the wage is fixed, so more hours worked do not pay off, despite the higher productivity. However, if the wage could be reset, it would be optimal to take the opportunity to work more hours. The employment status is defined in relation to the hypothetical labor supply if the wage could be updated, so that the numerator on the right-hand side of equation (1.20) goes down, while the denominator goes up, resulting in a lower em-

importance of whether the labor demand is assumed to be forcing or not well.

ployment status and some underemployment.³⁹ However, note that there is a big conceptual difference between this underemployment, explained by the *wage being lower* than motivated by the productivity, and the underemployment in the previous episode, explained by a suppressed labor demand due to a *higher wage* than what is motivated by the productivity. The underemployment in the second episode is what I refer to as *voluntary*.

Underemployment is voluntary if the labor demand is not binding, and involuntary if the labor demand is binding. One perspective of this is that, for any idiosyncratic productivity level, a wage stuck below the optimal wage (marked by yellow circles in figure 1.6) leads to voluntary underemployment, while a wage stuck above the optimal wage renders an involuntary deviation from full employment.⁴⁰ With this in mind, the slopes of the demand and desired supply curves hint at the severity of voluntary and involuntary underemployment.

The right-hand panel of figure 1.5 also reveals that the labor demand is never binding throughout this episode. Although not visible in the figure, it is reasonable to believe that the supplied labor is not far from the demanded labor at the beginning and the end of the episode, when the wage matches productivity quite well. But during the period of voluntary underemployment, the gap ought to widen, since the labor demand rises drastically, while the labor supply decreases slightly. This is indeed the case.

Moving on from labor-market outcomes and mechanisms, figure 1.7 shows the same episodes, with the same outcomes for the idiosyncratic productivity and the wage updating as figures 1.2-1.5, but the focus is on consumption and savings. The household enters the first episode with a high productivity and a matching wage, and hence her earnings are high. She enjoys a high and smooth level of consumption, and can at the same time afford to build up her wealth. When the negative shock hits her productivity, she enters the underemployment spell and her earnings drop. Her consumption also decreases, but only slightly. This is made

³⁹Equation (1.20) again:

$$\text{Employment status} = \frac{g^l(b, w_{-1}, e; f)}{g^l(b, w_{-1}, e; 1)}.$$

⁴⁰In the involuntary case, the deviation in employment status might actually be towards a slight overemployment.

possible by a change from wealth building to dis-saving. Her consumption keeps decreasing smoothly throughout the underemployment spell, as her wealth decreases. When the underemployment spell ends in period 382, her earnings increase, but only to about half the pre-spell level. This is explained by both a lower wage, and less hours worked. However, the new earnings level is enough to somewhat increase her consumption, and run a balanced budget.




Figure 1.7: Simulation of one household, selected illustrative episodes. The horizontal axis shows the simulation period. Vertical dotted lines indicate periods of wage updating, $f = 1$. Shaded areas indicate periods where labor demand is binding. The series shown are:

Idiosyncratic productivity (e), solid line
Wealth (b), dashed line
Consumption (c), dotted line
Earnings (w^{hl}), dash-dotted line

In the second episode of interest, with a mild voluntary underemployment, the household enters quite wealthy, and with an above-average productivity. She eats out of her wealth at a moderate pace, and her consumption is kept high and smooth. When the productivity rises, her outlooks for the future, in particular the next opportunity to reset her wage, improves. She enjoys this better outlook by a combination of more leisure and higher consumption. With a fixed wage and decreasing hours worked, her earnings drop slightly, and the increased consumption must be financed by more aggressive dis-saving, motivated by higher expected future earnings. When the opportunity to reset her wage appears in period 456, she is more productive and hence chooses a higher wage, and

her earnings go up enough to start building up more wealth again. Her expectations of higher future earnings, that motivated more dis-saving, materialize.

Intuition tells us that the idiosyncratic productivity and welfare should co-move positively, see the reasoning in section 1.4.1 above. Figure 1.8 shows the time series for the idiosyncratic productivity and the value of the simulated household in the two episodes of interest. It confirms the intuition — also for the second episode, which is not the case in version 1.⁴¹ The positive co-movement of the value with the idiosyncratic productivity is also highlighted, for the entire simulation, by the almost perfect positive contemporaneous correlation between productivity and value presented in appendix figure 1.19.⁴²




Figure 1.8: Simulation of one household, selected illustrative episodes. The horizontal axis shows the simulation period. Vertical dotted lines indicate periods of wage updating, $f = 1$. Shaded areas indicate periods where labor demand is binding. The series shown are:
Idiosyncratic productivity (e), solid line (left scale)
Value, $V(\cdot)$, dashed line (right scale)

Figure 1.19 in the appendix is difficult to read, but contains much information about the dynamics of version 2 of the model. Information that can be related to intuition. One example is that high idiosyncratic productivity today tends to yield high wages, not only today but also in

⁴¹Note also that a wage-updating opportunity can increase the value substantially, despite an unchanged productivity; see, e.g., period 382.
⁴²For version 1, this correlation is also positive but not as strong, see appendix figure 1.18.

the near future. High wages today tend to coincide with high contemporaneous savings, but have no clear co-variance with contemporaneous hours worked.⁴³ However, high wages today tend to yield negative future savings, due to a build-up of wealth that is dis-saved in the future. Moreover, high wages today yield low future hours worked, due to an increased risk for involuntary underemployment. This reasoning is simplified, but the analysis of lag-correlation structures can still provide a better understanding of the mechanisms at play. One last example is that in version 2, where savings are exclusively voluntary, net savings correlate positively with the contemporaneous value. This is not the case in version 1, where a large part of the savings is of a more forced nature.

To summarize version 2, a period of decreasing productivity leads to involuntary underemployment *if* the wage was last reset with the higher productivity. The underemployment can be more or less severe, but can bring down both hours worked and the employment status to close to zero. An underemployment spell can be escaped by either of two exogenous events; an opportunity to reset a lower wage, or a rising idiosyncratic productivity. This mechanism is similar to version 1, and is in line with the intuition.

The opposite case, of a rising idiosyncratic productivity and a fixed wage, is not symmetric, as it was in version 1. A mismatch where the wage is low in relation to the productivity does not lead to overemployment. Instead, it leads to mild voluntary underemployment where the household chooses to supply less labor while waiting for the opportunity to reset a more appropriate (higher) wage.

1.4.3 Version 3: sticky effective wage and a non-forcing labor demand

Version 3 differs from versions 1 and 2 in the nature of the wage friction. In version 3, it is the *effective* wage that is sticky, and remains unchanged in periods when a wage update is not possible. The effective wage refers to the compensation per effective output, as opposed to the hourly wage, which refers to the compensation per time spent working.

⁴³Two mechanisms cancel: on the one hand, the wage is positively correlated with productivity, and high productivity tends to increase the labor supply. On the other hand, (too) high wages tend to cause involuntary underemployment.

However, version 3 is similar to version 2 in the nature of the labor demand, which is not forcing. Instead, a household chooses how much labor to supply, but may never supply more than the demand she faces.

The analysis of version 3 is largely made by discussing the same type of simulation analysis as with versions 1 and 2, i.e., simulating one household for many periods, and focusing on a couple of interesting episodes in the simulation. I again use the same sequence of (stochastic) exogenous outcomes for idiosyncratic productivity and wage-updating opportunities, and also focus on the same episodes as in versions 1 and 2.

Some labor-market outcomes of the simulated household are shown in figure 1.9. Note that although it is the effective wage that is subject to the friction, the figure plots the hourly wage, mainly to be consistent with figures 1.2 and 1.5. This means that, between updates, the hourly wage moves one-for-one with the idiosyncratic productivity.

In the first episode, the wage is reset with a high idiosyncratic productivity in period 376. In the coming periods, the productivity drops in a few steps, while the effective wage remains fixed. An unchanged effective wage means that the firm's demand for effective output remains the same, in accordance with equation (1.8).⁴⁴ As the household is now less productive, i.e., produces less per hour worked, she would have to work more hours to meet the demand for effective labor. However, despite a change in her hourly wage, she is happy to keep labor supply fixed at the initial level. She chooses to deviate from the increased labor demand, which hence stops to bind.

When an opportunity to reset the wage appears in period 382, the household chooses to increase it slightly. The reason is that with a lower productivity, it takes more hours worked to meet the demand for effective labor, and she is compensated less per hour worked. She is not willing to meet the higher demand, and can hence maximize her earnings by allowing a lower labor demand. Throughout the first episode, her employment status remains very close to one.

⁴⁴Equation (1.8) again, with some reshuffling:

$$n^*(i) = \left(\frac{g^w(i)}{W_d} \right)^{-\delta} N_d e(i)^{-1} \iff \underbrace{e(i)n^*(i)}_{\text{Effective labor}} = \left(\frac{g^w(i)}{W_d} \right)^{-\delta} N_d.$$

Note that g^w denotes the effective wage in this version.

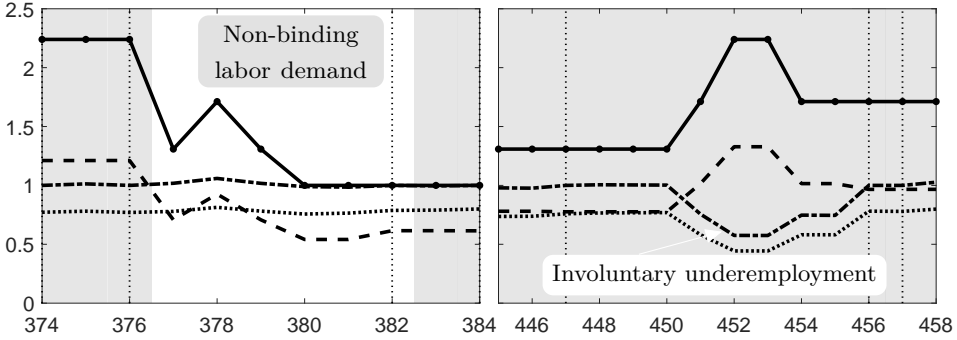






Figure 1.9: Simulation of one household, selected illustrative episodes. The horizontal axis shows the simulation period. Vertical dotted lines indicate periods of wage updating, $f = 1$. Shaded areas indicate periods where labor demand is binding. The series shown are:

Idiosyncratic productivity (e), solid 

Hourly wage (w^h), dashed 

Labor supply (l), dotted 

Employment status, see equation (1.20), dash-dotted 

The simulated household enters the second episode with an idiosyncratic productivity close to the economy-wide average, a wage that matches this well, and hence an employment status very close to one. Some periods into a spell of a fixed wage, her productivity rises, and hence so does her hourly wage. Since the effective wage is fixed, the demand for her effective labor output remains unchanged. However, it takes less labor input from her, in terms of hours worked, to meet this demand. This is sub-optimal for her, since at the higher hourly wage, she would prefer to work more hours, not fewer. Both the decrease in demanded hours worked and the increase in her desired hours worked point towards the employment status dropping, and she enters a spell of underemployment. Since the labor demand binds in this situation, the underemployment is involuntary. The spell is dampened when her productivity comes down somewhat in period 454, but does not end until a lower effective wage can be reset in period 456. Note, however, that her hourly wage is, in fact, higher than when she entered the episode.

The spell of involuntary underemployment in the second episode of

figure 1.9 brings both the hours worked and the employment status down, but not close to zero. In general, the employment status is less volatile with sticky effective wages than with sticky hourly wages. The reason is that with a sticky effective wage, the demand for effective labor from a household is fixed during a spell of fixed wage. This means that the entire variation in hours worked by a household during the spell comes from a variation in the labor input needed to meet that demand. The variation in labor input needed, in turn, stems from the varying idiosyncratic productivity. In that sense, no substitution between households of different skill types is taking place. However, when the hourly wage is fixed, as in versions 1 and 2, the substitution for other skill types is taking place within spells of a fixed wage. It is the substitution that can bring labor demand to close to zero, or to extremely high levels, for individual households. One conclusion is that the labor market is much more sensitive to the elasticity of substitution between skill types, δ , with sticky hourly wages than with sticky effective wages.

For the reason described above, the risk for severe involuntary underemployment is smaller in version 3 than in version 2. This results in less reason to choose a wage below what would result in coinciding labor supply and demand, as an “insurance” against involuntary underemployment.⁴⁵ This, in turn, makes the labor demand bind more often in version 3 than in version 2, which can be seen in a comparison of figure 1.9 with figure 1.5. Moreover, appendix figure 1.16 shows that wage-updating households with higher levels of idiosyncratic productivity choose labor supply such that the demand binds.

Shifting the focus to consumption and savings decisions, figure 1.10 shows the same episodes for the same household as figure 1.9. Just as in version 2, earnings go down as productivity drops after period 376.⁴⁶ However, the reason for the earnings drop is very different in version 3. In version 2, the hourly wage is fixed, so that the entire earnings loss is explained by less hours worked. In version 3, however, we learned that hours worked remain very flat during the fixed-wage spell, so the entire earnings loss is explained by the lower hourly wage.

The savings dynamics is also very different in versions 2 and 3. In

⁴⁵The insurance motive is discussed in more detail in section 1.4.2.

⁴⁶The version-2 equivalent of figure 1.10 is figure 1.7.

version 3, the simulated household enters the fixed-wage spell of the first episode with some, but not very much, wealth. As earnings drop, she starts to dis-save in order to keep the consumption smooth. The consumption is adjusted in the same fashion as earnings qualitatively, but is kept quite smooth with the aid of the wealth adjustments. When the spell ends in period 382, her productivity is less than half of the initial level, and both her earnings and consumption stabilize at lower levels than before the spell.








Figure 1.10: Simulation of one household, selected illustrative episodes. The horizontal axis shows the simulation period. Vertical dotted lines indicate periods of wage updating, $f = 1$. Shaded areas indicate periods where labor demand is binding. The series shown are:

Idiosyncratic productivity (e), solid 

Wealth (b), dashed 

Consumption (c), dotted 

Earnings (w^hl), dash-dotted 

In the beginning of the second episode of interest, productivity, earnings and consumption remain smooth. The consumption is kept up by the household dis-saving at a constant pace. When the positive productivity shock hits in periods 451–452, the increased hourly wage, and the decreased hours worked, cancel, so that the earnings remain flat. However, during this spell of involuntary underemployment, her consumption increases slightly. Since the earnings remain flat, this is financed by more aggressive dis-saving. The reason for this is that with a higher idiosyncratic productivity, her prospects for higher earnings in the future look better, and hence her precautionary motive to save is dampened. Less motive for precautionary

savings leaves room for more consumption on account of her wealth.

The underemployment spell ends in period 456, and the earnings jump up due to a higher hourly (and effective) wage than before the spell. Higher earnings give room for both a slightly higher consumption than before the spell, and positive savings.

Section 1.4.1 shows that with a sticky hourly wage, forcing labor demand may occasionally break the intuitive positive co-movement of idiosyncratic productivity and value. Figure 1.11 shows that, at least during the the two episodes in focus, there is a strong positive relationship between the productivity and the value at the household level. Figure 1.20 in the appendix confirms this for the entire simulation.




Figure 1.11: Simulation of one household, selected illustrative episodes. The horizontal axis shows the simulation period. Vertical dotted lines indicate periods of wage updating, $f = 1$. Shaded areas indicate periods where labor demand is binding. The series shown are: idiosyncratic productivity (e), solid line (left scale) Value, $V(\cdot)$, dashed line (right scale)

A further study of figure 1.20 in the appendix can reveal more about the mechanisms at play in version 3, and how they relate to versions 1 and 2. One example is that, as is also the case in the second episode of interest for the simulation, a hike in idiosyncratic productivity leads to a lower labor supply if the wage cannot be adjusted. When hours worked drop for this reason, an opportunity to adjust the wage is likely to occur in the near future. Due to persistence in the process governing the idiosyncratic productivity, it is likely to be higher also when a wage-setting opportunity occurs, and thus lead to a higher (hourly) wage and higher earnings in the

future. For this reason, labor supply is negatively correlated with future earnings, consumption and net savings in version 3. This is entirely an consequence of the wage friction being on the effective, not the hourly, wage, and is hence not the case in version 2.

To summarize version 3, involuntary underemployment is caused by increasing idiosyncratic productivity during a fixed-wage spell. As a proxy for unemployment, this makes version 3 a poor choice. This is further highlighted by an increasing consumption during an involuntary underemployment spell. Moreover, a fixed effective wage at the household level also fixates the demand for effective labor supply. This makes substitution between households with different skill types, as a result of the wage friction, much less significant than when the hourly wage is sticky. This result is independent of the elasticity of substitution between the skill types. As a consequence, hours worked and the employment status vary a great deal less at the individual level with sticky effective wages than with sticky hourly wages.

1.4.4 Summary and version 4

The above sections have, quite thoroughly, discussed the mechanisms driving household behavior in three versions of a model with incomplete markets, idiosyncratic shocks to productivity, and wage stickiness at the household level. An overall conclusion to be drawn is that the assumptions regarding the nature of the wage friction, and the rules in the labor market, are of first-order importance for the study of the household behavior. For example, assuming that the labor demand is forcing when the friction applies to the hourly wage causes households to occasionally work many times more than desired. On the other hand, assuming that the friction applies to the effective wage results in high-productive households being the group suffering most from underemployment in the sense of the labor demand binding at a level significantly lower than the desired supply.

In light of the above, I conclude that version 2, assuming a sticky hourly wage and the labor supply not being forced by demand at the household level, is the preferred choice. That version is in line with the intuition regarding the labor-market dynamics in general at the household level, and in particular regarding which type of households suffer from involuntary underemployment.

Version 4 of the model, assuming sticky effective wages and a forcing labor demand, has so far been omitted from the discussion in section 1.4. The reason is that it deviates in both dimensions from the preferred version 2. Broadly, the critiques against both versions 1 and 3 also apply to version 4, making it the least interesting one to discuss. Hence I choose to not present any results here; they are, however, available upon request.

In the accompanying chapter Åhl (2020), I present some sensitivity analysis regarding some of the key model parameters. However, this is only carried out for the preferred version, i.e., version 2.

1.5 Conclusions

Wage frictions are realistic and known to improve the mechanisms and the performance of macroeconomic models. However, the interaction with household heterogeneity remains unexplored. In this paper, I have introduced a wage friction at the household level in an otherwise standard macroeconomic model with heterogeneous households. When the productivity varies between households, and over time for each household, there are two different wages to consider: the compensation per time worked (hourly wage), and the compensation per output produced (effective wage). I investigated the consequence of letting the wage friction apply to each of these, and concluded that in terms of realism at the micro level, a sticky hourly wage is preferred. However, the common assumption of a forcing labor demand in combination with wage stickiness turns out to be far from innocuous in this setting — a consequence of highly volatile idiosyncratic productivity. Hence, I also investigated the impact of relaxing the assumption of a forcing labor demand, which turned out to improve the model further in terms of realism.

The main conclusions, besides sticky hourly wages and a non-forcing labor demand being the preferred modeling choices, are about micro mechanisms. A wage-productivity mismatch at the individual level leads to involuntary underemployment if the wage is too high, and to a potentially large gap between supplied and demanded labor if the wage is too low. The involuntary underemployment, caused by a low demand for a specific household's labor services, can be severe and push hours worked down to almost zero. It is natural to view this as a proxy for real-world

unemployment.

The analysis was carried out with an exogenous aggregate wage and interest rate, and thus there is no clearing of the labor market and the market for government bonds. This makes the analysis suitable for studying micro behavior, but not the outcome of aggregate variables. Closing the model to general equilibrium is a natural extension, which I pursue in a separate paper, Åhl (2020). In this paper, I have also showed that ag forcing labor demand yields unrealistic results in the presence of idiosyncratic productivity shocks and sticky wages. Nevertheless, the common assumption in models with sticky prices is that demand is forcing in the goods market. Comparing forcing and non-forcing goods demand in an economy with price stickiness and idiosyncratic productivity shocks in the firm sector would be interesting.

References

- Åhl, M. (2020). Wage stickiness and household heterogeneity in general equilibrium. Thesis chapter 2, Institute for International Economic Studies.
- Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, 109(3):659–684.
- Arrow, K. J. and Debreu, G. (1954). Existence of an Equilibrium for a Competitive Economy. *Econometrica*, 22(3):265–290.
- Auclert, A. (2017). Monetary Policy and the Redistribution Channel. Working paper.
- Auclert, A. (2019). Monetary Policy and the Redistribution Channel. *American Economic Review*, 109(6):2333–67.
- Auclert, A., Rognlie, M., and Straub, L. (2020). Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model. Working paper.
- Barattieri, A., Basu, S., and Gottschalk, P. (2014). Some Evidence on the Importance of Sticky Wages. *American Economic Journal: Macroeconomics*, 6(1):70–101.
- Barillas, F. and Fernández-Villaverde, J. (2007). A generalization of the endogenous grid method. *Journal of Economic Dynamics and Control*, 31(8):2698 – 2712.
- Bayer, C., Born, B., and Luetticke, R. (2020). Shocks, Frictions, and Inequality in US Business Cycles. Discussion Paper DP14364, CEPR.
- Bellman, R. (1957). A Markovian decision process. *Journal of mathematics and mechanics*, pages 679–684.
- Bewley, T. (1986). Stationary monetary equilibrium with a continuum of independently fluctuating consumers. *Contributions to mathematical economics in honor of Gérard Debreu*, 79.

- Broer, T., Harbo Hansen, N.-J., Krusell, P., and Öberg, E. (2019). The New Keynesian Transmission Mechanism: A Heterogeneous-Agent Perspective. *The Review of Economic Studies*, 87(1):77–101.
- Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383 – 398.
- Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312 – 320.
- Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, 113(1):1–45.
- Dixit, A. K. and Stiglitz, J. E. (1977). Monopolistic Competition and Optimum Product Diversity. *The American Economic Review*, 67(3):297–308.
- Druehl, J. and Jørgensen, T. H. (2017). A general endogenous grid method for multi-dimensional models with non-convexities and constraints. *Journal of Economic Dynamics and Control*, 74:87 – 107.
- Erceg, C. J., Henderson, D. W., and Levin, A. T. (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics*, 46(2):281 – 313.
- Floden, M. and Lindé, J. (2001). Idiosyncratic risk in the United States and Sweden: Is there a role for government insurance? *Review of Economic dynamics*, 4(2):406–437.
- Galí, J. (2008). Introduction to Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework. In *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Introductory Chapters. Princeton University Press.
- Gornemann, N., Kuester, K., and Nakajima, M. (2016). Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy. International Finance Discussion Papers 1167, Board of Governors of the Federal Reserve System (U.S.).

- Hagedorn, M., Luo, J., Manovskii, I., and Mitman, K. (2019a). Forward guidance. *Journal of Monetary Economics*, 102:1 – 23. Carnegie-Rochester-NYU Conference Series on Public Policy “A Conference Honoring the Contributions of Charles Plosser to Economics” Held at the University of Rochester Simon Business School April 20-21, 2018.
- Hagedorn, M., Manovskii, I., and Mitman, K. (2019b). The Fiscal Multiplier. Working Paper 25571, National Bureau of Economic Research.
- Hintermaier, T. and Koeniger, W. (2010). The method of endogenous gridpoints with occasionally binding constraints among endogenous variables. *Journal of Economic Dynamics and Control*, 34(10):2074 – 2088.
- Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5):953 – 969.
- Huo, Z. and Ríos-Rull, J.-V. (2020). Sticky Wage Models and Labor Supply Constraints. *American Economic Journal: Macroeconomics*, 12(3):284–318.
- Imrohoroglu, A. (1989). Cost of Business Cycles with Indivisibilities and Liquidity Constraints. *Journal of Political Economy*, 97(6):1364–1383.
- Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary Policy According to HANK. *American Economic Review*, 108(3):697–743.
- Kaplan, G. and Violante, G. L. (2018). Microeconomic Heterogeneity and Macroeconomic Shocks. *Journal of Economic Perspectives*, 32(3):167–94.
- Krueger, D., Mitman, K., and Perri, F. (2016). Chapter 11 - Macroeconomics and Household Heterogeneity. volume 2 of *Handbook of Macroeconomics*, pages 843 – 921. Elsevier.
- Krusell, P. and Smith, A. A. (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy*, 106(5).
- Ludwig, A. and Schön, M. (2018). Endogenous Grids in Higher Dimensions: Delaunay Interpolation and Hybrid Methods. *Computational Economics*, 51:463 – 492.

- Luetticke, R. (2020). Transmission of Monetary Policy with Heterogeneity in Household Portfolios. *American Economic Journal: Macroeconomics*, forthcoming.
- MaCurdy, T. E. (1981). An Empirical Model of Labor Supply in a Life-Cycle Setting. *Journal of Political Economy*, 89(6):1059–1085.
- McKay, A., Nakamura, E., and Steinsson, J. (2016). The Power of Forward Guidance Revisited. *American Economic Review*, 106(10):3133–58.
- Pissarides, C. A. (1985). Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages. *The American Economic Review*, 75(4):676–690.
- Rotemberg, J. J. (1982). Monopolistic Price Adjustment and Aggregate Output. *The Review of Economic Studies*, 49(4):517–531.
- Rouwenhorst, G. (1995). ‘Asset Pricing Implications of Equilibrium Business Cycle Models’ Chapter 10, *Frontiers of Business Cycle Research*, T. Cooley.
- Schmitt-Grohé, S. and Uribe, M. (2005). Optimal Fiscal and Monetary Policy in a Medium-Scale Macroeconomic Model. *NBER Macroeconomics Annual*, 20:383–425.
- Smets, F. and Wouters, R. (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, 97(3):586–606.
- Tauchen, G. (1986). Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2):177 – 181.
- White, M. N. (2015). The method of endogenous gridpoints in theory and practice. *Journal of Economic Dynamics and Control*, 60:26 – 41.

1.A Appendix

1.A.1 Derivation of labor demand

Here I lay out some details of the solution to the representative firm's problem (1.2), restated below, in versions 1 and 4 of the model, i.e., when the labor demand is forcing.

$$\min_{n(i)} \left\{ \int_0^1 g^w(i) n(i) di \right\} \quad \text{subject to} \quad \left(\int_0^1 [e(i) n(i)]^{\frac{\delta-1}{\delta}} di \right)^{\frac{\delta}{\delta-1}} \geq N,$$

Using Lagrange's method:

$$\mathcal{L} = \int_0^1 g^w(i) n(i) di - \lambda \left[\left(\int_0^1 [e(i) n(i)]^{\frac{\delta-1}{\delta}} di \right)^{\frac{\delta}{\delta-1}} - N \right],$$

the first-order conditions $\frac{\partial \mathcal{L}}{\partial n(i)} = 0$ are

$$g^w(i) - \lambda \frac{\delta}{\delta-1} \left(\int_0^1 [e(i) n(i)]^{\frac{\delta-1}{\delta}} di \right)^{\frac{\delta}{\delta-1}-1} e(i)^{\frac{\delta-1}{\delta}} \frac{\delta-1}{\delta} n(i)^{\frac{\delta-1}{\delta}-1} = 0.$$

Noting that $\frac{\delta}{\delta-1} - 1 = \frac{1}{\delta-1}$, $\frac{\delta-1}{\delta} - 1 = -\frac{1}{\delta}$, and $\left(\int_0^1 [e(i) n(i)]^{\frac{\delta-1}{\delta}} di \right)^{\frac{\delta}{\delta-1}} = N^{\frac{1}{\delta}}$, the above can be expressed as

$$g^w(i) - \lambda N^{\frac{1}{\delta}} e(i)^{\frac{\delta-1}{\delta}} n(i)^{-\frac{1}{\delta}} = 0.$$

Now, the Lagrange multiplier λ is the cost of one additional unit of aggregate labor, i.e., the shadow wage. By definition (1.4), this is equal to the aggregate wage index W . Applying this and rearranging yields

$$n(i) = \left(\frac{g^w(i)}{W} \right)^{-\delta} N e(i)^{\delta-1},$$

i.e., the labor-demand function (1.3).

1.A.2 Numerical solution to the household's problem

The main challenge to solving the household's problem lies in the fact that there are two endogenous state variables: savings (b'), and the wage (w^h or w^e). The algorithm I develop is a combination of the endogenous-gridpoint method (EGM) by Carroll (2006), and a value-function iteration method (Bellman, 1957).⁴⁷ Below I describe the algorithm in some detail for version 2.⁴⁸

The numerical solution requires a discretization of the state space. The grid for bonds, \mathcal{G}^b , consists of 80 grid points on $[0, 32]$, with an increasing grid space by a factor 1.3. The wage grid, \mathcal{G}^w , consists of 25 grid points on $[0.05, 2.85]$, with an increasing grid space by a factor 1.2. The grid for idiosyncratic productivity, \mathcal{G}^e , consists of 15 log-linearly spaced points on the interval $[0.15, 6.65]$, given by the Rouwenhorst (1995) method.

The rough idea of the solution method is to begin with the case of non-updating households ($f = 0$), and solve it by EGM, obtaining the value for all possible wage levels. Thereafter, move to the case of updating households ($f = 1$), and solve it by choosing the wage that maximizes the value. The steps of the algorithm are outlined here:⁴⁹

1. Guess $g^c(\cdot)$ and $V(\cdot)$.
2. Use the Euler equation (1.19), with $c' = g^c(b', w^h \equiv w_{-1}^h, e', f')$, to calculate $c^*(b', w_{-1}^h, e; 0)$ for all $(b', w_{-1}^h, e) \in \mathcal{G}^b \times \mathcal{G}^w \times \mathcal{G}^e$.⁵⁰

Case $f = 0$ (not able to update the wage), i.e., $g^w(\cdot) \equiv w_{-1}^h$

⁴⁷There are several generalizations of EGM, e.g., Barillas and Fernández-Villaverde (2007) and Druedahl and Jørgensen (2017). However, none of these methods are directly applicable to the model presented here. Barillas and Fernández-Villaverde (2007) fails because it applies to a second control variable which is not also a state variable. Druedahl and Jørgensen (2017) fails because it requires closed-form optimality conditions for all endogenous state variables, which is not the situation in my model. The solution method I use is closest to the hybrid methods suggested in Hintermaier and Koeniger (2010), or Ludwig and Schön (2018).

⁴⁸Solving versions 2 and 3 is more demanding than solving versions 1 and 4, due to the extra choice variable in labor supply. However, the algorithm is easily adjusted for any of the other versions.

⁴⁹The following aggregate prices and quantities are assumed to be known: W , R , N_d , W_d , D and T .

⁵⁰In the numerical exercises, I assume utility to be separable in consumption and leisure, so that $\frac{\partial U}{\partial c}$ is independent of l , see section 1.3. Note also that expectations have to be taken over e' and f' .

3. For all w_{-1}^h and e , calculate the labor demand, $n^*(w_{-1}^h, e)$, from (1.7).
4. Create (equally spaced) grids, $\mathcal{G}^l(w_{-1}^h, e)$, on $[0, n^*(w_{-1}^h, e)]$ for all w_{-1}^h and e .⁵¹
5. For all possible labor supply $l \in \mathcal{G}^l(w_{-1}^h, e)$, use EGM to update $\tilde{g}^c(b \in \mathcal{G}^b, w_{-1}^h, e; f = 0, l \in \mathcal{G}^l(w_{-1}^h, e))$:
 - (a) Calculate the wealth, $b^*(b', w_{-1}^h, e; 0, l)$, that corresponds to $c^*(b', w_{-1}^h, e; 0)$, from the budget constraint (1.16). Note that this is done for all $l \in \mathcal{G}^l(w_{-1}^h, e)$.
 - (b) Interpolate $c^*(b', w_{-1}^h, e; 0)$ from $b^*(b', w_{-1}^h, e, l; 0)$ to \mathcal{G}^b .⁵²
 - (c) For $\{b \in \mathcal{G}^b : b < \min \{b^*(b', w_{-1}^h, e; 0, l)\}\}$, i.e., households that would borrow if allowed, let $b' = 0$. Use this in the budget constraint (1.16) to solve for the corresponding c .
 - (d) Combine the two previous steps to define an auxiliary consumption decision rule $c = \tilde{g}^c(b, w_{-1}^h, e; f = 0, l)$.
 - (e) Calculate the corresponding auxiliary savings decision rule, $\tilde{g}^b(b, w_{-1}^h, e; f = 0, l)$, from the budget constraint (1.16).
6. For each $l \in \mathcal{G}^l$, use the Bellman equation (1.15), with the maximization replaced by $c = \tilde{g}^c(b, w_{-1}^h, e; f = 0, l)$, $w^h = w_{-1}^h$, and $b' = \tilde{g}^b(b, w_{-1}^h, e; f = 0, l)$ to define $\tilde{V}(b, w_{-1}^h, e; f = 0, l)$.⁵³
7. Using interpolation, let⁵⁴

$$V(b, w_{-1}^h, e; 0) = \max_{l \in \mathcal{G}^l(w_{-1}^h, e)} \left\{ \tilde{V}(b, w_{-1}^h, e; 0, l) \right\}$$

$$g^l(b, w_{-1}^h, e; 0) = \arg \max_{l \in \mathcal{G}^l(w_{-1}^h, e)} \left\{ \tilde{V}(b, w_{-1}^h, e; 0, l) \right\}.$$

⁵¹I use 0.1 spacing, or denser if required to reach a minimum of 11 grid points, for $\mathcal{G}^l(\cdot)$.

⁵²I use a piece-wise cubic hermite interpolating polynomial (PCHIP).

⁵³Note that $g^w(\cdot)$ need not be known to calculate the continuation value in the right-hand side of the Bellman equation. For the case of $f' = 0$, it is known that $w^h = w_{-1}^h$; and for the case of $f' = 1$, $V(\cdot)$ is independent of w^h .

⁵⁴I use a cubic spline with not-a-knot end conditions. An involuntary underemployed household ends up in a corner solution in this step, constrained by labor demand.

8. Let the decision rules for consumption and savings be the corresponding choices, i.e.,

$$\begin{aligned} g^c(b, w_{-1}^h, e; 0) &= \tilde{g}^c(b, w_{-1}^h, e; 0, g^l(b, w_{-1}^h, e; 0)) \\ g^b(b, w_{-1}^h, e; 0) &= \tilde{g}^b(b, w_{-1}^h, e; 0, g^l(b, w_{-1}^h, e; 0)). \end{aligned}$$

Case $f = 1$ (able to update the wage)

9. Using interpolation, let⁵⁵

$$\begin{aligned} V(b, w_{-1}^h, e; 1) &= \max_{w_{-1}^h} \left\{ V(b, w_{-1}^h, e; 0) \right\} \\ g^w(b, w_{-1}^h, e; 1) &= \arg \max_{w_{-1}^h} \left\{ V(b, w_{-1}^h, e; 0) \right\}. \end{aligned}$$

10. Evaluate $|V(\cdot) - V_{\text{old}}(\cdot)|$ and $|g^c(\cdot) - g_{\text{old}}^c(\cdot)|$, and if above some tolerance, return to step 2.⁵⁶

This concludes the algorithm, which is robust to some changes of interpolation methods and tolerance levels. Note from step 9 that, all else equal, the following two things are equivalent:

- to enter a non-updating period ($f = 0$) with exactly the wage that would have been chosen if the wage could have been reset,
- to be able to update the wage ($f = 1$),

which is intuitive.

1.A.3 Decision rules

This section explicitly illustrates some of the decision rules governing the households' behavior.

⁵⁵I use a cubic spline with not-a-knot end conditions.

⁵⁶The tolerances I use are 10^{-4} for $V(\cdot)$, and 10^{-6} for $g^c(\cdot)$, for the average over the state space.

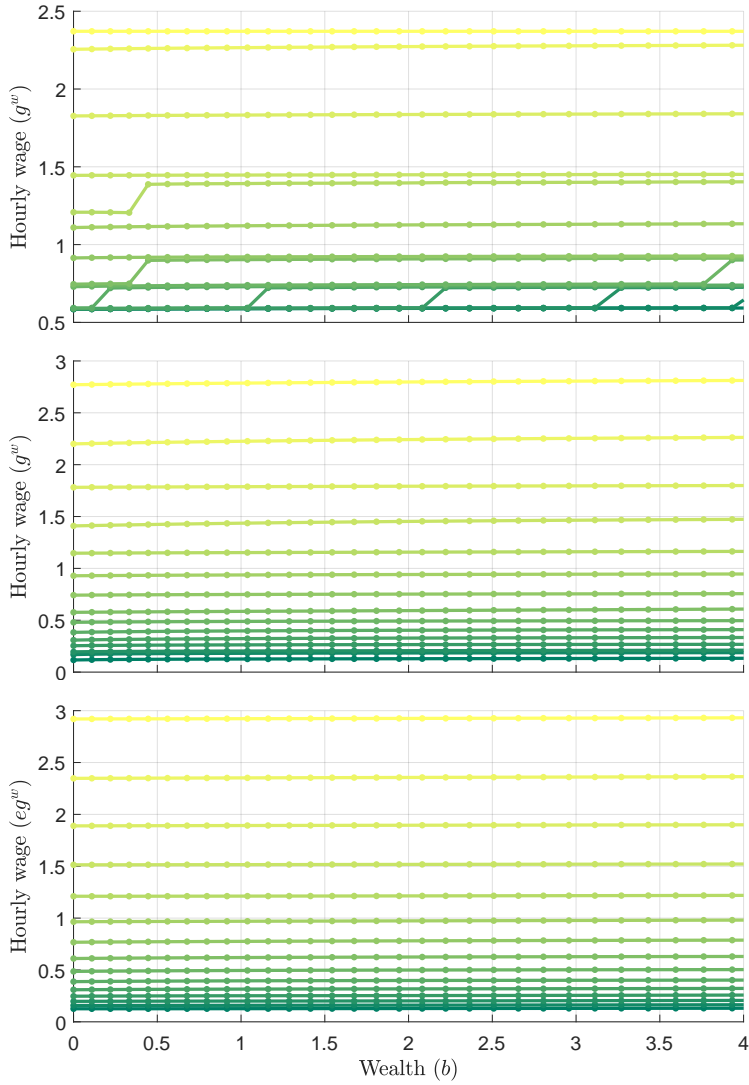


Figure 1.12: Wage decision rules of wage-updating households, $f = 1$. The top row is version 1, the middle row is version 2, and the bottom row is version 3. Line colors indicate level of idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$). Note that for version 3, g^w refers to the effective wage, so for comparison, equation (1.11) is used to transform into hourly wage.




Figure 1.13: Labor-supply decision rules of wage-updating households, $f = 1$. The top row is version 1, the middle row is version 2, and the bottom row is version 3. Line colors indicate the level of idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$).




Figure 1.14: Decision rules for the net savings of households unable to update the wage, $f = 0$. The top row is version 1, the middle row is version 2, and the bottom row is version 3. The columns represent household productivity levels. Line colors indicate the wage, from **lowest** to **highest**.




Figure 1.15: Decision rules for the consumption of households unable to update the wage, $f = 0$. The top row is version 1, the middle row is version 2, and the bottom row is version 3. The columns represent household productivity levels. Line colors indicate the wage, from **lowest** to **highest**.




Figure 1.16: Labor supply-to-demand ratio of wage-updating households, $f = 1$. The top row is version 2, and the bottom row is version 3. Line colors indicate the level of idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$).

1.A.4 Value




Figure 1.17: Value as a function of the state variables, $V(b, w_{-1}^e, e; f)$, of households unable to update their wage, $f = 0$. The top row is version 1, the middle row is version 2, and the bottom row is version 3. The columns represent household wage levels: low wage to the left, average wage in the middle, and high wage to the right. The exact wage levels vary between the three versions.

1.A.5 Time correlations

This section presents a correlation analysis of household-level variables, including eight lags.

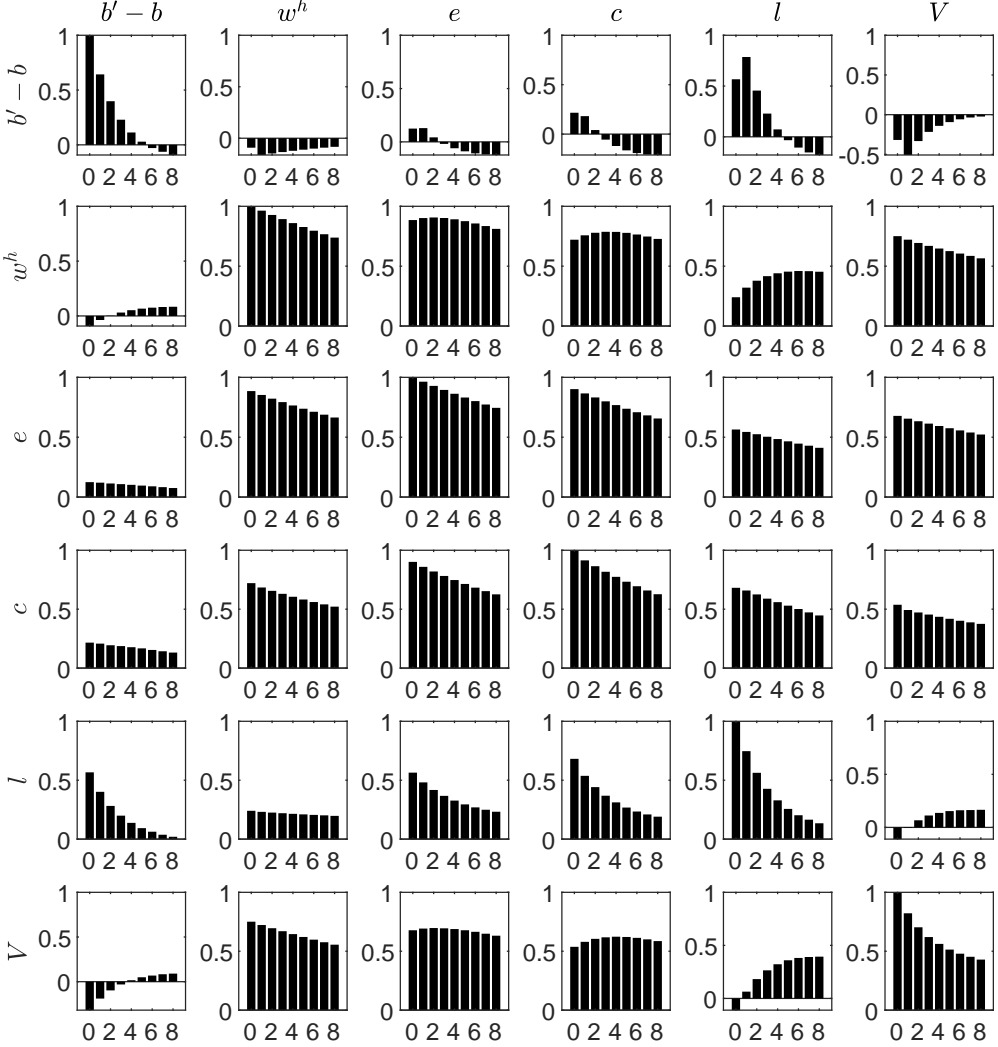


Figure 1.18: Time correlations of individual variables in version 1. Panel rows and columns correspond to labeled household-level variables. The horizontal axis in each panel shows lag/lead periods, while the vertical axis shows the correlation coefficient: $\text{corr}(\text{column}_t, \text{row}_{t+k})$, where k is the value on the horizontal axis. As an example, the panel in row 1 ($b' - b$), column 4 (c) shows that $\text{corr}(c_t, b_{t+7} - b_{t+6}) < 0$, while the panel in row 4 (c), column 1 ($b' - b$) shows that $\text{corr}(b_{t+1} - b, c_{t+6}) > 0$.

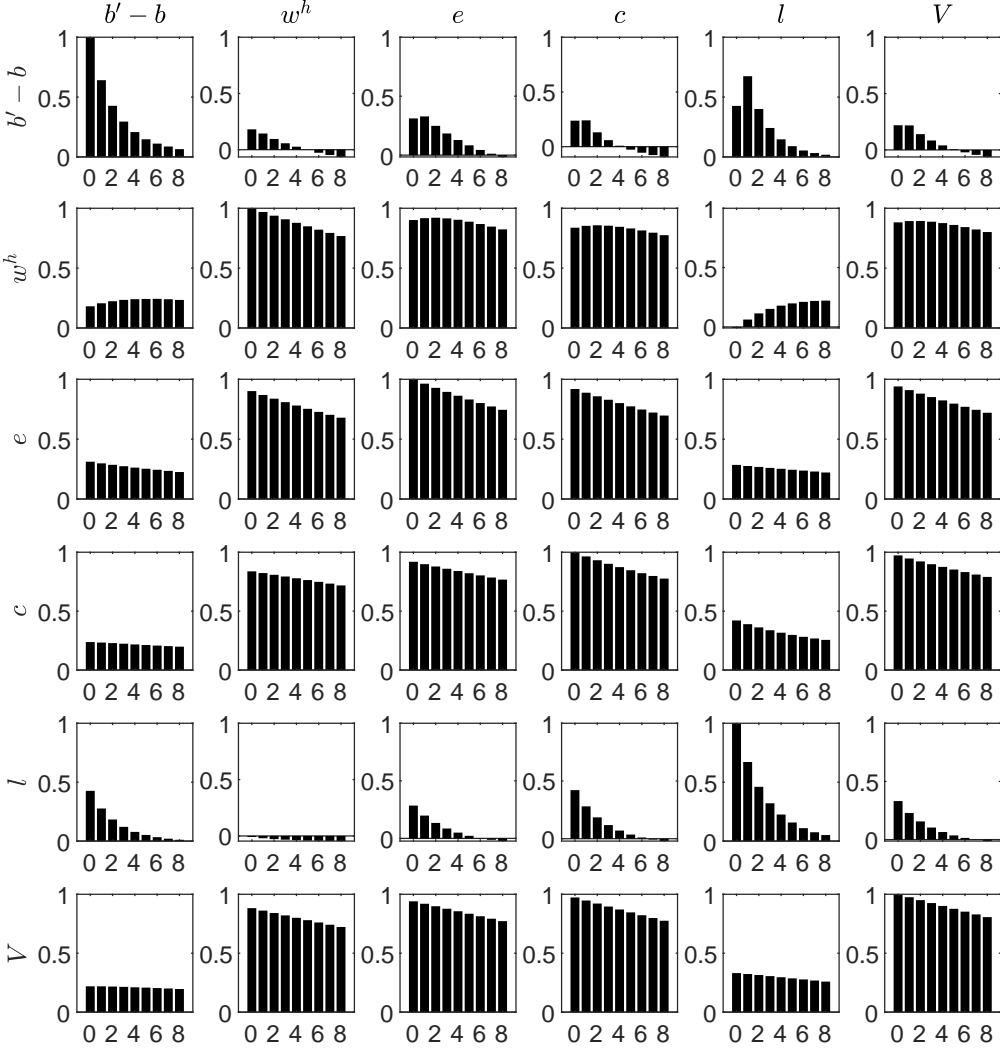


Figure 1.19: Time correlations of individual variables in version 2. Panel rows and columns correspond to labeled household-level variables. The horizontal axis in each panel shows lag/lead periods, while the vertical axis shows the correlation coefficient: $\text{corr}(\text{column}_t, \text{row}_{t+k})$, where k is the value on the horizontal axis. As an example, the panel in row 2 (w^h), column 5 (l) shows that $\text{corr}(l_t, w^h_{t+6}) > 0$, while the panel in row 5 (l), column 2 (w^h) shows that $\text{corr}(w^h_t, l_{t+6}) < 0$.

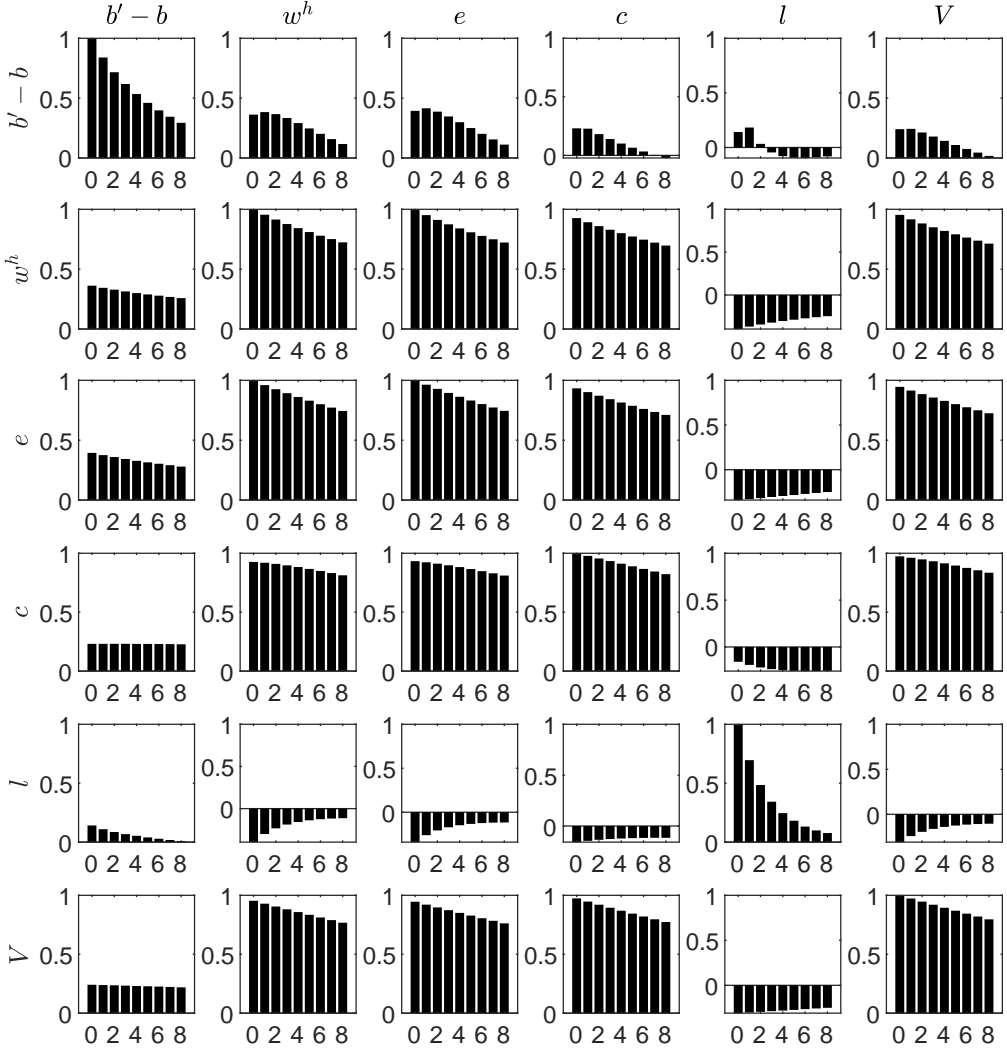


Figure 1.20: Time correlations of individual variables in version 3. Panel rows and columns correspond to labeled household-level variables. The horizontal axis in each panel shows lag/lead periods, while the vertical axis shows the correlation coefficient: $\text{corr}(\text{column}_t, \text{row}_{t+k})$, where k is the value on the horizontal axis. As an example, the panel in row 1 ($b' - b$), column 5 (l) shows that $\text{corr}(l_t, b_{t+5} - b_{t+4}) < 0$, while the panel in row 5 (l), column 1 ($b' - b$) shows that $\text{corr}(b_{t+1} - b, l_{t+4}) > 0$.

Chapter 2

Wage stickiness and household heterogeneity in general equilibrium^{*}

^{*}I would like to thank Susanto Basu, Timo Boppart, Tobias Broer, Richard Foltyn, John Hassler, Karin Kinnerud, Per Krusell, Kurt Mitman, Jonna Olsson, Christina Patterson, Kathrin Schlafmann, David Vestin, and participants in the IIES Macro Group for helpful discussions.

2.1 Introduction

In chapter 1 of this thesis, Åhl (2020), I develop a macroeconomic model with rich household heterogeneity. In particular, I model wage stickiness at the individual-household level, and analyze how that interacts with household behavior. The main focus is on labor-market outcomes at the micro level. In this chapter, I close the model by explicitly modeling the firm side of the economy, and endogenizing the aggregate wage and the interest rate to clear the markets for labor and government bonds.

A general equilibrium allows me to study how the aggregate outcomes are affected by micro-level wage stickiness. How is the precautionary motive to save affected by an underemployment risk, i.e., the risk of working less than desired due to a suppressed labor demand, and what does this mean for the equilibrium interest rate? Does rationing in the labor market, caused by idiosyncratic wage-productivity mismatches, affect the aggregate production?

More precisely, I take the standard HANK model and add two elements: a Calvo (1983) friction to the individual-household wages, and a relaxation of the assumption that demand is forcing in the labor market.¹ In addition to answering questions about aggregate outcomes, this framework is suitable for addressing distributional questions. What does the wealth distribution look like in this economy? How is the underemployment, caused by wage stickiness, distributed among households? How do marginal propensities to consume and work interact with the wage stickiness? And what about welfare? An important purpose of this chapter is also to test how sensitive these results are to variations in key model parameters.

The main mechanism driving my results is that sticky individual wages give rise to a potential mismatch between the wage and the idiosyncratic productivity. If the mismatch is in the direction of a too high wage, the labor demand drops and the household becomes underemployed. If, on the other hand, the mismatch is towards a too low wage, the labor demand can be very high. However, the household will not meet the high demand for her labor services and hence, there is rationing in the labor market. The idiosyncratic demand-supply gap arising from the wage stickiness

¹This is version 2 in Åhl (2020), i.e., the preferred version.

gives rise to an unresponsiveness by some households to a change in labor demand. This is taken into account when the firms solve for optimal labor demand, which matters significantly in a general equilibrium.

My main findings are that the model, with standard parameter values, gives rise to micro distributions that qualitatively match the US data. The risk of underemployment, which is economically significant, amplifies the precautionary motive to save. Compared to the model with flexible wages, this results in a lower equilibrium interest rate. The wage friction also distorts the labor supply, resulting in a lower aggregate labor, and hence a lower output. Another consequence of the dispersed labor supply, mainly caused by underemployment, is a larger earnings inequality. However, the precautionary savings behavior actually results in a more concentrated wealth distribution. In terms of mechanisms, the underemployed tend to be wealth-poor and low-productive, in line with micro data on unemployment.

I show that the findings are robust to variations in the key parameters in the model: higher wage stickiness, less specialization in the labor market, a non-uniform distribution of firm profits, and variations in the process for idiosyncratic productivity. The different parameter assumptions mainly affect the precautionary motive to save, and the results are in line with the intuition.

The structure of the rest of the paper is the following: Immediately below, I comment briefly on the related literature, section 2.2 presents the model in some detail, section 2.3 briefly discusses parameter values, section 2.4 presents data distributions of the variables with a model counterpart, section 2.5 presents and discusses some results and sensitivity exercises, and finally section 2.6 concludes the paper.

Related literature The foundations of the model employed in this chapter are incomplete markets, and household heterogeneity, which was introduced by Bewley (1986), Imrohoroglu (1989), Huggett (1993) and Aiyagari (1994), in combination with wage stickiness, which was introduced in a microfounded macroeconomic model by Erceg et al. (2000). This, together with price stickiness, places this paper in the heterogeneous-agents new-Keynesian literature (HANK, as coined by Kaplan et al. (2018)). Since the model overlaps largely with the one developed in chapter 1 of this thesis, Åhl (2020), I refer to the literature discussion therein for details.

However, what is not covered in Åhl (2020) are aspects related to a general equilibrium and aggregate outcomes. For an overview of the general-equilibrium HANK literature, see Kaplan and Violante (2018). In this chapter, I also discuss how the marginal propensity to consume and supply labor co-varies with other characteristics — a discussion that gained attention with Kaplan and Violante (2014). I also relate to the literature on mechanisms explaining wealth inequality; see Hubmer et al. (2020) for an overview and discussion.

2.2 Model

This section introduces the model used in the analysis of this paper. It belongs to the class of idiosyncratic productivity and incomplete markets models. A novel feature of this model is the rich household heterogeneity resulting from a wage friction effective at the household level.

The model is essentially the same as the preferred version in Åhl (2020), with the main difference that here I analyze a general equilibrium, in which the markets for goods, labor and bonds clear. The main features of the model are incomplete markets, wage stickiness at the household level, and a non-forcing demand in the labor market. On the firm side prices are sticky, thus making the model new-Keynesian.

There are two natural model candidates to which the results can be compared: a similar model with complete markets, and a similar model with fully flexible wages. Those are handled in appendices 2.A.3 and 2.A.4, respectively. The remainder of this section presents details on the agents in the model and how an equilibrium is found.

2.2.1 Households

As in many other macroeconomic models, the households are at the core of the analysis. I first present some assumptions about the environment, and then focus on the households' problem and optimal decisions. Some details are omitted, but are found in Åhl (2020).

Idiosyncratic productivity

Idiosyncratic productivity, denoted e , is the amount of output, as valued by a firm, a household produces while supplying labor in one (unspecified) unit of time. A doubling of e results in twice as much effective labor in every hour worked by the household in question. However, e does not affect the disutility suffered from working. Two households working an equal number of hours, but with different productivity, suffer the same disutility.²

I assume that e is exogenous to the household, and follows an $AR(1)$ process (in logarithm) with stochastic innovations:

$$\log(e_t) = \rho_e \log(e_{t-1}) + \varepsilon_t^e, \quad \varepsilon_t^e \sim N(0, \sigma_e). \quad (2.1)$$

ρ_e is the persistence of the process and ε_t^e is the innovation to the process at time t , which is normally distributed with mean 0 and standard deviation σ_e . Note that although the innovations are idiosyncratic, all households share the same parameters of the process, and are hence ex-ante identical.

The more persistent is the idiosyncratic productivity process, i.e., the higher is ρ_e , the more it influences future earnings. In that sense, the idiosyncratic productivity can be seen as a measure of human capital, as opposed to the financial wealth in government bonds, b , to be introduced below.

Equation (2.1) suggests a continuous support for e . However, in the numerical exercises the stochastic process for e has a discrete support and follows a Markov chain by the method of Rouwenhorst (1995).

Financial market structure

I assume that markets are incomplete, i.e., there is no market for Arrow and Debreu (1954)-type securities. Households can trade in government bonds to insure against bad future idiosyncratic shocks and smooth consumption. However, I also assume a zero-borrowing limit, so negative bond holdings are not feasible.

²This assumption can be relaxed. Moreover, an alternative to idiosyncratic productivity as the source of heterogeneity is to assume stochastic idiosyncratic preferences for leisure. That is not within the scope of this paper, but preliminary results are available upon request.

There is no capital in the economy. Households own the firms, which make profits paid out as dividends, but the firm shares are not traded. To conclude, government bonds is the only traded asset.

Wage friction

Wage setting is individual in this economy, meaning that each household sets her own wage. This is done subject to a Calvo (1983)-type friction, i.e., there is a probability $1 - \theta_w$ each period that she may update her wage. In periods when she cannot update the wage, it remains the same as in the previous period. Note that the probability of not updating the wage, θ_w , is the same for all households at all times, and it is hence independent of the state of the household, and of the time since the last update. One view is that households are subject to idiosyncratic wage-updating shocks that are independent and identically distributed.

With fluctuating idiosyncratic productivity, it is important to distinguish between two possible wages: the compensation per time spent working (regardless of the output), and the compensation per produced output (regardless of the time spent doing so). For a more detailed discussion, see Åhl (2020). In this paper, *wage* refers to compensation per time unit worked, which I refer to as hours worked.

In an economy where the inflation is not constant and zero, it is important to distinguish between the real and the nominal wage.³ In this paper, I limit the analysis to a steady state where I assume inflation to be zero. This means that the real and nominal wages coincide by assumption. However, if one wants to study a response to an aggregate shock, or a steady state with a non-zero inflation, one has to take a stand on which wage is sticky; the nominal or the real wage.

The households' problem

The time horizon is infinite, which is motivated by households being dynastic and hence taking all future generations into consideration, but with discounting. A household has three sources of income: the earnings from labor, the interest on holdings of government bonds, and the dividends

³The same is true for the government bond, which can be considered as nominal or real.

from non-traded ownership of firms. If the interest rate is negative, the household will also receive a transfer from the government (see section 2.2.3), but in that case, her income from government bonds is negative. Her recursive optimization problem is:

$$V(b, w_{-1}, e; f) = \max_{c, w, l, b'} \{U(c, l) + \beta E[V(b', w, e'; f')]\} \quad (2.2)$$

subject to a budget constraint

$$c + \frac{b'}{R} = wl + b + D - T, \quad (2.3)$$

a non-forcing labor-demand function

$$l \leq n^*(w, e), \quad (2.4)$$

and the wage-updating friction

$$w = w_{-1} \quad \text{if } f = 0. \quad (2.5)$$

In general, upper-case Latin variables denote aggregate variables (or functions), while lower-case Latin letters denote household-level variables. $V(\cdot)$ is the value function, and $U(\cdot)$ is the instant-utility function. b is one-period government bonds (bought in the last period, and maturing today), w is the wage, and e is the exogenous idiosyncratic productivity. A prime ($'$) indicates next period, and a minus-one sub index ($_{-1}$) indicates the last period (or, equivalently, the beginning of the period). (b, w_{-1}, e) is the state vector of a household. $f \in \{0, 1\}$ is an indicator of the possibility to update the wage ($f = 1$) or not ($f = 0$). The choice variables are c , consumption, and l , labor supply. R is the gross interest rate on government bonds, D is lump-sum dividends from firm profits, and T is a lump-sum tax to the government (or a transfer from the government if $R < 1$). $n^*(\cdot)$ is a function for the labor demand, which is described in detail in section 2.2.2.⁴

Optimal household behavior is described by decision rules: $c = g^c(b, w_{-1}, e; f)$, $w = g^w(b, w_{-1}, e; f)$, $l = g^l(b, w_{-1}, e; f)$, and

⁴Note that labor supply is not forced to meet the demand at the individual-household level. This is discussed in more detail in Åhl (2020).

$b' = g^b(b, w_{-1}, e; f)$. Note that, first, since a household who cannot update the wage is stuck with the wage from the last period, we trivially have $g^w(b, w_{-1}, e; f = 0) = w_{-1}$ for all (b, w_{-1}, e) . Second, the beginning-of-period wage, w_{-1} , is irrelevant and will have no effect on a household who can update her wage. Hence, the argument w_{-1} is superfluous in the value function and in the decision rules for a household with the opportunity to update her wage in the current period ($f = 1$). However, I will keep it for convenience.

Demand for goods

The consumption variable c represents a basket of differentiated consumption goods, each produced by a specialized firm; see more details in section 2.2.2. Individual consumption goods are substitutable with elasticity ϵ , so that the bundle c is put together as a Dixit and Stiglitz (1977) aggregator:

$$c = \left(\int_0^1 c(i)^{\frac{\epsilon-1}{\epsilon}} di \right)^{\frac{\epsilon}{\epsilon-1}},$$

where the different types of consumption goods have been indexed by $i \in [0, 1]$, and $\epsilon > 0$ is the constant elasticity of substitution between different types of consumption goods. Bundling the different types of consumption goods optimally yields the following demand function for each type:⁵

$$c(i) = \left(\frac{p(i)}{P} \right)^{-\epsilon} c, \quad (2.6)$$

where $c(i)$ is the demand for the type- i consumption good, $p(i)$ is the price of good i , and P is the aggregate consumer price index (CPI) defined in equation (2.15) below. I refer to $\frac{p(i)}{P}$ as the relative price of good i . Note that the demand for any good i is a continuous and decreasing function of the relative price of that good. Note also that the proportion of each good in the consumption basket is independent of the level of consumption, which is known to be empirically wrong but a simplification often used in this class of models.

⁵This is a standard result in the new-Keynesian literature, see, e.g., Galí (2008).

Employment status

I define the employment status of a household as a continuous measure of her actual labor supply relative to the hypothetical labor she would supply if she had the opportunity to reset her wage:

$$\text{Employment status} = \frac{g^l(b, w_{-1}, e; f)}{g^l(b, w_{-1}, e; 1)}. \quad (2.7)$$

Note that for a wage-updating household ($f = 1$), the employment status is 100% by definition.

Although an employment status above one is possible in the model, it is rare and will not be analyzed further. I refer to an employment status below one as *underemployment*, and distinguish between two cases. *Involuntary underemployment* is when the labor supply is kept below the desired level due to the labor demand being low and binding, i.e., $g^l(b, w_{-1}, e; f) = n^*(w, e)$. This tends to occur when the wage is *high* in relation to the idiosyncratic productivity, and can be seen as a proxy for unemployment. *Voluntary underemployment* is when the labor supply is lower than the desired level because this is the optimal choice by the household.⁶ In this case, labor demand is not binding; $g^l(b, w_{-1}, e; f) < n^*(w, e)$. This tends to occur when the wage is *low* in relation to the idiosyncratic productivity, and is somewhat related to a suppressed labor-force participation.

The distribution of households

To solve the model in general equilibrium, aggregate labor supply, savings, and consumption must be known. Hence, also knowledge of how the households are distributed over the state space, (b, w_{-1}, e) , is necessary. Numerically, this is carried out by discretizing each dimension of the state space and keeping track of transitions between the resulting bins, which is further described in section 2.3. I denote the distribution of households over the state space by $\Omega(b, w_{-1}, e)$, which is to be interpreted as the probability density function (PDF) of households, or the probability mass

⁶If given the chance, she would choose a higher wage, which would result in a stronger incentives to work more — i.e., the substitution effect dominates the income effect.

function in the discretized case.

The total mass of households is assumed to be unity, i.e.,

$$\int \Omega(b, w_{-1}, e) \, db \, dw_{-1} \, de = 1.$$

For compactness, I denote the integral of $\Omega(b, w_{-1}, e)$ over all three arguments by \int_{Ω} henceforth. The marginal distributions of wealth (b) and wages (w_{-1}) are endogenous, but the marginal distribution of idiosyncratic productivity (e), $\int \Omega(b, w_{-1}, e) \, db \, dw_{-1}$, is exogenous and given by the ergodic properties of the idiosyncratic productivity process (2.1). Moreover, note that the ability to update the wage is independent of the state. By the law of large numbers, a share $1 - \theta_w$ of the households in *each* state can update their wage in each period, while the remaining θ_w share is unable to do so.

Summary of the households Households are subject to an uninsurable idiosyncratic risk to productivity. They choose how much to consume and how much to save, subject to a borrowing constraint. They can also affect their earnings and labor demand by occasionally resetting their wage, but when doing so, they do not know for how long this wage will be valid. They choose how much to work, capped by labor demand.

Households have preferences for all consumption goods, yielding a demand function that is downward sloping in the price of an individual consumption good. It might be interesting to study how the wage friction affects the actual labor supply relative to the desired supply if the friction is not binding. Household heterogeneity gives rise to a three-dimensional distribution of households over the state space. The distribution is important for the aggregate outcomes of the model in general equilibrium.

2.2.2 Firms

There is a continuum of firms in the economy, each producing one differentiated consumption good. The firms are using their monopolistic power to compete with prices. I let $i \in [0, 1]$ interchangeably denote good i , or the firm producing it. Each firm hires labor services from each household. There is no capital, so labor is the only input to production. Each firm chooses how much labor to demand from each household in order to

minimize the wage-bill costs subject to the current level of output.⁷ When an opportunity occurs, the firm also chooses a price that maximizes the expected future profits as long as that price is valid, i.e., subject to a Calvo (1983)-type price friction.

Production technology

Each household possess a specific labor-skill type. The skill types are symmetric and the labor supply is Dixit and Stiglitz (1977) aggregated, as the product of time worked and the idiosyncratic productivity, to effective labor as:

$$n(i) = \left(\int_0^1 [e(j)l(j, i)]^{\frac{\delta-1}{\delta}} dj \right)^{\frac{\delta}{\delta-1}}, \quad (2.8)$$

where households are indexed by $j \in [0, 1]$, and $l(j, i)$ denotes the hours worked by household j in firm i .⁸ δ is the elasticity of substitution between skill types.

The effective labor is transformed into output by a production function:

$$y(i) = n(i)^{1-\alpha}, \quad (2.9)$$

where $y(i)$ is the output of good i . α is a parameter determining the output elasticity of labor. In equilibrium, the output must equate the consumption for each good, i.e., $y(i) = c(i)$. Given the technology, each firm faces two problems; one static, and one dynamic; which I go through in detail below.

Intratemporal cost minimization

From equation (2.9), it is easy to realize that each output level $y(i)$ of a firm translates into a unique level of effective labor input $n(i)$. However, this effective labor input can be achieved by an infinite number of combinations of labor input from individual households $j \in [0, 1]$. But individual wages

⁷Note that the firm chooses how much labor services to demand from a household, but the household may choose to supply less.

⁸Note that $n^*(\cdot)$ denotes the hypothetical demand (for hours worked), while $n(\cdot)$ denotes the actual effective labor employed.

are exogenous to the firm, so the different combinations differ in wage-bill costs, and one unique combination is strictly cheaper to the firm than all the rest. Hence, a firm's static problem is to minimize the costs subject to an output target, which is easily translated into an effective-labor target.

However, when the labor demand is not forcing, not all households are responsive to a marginal change in labor demand. The supply from households that are unaffected by the demand is exogenous to the firm, which must be taken into consideration in the firm's static problem:

$$\min_{(\mathbf{1}_d(j)n(j))} \left\{ \int_0^1 w(j)n(j)dj \right\} \quad \text{subject to} \quad \left(\int_0^1 [e(j)n(j)]^{\frac{\delta-1}{\delta}} dj \right)^{\frac{\delta}{\delta-1}} \geq N.$$

$\mathbf{1}_d(j) \in \{0, 1\}$ is an indicator of whether the labor demand is binding or not for household j , i.e., whether household j is responsive or not to a marginal change in the labor demand. $w(j)$ is the wage of household j , which is exogenous to the firms.

The solution to this problem is a labor-demand function for household j as

$$n^*(j) = \left(\frac{w(j)}{W_d} \right)^{-\delta} N_d e(j)^{\delta-1}, \quad (2.10)$$

where

$$N_d = \left(N^{\frac{\delta-1}{\delta}} - \int_0^1 ([1 - \mathbf{1}_d(j)]e(j)n(j))^{\frac{\delta-1}{\delta}} dj \right)^{\frac{\delta}{\delta-1}} \quad (2.11)$$

is the aggregate labor needed to meet the output target, on top of the labor supplied by households that are unresponsive to demand changes.

$$W_d = \frac{1}{N_d} \int_0^1 \mathbf{1}_d(j)w(j)n(j)dj \quad (2.12)$$

is an auxiliary wage index that relates only to households responsive to a labor-demand change. The details of the derivation are found in Åhl (2020).

Intertemporal price setting

The price setting is a dynamic problem to the firms because the Calvo (1983) friction makes it possible that a price set today is still valid in the future. A firm that is faced with an opportunity to reset the price of its differentiated good does so in order to maximize the discounted expected stream of future profits. However, only future outcomes where the price set today is still valid need to be considered in today's problem. Formally, firm i 's problem is:

$$\max_{p_t(i)} \left\{ E_t \left[\sum_{k=0}^{\infty} (\beta \theta_p)^k \tilde{d}_{t+k|t}(i) \right] \right\}, \quad (2.13)$$

where $\tilde{d}_{t+k|t}(i) = p_t(i)y_{t+k|t}(i) - W_{t+k}P_{t+k}n_{t+k|t}(i)$ denotes the nominal dividends in period $t+k$, $\forall k \geq 0$, conditional on the price last being updated in period t (implying $p_{t+k}(i) = p_t(i)$). $y_{t+k|t}(i)$ and $n_{t+k|t}(i)$ are production and labor input in firm i in period $t+k$ conditional on the price last being updated in period t . The firms' problem (2.13) is subject to the production technology (2.9), the goods-demand function (2.6), and the goods-specific clearing condition $y_t(i) = c_t(i)$.

$$W = \frac{1}{N} \int_0^1 w(j)n(j)dj \quad (2.14)$$

is an aggregate wage index,

$$N = \int_0^1 n(i)di$$

is the aggregate effective labor employed by all firms,

$$P = \frac{1}{C} \int_0^1 p(i) \left(\int_0^1 c(j,i)dj \right) di \quad (2.15)$$

is the definition of aggregate CPI, and

$$C = \int_0^1 c(j)dj \quad (2.16)$$

is the aggregate consumption. This makes the product WP the aggregate nominal wage index.

The firms are owned, as non-tradable shares, by the households, who care about profit maximization. However, with household heterogeneity the relevant stochastic discount factor (SDF) differs between households. Which is the “correct” SDF to use is up for debate, but I choose to use the time preference β . This choice is common in the literature of heterogeneous households, see Kaplan and Violante (2018) for a brief discussion, and it is sometimes motivated by introducing a risk-neutral mutual fund managing the firms, e.g., in McKay et al. (2016).

The solution to the firms’ problem can be expressed recursively in three equations:⁹

$$J_I = \mathcal{M}J_{II} \quad (2.17)$$

$$J_I = Cp^{*1-\epsilon} + \beta\theta_p E \left[\left(\frac{p^*}{p^{*I}} \right)^{1-\epsilon} \Pi'^\epsilon J'_I \right] \quad (2.18)$$

$$J_{II} = WC^{\frac{1}{1-\alpha}} p^{*- \frac{\epsilon}{1-\alpha}} + \beta\theta_p E \left[\left(\frac{p^*}{p^{*I}} \right)^{- \frac{\epsilon}{1-\alpha}} \Pi'^{1+\frac{\epsilon}{1-\alpha}} J'_{II} \right], \quad (2.19)$$

where J^I and J^{II} are auxiliary variables without a meaningful economic interpretation, p^* is the optimal relative price of those firms that are re-optimizing their prices in the current period (which by symmetry is the same for all such firms), $\Pi \equiv \frac{P}{P_{-1}}$ is gross inflation, and $\mathcal{M} \equiv \frac{\epsilon}{(\epsilon-1)(1-\alpha)}$ is the markup.

Aggregate profits The aggregate profits constitute an important source of income in the households’ budget, see equation (2.3). It is

⁹Details in appendix 2.A.1.

computed by summing all firms' nominal profits:

$$\begin{aligned}
 PD &\equiv \int_0^1 \tilde{d}(i) di \\
 &= \int_0^1 p(i)y(i) - WPn(i) di \\
 &= \underbrace{\int_0^1 p(i)y(i) di}_{\equiv PY} - WP \underbrace{\int_0^1 n(i) di}_{\equiv N} \\
 &= PY - WPN \qquad \Longleftrightarrow \\
 D &= Y - WN.
 \end{aligned} \tag{2.20}$$

I.e., aggregate (real) profits constitute the difference between aggregate output and aggregate (real) wage-bill costs.

The distribution of firms The firms are ex-ante identical, but are heterogeneous by idiosyncratic price-updating opportunities. This gives rise to a distribution of firms over a space of possible goods prices. In steady state, which is the scope of this paper, the distribution degenerates to a single point, since all resetting firms choose the same price at all times. However, when not degenerate, the distribution is important since it distorts the aggregate production; see, e.g., Nakamura et al. (2018) for a discussion.¹⁰ Using individual-goods market clearing and the aggregate market clearing, $Y = C$ (see appendix equation (2.36)), the aggregate

¹⁰The distortion is of second order, and thus disappears when perturbation methods of the first order are used to solve these models.

production function is

$$\begin{aligned}
N &= \int_0^1 n(i) di \\
&\stackrel{\text{by (2.9)}}{=} \int_0^1 y(i)^{\frac{1}{1-\alpha}} di \\
&\stackrel{\text{by (2.6)}}{=} \int_0^1 \left(\left(\frac{p(i)}{P} \right)^{-\epsilon} Y \right)^{\frac{1}{1-\alpha}} di \\
&= Y^{\frac{1}{1-\alpha}} \underbrace{\int_0^1 \left(\frac{p(i)}{P} \right)^{-\frac{\epsilon}{1-\alpha}} di}_{\equiv S^{\frac{1}{1-\alpha}}} \\
&= (SY)^{\frac{1}{1-\alpha}} \quad \Longleftrightarrow \\
SY &= N^{1-\alpha}, \tag{2.21}
\end{aligned}$$

where S is a measure of how much the price dispersion distorts the aggregate output.

The point of the above is that, to solve the model, it is not necessary to keep track of the entire distribution of firms. It is sufficient to keep track of the distortion factor S , which evolves as

$$\begin{aligned}
S^{\frac{1}{1-\alpha}} &= \int_0^1 \left(\frac{p(i)}{P} \right)^{-\frac{\epsilon}{1-\alpha}} di \\
&= \theta_p \int_0^1 \left(\frac{p_{-1}(i)}{P} \right)^{-\frac{\epsilon}{1-\alpha}} di + (1 - \theta_p) \int_0^1 p^{*- \frac{\epsilon}{1-\alpha}} di \\
&= \theta_p \int_0^1 \left(\frac{p_{-1}(i)}{P_{-1}} \frac{1}{\Pi} \right)^{-\frac{\epsilon}{1-\alpha}} di + (1 - \theta_p) p^{*- \frac{\epsilon}{1-\alpha}} \\
&= \theta_p S_{-1}^{\frac{1}{1-\alpha}} \Pi^{\frac{\epsilon}{1-\alpha}} + (1 - \theta_p) p^{*- \frac{\epsilon}{1-\alpha}}. \tag{2.22}
\end{aligned}$$

As mentioned above, and handled in more detail in section 2.2.4, the distortion is zero (i.e., $S = 1$) in a zero-inflation steady state.

Summary of the firms Firms solve two problems. A static problem of finding the optimal skill-type composition, subject to some households being unresponsive to labor demand, yields the labor-demand function that by the households face. A dynamic problem of optimally choosing

the goods price yields a set of optimality conditions that must hold in equilibrium. To solve the model in general equilibrium, we must also calculate the aggregate profits and the distortion to output caused by the price dispersion.

2.2.3 Government

A government provides the market for bonds, which are the traded means of savings. The structure is as simple as possible. The government issues one-period bonds bearing an interest. The interest payments are financed by taxing the households lump-sum, so that the government's budget is balanced in every period. Should the interest rate be negative, it constitutes a revenue to the government, which is then transferred to the households lump-sum, still with the government's budget in balance.

There is no active government policy and no government spending exists in the model. The stock of government bonds is kept fixed, making the government even more passive.¹¹ The balanced budget is:

$$\begin{aligned} T_t &= B_t - \frac{B_{t+1}}{R_t} = \bar{B} \left(1 - \frac{1}{R_t} \right) & \implies \\ T &= \bar{B} \left(1 - \frac{1}{R} \right), \end{aligned} \tag{2.23}$$

where \bar{B} is the fixed stock of bonds.

Models with nominal frictions are often used to study monetary policy, and thus include a central bank conducting monetary policy by setting a short-term nominal interest rate. A common assumption is that the central bank uses a Taylor (1993)-type rule to conduct monetary policy. However, in steady state (as well as in the long run), the equilibrium real rate is determined by fundamental endogenous factors, which cannot be affected by the central bank.¹² Since the analysis conducted here only concerns steady state, i.e., absent aggregate shocks, there is no need to explicitly model a central bank.

¹¹In a zero-inflation steady state, there is no need to distinguish between nominal and real bonds. However, if this assumption is relaxed the distinction needs to be specified more carefully.

¹²The steady-state nominal rate is determined by the real rate and the inflation target by the Fisher (1930) equation: $R^n = R\Pi$, where R^n is the gross nominal equilibrium rate.

2.2.4 A steady-state equilibrium

A steady state in the economy presented above is a situation without fluctuations in any aggregate variables. It is not steady at the household level, in the sense that individual households are subject to shocks to their idiosyncratic productivity, causing them to move around in the distribution over the state space, Ω . However, the distribution is steady in the sense that it is constant over time.

Aggregate variables

Before addressing the equilibrium definition, it is convenient to define some aggregate counterparts to the household-level variables. Equation (2.16) defines the aggregate consumption, i.e., all households' consumption of all goods varieties. Taking decision rules and the distribution of households into account, it can be expressed as

$$C = \int_{\Omega} \theta_w g^c(b, w_{-1}, e; 0) + (1 - \theta_w) g^c(b, w_{-1}, e; 1), \quad (2.24)$$

which is more useful for the numerical solution of the model.

The aggregate savings simply constitute the sum of all households' savings,

$$B' = \int_{\Omega} \theta_w g^b(b, w_{-1}, e, 0) + (1 - \theta_w) g^b(b, w_{-1}, e, 1). \quad (2.25)$$

It may also be convenient, although not necessary, to define the aggregate effective labor supply as

$$L = \left(\int_{\Omega} \theta_w \left[e g^l(b, w_{-1}, e, 0) \right]^{\frac{\delta-1}{\delta}} + (1 - \theta_w) \left[e g^l(b, w_{-1}, e, 1) \right]^{\frac{\delta-1}{\delta}} \right)^{\frac{\delta}{\delta-1}}. \quad (2.26)$$

The aggregate wage index, as suggested by equation (2.14), is defined

as the total wage-bill costs divided by the aggregate effective labor,

$$W = \frac{1}{N} \left(\int_{\Omega} \theta_w g^w(b, w_{-1}, e; 0) g^l(b, w_{-1}, e; 0) \right. \quad (2.27) \\ \left. + (1 - \theta_w) g^w(b, w_{-1}, e; 1) g^l(b, w_{-1}, e; 1) \right)$$

The auxiliary aggregate labor demand and the wage index for responsive households are, when expressed in terms of decision rules and the distribution of households,

$$N_d = \left(N^{\frac{\delta-1}{\delta}} - \int_{\Omega} \left(\theta_w (1 - \mathbf{1}_d(b, w_{-1}, e; 0)) \left[e g^l(b, w_{-1}, e; 0) \right]^{\frac{\delta-1}{\delta}} \right. \right. \quad (2.28) \\ \left. \left. + (1 - \theta_w) (1 - \mathbf{1}_d(b, w_{-1}, e; 1)) \left[e g^l(b, w_{-1}, e; 1) \right]^{\frac{\delta-1}{\delta}} \right) \right)^{\frac{\delta}{\delta-1}},$$

$$W_d = \frac{1}{N_d} \int_{\Omega} \left(\theta_w \mathbf{1}_d(b, w_{-1}, e; 0) g^w(b, w_{-1}, e; 0) g^l(b, w_{-1}, e; 0) \right. \quad (2.29) \\ \left. + (1 - \theta_w) \mathbf{1}_d(b, w_{-1}, e; 1) g^w(b, w_{-1}, e; 1) g^l(b, w_{-1}, e; 1) \right),$$

respectively, c.f. equations (2.11) and (2.12). Note that the labor-demand function (2.10) can be expressed as a function of the state vector,

$$n^*(b, w_{-1}, e; f) = \left(\frac{g^w(b, w_{-1}, e; f)}{W_d} \right)^{-\delta} N_d e^{\delta-1}, \quad (2.30)$$

and formally,

$$\mathbf{1}_d(b, w_{-1}, e; f) = \begin{cases} 1 & \text{if } g^l(b, w_{-1}, e; f) = n^*(b, w_{-1}, e; f) \\ 0 & \text{if } g^l(b, w_{-1}, e; f) < n^*(b, w_{-1}, e; f) \end{cases} \quad (2.31)$$

describes the indicator for binding labor demand (i.e., a responsive household) as a function of the state vector.

Analytical steps to finding a steady-state equilibrium

The assumption of a steady state is not necessary to define an equilibrium in the economy, but simplifies the analysis somewhat. First, I assume that a credible zero-inflation target is in place, so $\Pi = 1$. The appendix

equation (2.35) then trivially gives that $p^* = 1$, meaning that all firms set the same price, and that price coincides with the aggregate price index. With all firms setting the same price at all times, the price dispersion is trivially zero, and hence the distortion to the output is zero. This can be verified in equations (2.22) and (2.21), yielding $S = 1$, and thus

$$Y = N^{1-\alpha}. \quad (2.32)$$

Using $p^* = \Pi = 1$ in the steady-state version of the equilibrium conditions for the firms' price-setting problem yields, for (2.18)

$$J_I = \frac{C}{1 - \beta\theta_w},$$

and for (2.19)

$$J_{II} = \frac{\mathcal{M}WC^{\frac{1}{1-\alpha}}}{1 - \beta\theta_w}.$$

Combining these in (2.17) gives the steady-state equilibrium condition for the firms' dynamic problem:

$$\mathcal{M}W = C^{-\frac{\alpha}{1-\alpha}}. \quad (2.33)$$

Equilibrium definition

A recursive competitive equilibrium is given by

- Decision rules and value functions: $g^c(b, w_{-1}, e; f)$, $g^b(b, w_{-1}, e; f)$, $g^w(b, w_{-1}, e; f)$, $g^l(b, w_{-1}, e; f)$, and $V(b, w_{-1}, e; f)$
- Labor-demand functions $n^*(b, w_{-1}, e; f)$ and indicator functions $\mathbf{1}_d(b, w_{-1}, e; f)$
- A distribution $\Omega(b, w_{-1}, e)$ over the state space
- Aggregate quantities: $R, W, C, Y, N, D, T, N_d, W_d$

such that

1. The decision rules and value functions solve the households' problem (2.2).

2. For all relevant Borel sets \mathcal{B} and \mathcal{W} ,

$$\Omega(\mathcal{B}, \mathcal{W}, e') = \sum_e Pr(e'|e) \left[\theta_w \int_{\left\{ b, w_{-1}: \begin{matrix} g^b(b, w_{-1}, e, 0) \in \mathcal{B} \\ g^w(b, w_{-1}, e, 0) \in \mathcal{W} \end{matrix} \right\}} \Omega(b, w_{-1}, e) db dw_{-1} + \right. \\ \left. (1 - \theta_w) \int_{\left\{ b, w_{-1}: \begin{matrix} g^b(b, w_{-1}, e, 1) \in \mathcal{B} \\ g^w(b, w_{-1}, e, 1) \in \mathcal{W} \end{matrix} \right\}} \Omega(b, w_{-1}, e) db dw_{-1} \right],$$

i.e., the distribution of households generates itself through the policy functions and the exogenous idiosyncratic process.

3. The definitions (2.28), (2.29), (2.30), and (2.31); the computational equilibrium conditions (2.24), and (2.27); and the analytical equilibrium conditions (2.20), (2.23), (2.32), (2.33), and (2.36) are all satisfied.

Finding the equilibrium numerically

The equilibrium described above cannot be found analytically, unless some simplifying and often unrealistic assumptions are made; see, e.g., Krusell et al. (2011) and Ravn and Sterk (2017). Instead, numerical methods have to be employed. This section outlines the algorithm I use to find market-clearing prices in the steady state.¹³

1. Guess R , W , N_d and W_d .
2. Calculate T from (2.23), N from (2.34), Y from (2.32), and D from (2.20).
3. Solve the households' problem (2.2) for all relevant b , w_{-1} , and e ; and $f \in \{0, 1\}$.

¹³Note that equations (2.32), (2.33), and (2.36) can be combined into

$$N = (\mathcal{M}W)^{-\frac{1}{\alpha}}. \tag{2.34}$$

4. Use $g^b(\cdot)$ and $g^w(\cdot)$ to calculate the stationary Ω .
5. Use $g^l(\cdot)$, $g^w(\cdot)$, and $\mathbf{1}_d(\cdot)$ to update N_d from (2.28) and W_d from (2.29).
6. Use $g^b(\cdot)$ to calculate B' from (2.25), $g^c(\cdot)$ to calculate C from (2.24), $g^l(\cdot)$ to calculate L from (2.26), and $g^w(\cdot)$ to calculate \widehat{W} from (2.27).
7. Evaluate the following market-clearing conditions: $B' \stackrel{?}{=} \overline{B}$, $C \stackrel{?}{=} Y$, and $L \stackrel{?}{=} N$. Also evaluate if the implied aggregate wage is consistent with the guess: $\widehat{W} \stackrel{?}{=} W$.
8. If any of the criteria does not hold (to a tolerance), use net bond supply ($B' - \overline{B}$) to update R , and update W as a convex combination of \widehat{W} and the old guess.¹⁴ Return to step 2.

Step 3 above is arguably the most difficult one. The algorithm used to solve the households' problem numerically is a combination of the endogenous gridpoint method by Carroll (2006) and a value-function iteration method (Bellman, 1957). The details are found in Åhl (2020).

2.3 Choice of parameter values

To address quantitative questions with an economic model, one has to be careful about how the model parameters are calibrated to match features of the real-world data. This paper uses a quantitative model, but does not really address any quantitative questions directly. Instead, the aim of this paper is to document mechanisms in a type of model which is new in the literature. In that sense, this analysis gives more of a guidance to future quantitative work than it gives answers to quantitative questions. Hence, the model is not really calibrated to match any features of the data, but rather I choose parameter values that are widely accepted in

¹⁴A market is considered to clear if the absolute surplus is smaller than 10^{-4} for 16 consecutive iterations. The following updating rules are used: $\widehat{R} = R^{old} \left(1 - \frac{B' - \overline{B}}{200}\right)$, $R^{new} = .4\widehat{R} + .6R^{old}$, and $W^{new} = .4\widehat{W} + .6W^{old}$. Admittedly, the tolerance and the updating rules are arbitrary. However, they mainly affect the speed of convergence to a steady-state equilibrium. The tolerance can be set tighter and the updating rules can be changed somewhat without affecting the results.

the existing literature. This section, nevertheless, reports and comments on the chosen parameter values.

In the numerical exercises that follow, I need to be explicit on the utility function used. The choice is a MaCurdy (1981) utility function, separable in consumption and labor effort, and standard in the new-Keynesian literature:

$$U(c, l) = \frac{c^{1-\sigma} - 1}{1 - \sigma} - \frac{l^{1+\varphi}}{1 + \varphi}.$$

Note that, despite being unrealistic, there is no upper limit to labor supply (although there is an increasing marginal disutility). This might be a poor choice in a more quantitatively interesting setting.

Table 2.1 shows the parameter values that are used in the quantitative exercises of section 2.5. As stated above, these are fairly standard in the existing literature, but some comments follow.

Table 2.1: Parameter values

<i>Param.</i>	<i>Value</i>	<i>Interpretation</i>	<i>Comment</i>
β	0.99	Discount factor	Quarterly, standard
δ	6	Subst. elast., skill types	Galí (2008), chapter 6
θ_w	0.75	$1 - Pr(\text{update wage})$	Galí (2008), chapter 6
ϵ	5	Subst. elast., goods	20% markup, McKay et al. (2016)
α	1/3	Output elast. of labor	Standard value
θ_p	0.75	$1 - Pr(\text{update price})$	Galí (2008), chapters 6 and 3
<i>Idios. prod.</i>			
ρ_e	0.966	Persistence, idios. prod.	McKay et al. (2016)
σ_e	0.13	St. dev., idios. prod. innov.	McKay et al. (2016)
<i>Preferences</i>			
σ	2	Risk aversion	McKay et al. (2016)
φ	2	Inverse Frisch elasticity	McKay et al. (2016)
<i>Steady state</i>			
\bar{B}	1	Government debt	28% of annual GDP
Π	1	Gross inflation target	Zero net inflation

The time-discount factor $\beta = 0.99$ means that a household values utility one period into the future one percent less than utility now. One

period is to be interpreted as one quarter of a year. $\theta_w = 3/4$ corresponds to, on average, one year between the wage updates, which is in line with the literature.¹⁵

The parameters of the process for idiosyncratic productivity are taken from McKay et al. (2016). The stationary ergodic distribution is illustrated in figure 2.1. See also Åhl (2020) for more comments. The preference parameters for risk aversion (σ) and wage sensitivity (φ , inverse Frisch elasticity) are standard in the literature.

Finally, I have chosen a stock of government bonds that corresponds to slightly more than one quarter of annual GDP. Since fiscal policy is passive in the model economy, this should be viewed as means of liquidity rather than the level of indebtedness of the government. As a comparison, the US money supply according to the measures M1 and M2 have, on average, been 13% and 55% of GDP since 1980.¹⁶ Nevertheless, my choice for the stock of government bonds might be on the low side of what is realistic. For example, McKay et al. (2016) choose liquidity to be 140% of annual GDP. The qualitative results are robust to a higher stock of government bonds.

¹⁵See, e.g., Barattieri et al. (2014).

¹⁶Source: Federal Reserve Bank of St. Louis.

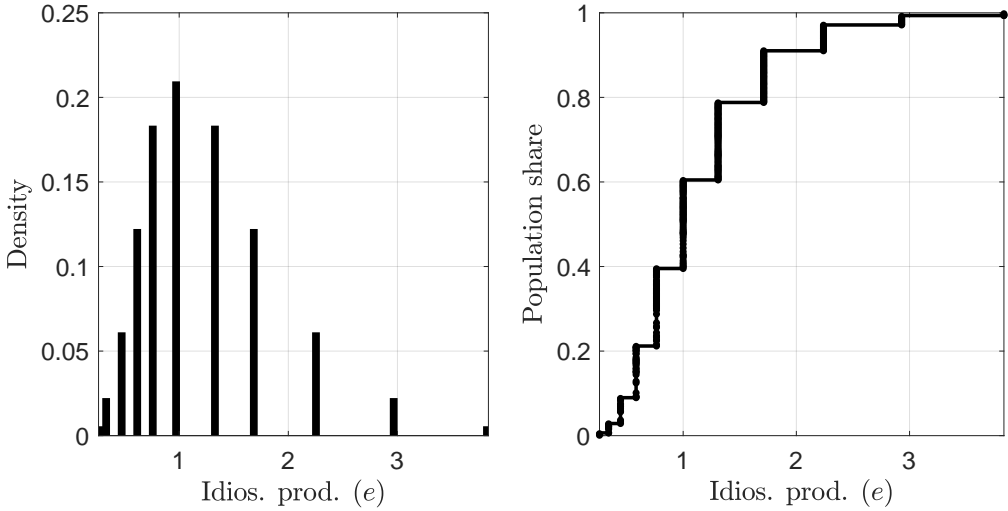


Figure 2.1: The ergodic distribution of the idiosyncratic productivity, e . Left is the probability mass function approximated by a histogram. Right is the corresponding cumulative distribution function. Limits of the horizontal axes are set to exclude at most 0.5 percent of households at each tail of the distribution.

2.4 Micro-level data (PSID 2017)

Macroeconomic models always contain simplifying and more or less unrealistic assumptions. Nevertheless, the models can be useful for answering quantitative questions. When doing so, it is important that the model can capture some moments or features of the real-world data well. The objective of this paper, however, is not to answer specific quantitative questions, but rather to highlight some model features and assumptions that might be important to include, depending on which questions one wants to answer. Hence, a good quantitative replication of the data is not necessary. However, if the mechanisms of the model make sense, the model might be able to *qualitatively* reenact some features of the data. Therefore, this section presents and briefly discusses some cross-sectional US household data, which has a counterpart in the model.

The data I use is the University of Michigan's *Panel Study of Income Dynamics* (PSID), version 2017. The PSID is a biennial survey covering almost 10,000 representative households, asking a broad range of questions

of economic relevance. I do not make use of the panel dimension of the survey, but look at the cross section of the latest available version, which is from 2017. I focus on the distributions for wealth, wages, labor supply, earnings, and consumption, since they have clear counterparts in the model.

I exclude households that are neither salaried nor paid by the hour, leaving a total of 5,101 households. Among the excluded households are many unemployed, where a wage cannot be computed. Below I comment on each variable separately.




Figure 2.2: US distribution of wealth, in thousand USD. Left is the probability density function approximated by a histogram. Right is the corresponding cumulative distribution function. The limits of the horizontal axes are set to exclude at most 0.5 percent of households at each tail of the distribution.

Figure 2.2 shows wealth excluding equity, net of debt. Excluding equity makes a fairer comparison with the model, where firm shares are not traded and not included in the wealth. More than half of the households in the data are bunched at, or very close to, a zero net wealth. In the data, some households have a negative net wealth, which is not possible in the model, where the zero-borrowing constraint prevents debt. The wealth inequality in the data is large, and a few households hold many times the median wealth, thus making the distribution highly skewed.

Figure 2.3 shows the distribution of hourly wages, where tips and commissions have been excluded because of measuring issues. The legislated federal minimum wage was USD 7.25 in 2017, which is hinted in the figure, but varies between states. The mode wage is not far above the federal minimum, but many households have a wage far above these




Figure 2.3: US distribution of hourly wages. Left is the probability density function approximated by a histogram. Right is the corresponding cumulative distribution function. The limits of the horizontal axes are set to exclude at most 0.5 percent of households at each tail of the distribution.

levels, making the distribution heavily skewed.




Figure 2.4: US distribution of hours worked in one year. Left is the probability density function approximated by a histogram. Right is the corresponding cumulative distribution function. The limits of the horizontal axes are set to exclude at most 0.5 percent of households at each tail of the distribution.

Figure 2.4 shows the number of hours worked in 2016, the year before the survey was conducted, calculated as the product of the number of employed weeks and the average hours worked per week.¹⁷ There is a clear spike around 2,000 hours, corresponding to around 40 hours per week, i.e., full-time work. It is also quite common to work substantially

¹⁷The fact that some of the PSID variables of interest refer to the previous year is not optimal, but still presents the best option.

more than full time, with some households actually supplying more than twice that of full-time labor. The figure also shows numerous part-time workers, spread out between zero and full time.




Figure 2.5: US distribution of earnings, in thousand USD. Left is the probability density function approximated by a histogram. Right is the corresponding cumulative distribution function. The limits of the horizontal axes are set to exclude at most 0.5 percent of households at each tail of the distribution.

Figure 2.5 shows pre-tax earnings in 2016. This is a separate variable in the data, while in the model, it is calculated as the product of the wage and the hours worked. The distribution is positively skewed, with some households very close to zero earnings. Since the question concerns 2016, and households unemployed in 2017 have been excluded, the zero earners could be partly explained by unemployment in 2016.

Figure 2.6 shows the consumption expenditures excluding mortgage costs. Excluding mortgages is motivated by these rather being a financial cost — a reasoning that could be questioned. The distribution is similar to that for earnings, with the difference that no households have zero consumption expenditures, and that the highest earnings are higher than the highest consumption expenditures, suggesting that high-earners are net savers. These two observations are consistent with the model behavior, which is further discussed in section 2.5.6.

A natural next step in the data analysis would be to look at interactions between the variables and compare these with the model counterparts. Since the model is not calibrated, I merely point this possibility out and do not pursue it here. Another way to proceed with the data analysis would be to make use of the panel structure and look at auto-correlations




Figure 2.6: US distribution of consumption expenditures, in thousand USD. Left is the probability density function approximated by a histogram. Right is the corresponding cumulative distribution function. The limits of the horizontal axes are set to exclude at most 0.5 percent of households at each tail of the distribution.

and time correlations, which could be compared to the model counterparts; see, e.g., figure 2.38 in the appendix.

2.5 Results

This section presents the main results of this paper. The focus is on the mechanisms driving the households' behavior and outcomes. Section 2.5.7 presents sensitivity analysis with respect to parameter values and assumptions.

2.5.1 The distribution of households, Ω

The households in the economy have three state variables: wealth (b), (beginning-of-period) wage (w_{-1}), and idiosyncratic productivity (e). The wealth and the wage are endogenous, in the sense that households influence what to bring to the next period, while the dynamics of productivity is governed by an exogenous stochastic process. In steady state, an individual household moves around in the state space, but with a continuum of households, the law of large numbers ensures that the distribution of households remains fixed. What this distribution looks like is an essential outcome of the model, worth investigating further.

The joint distribution over wealth and wages is shown in the bottom-




Figure 2.7: Steady-state distribution of households, $\Omega(b, w_{-1}, e)$. The bottom-left panel illustrates the 2-dimensional density over wealth (b) and beginning-of-period wage (w_{-1}), where the dot size represents the concentration of households. The top-left panel is the marginal density of wealth, and the lower-right panel is the marginal density of hourly wage. Colors represent idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$). Axes limits are set to exclude at most 0.5 percent of households at each tail of each dimension.

left panel of figure 2.7, where the productivity dimension is collapsed but still visible as shades of colors. The figure also shows the density (or more correctly, the mass) functions of the marginal distributions for wealth and wages, approximated by histograms. These marginal distributions are endogenous outcomes of the model, while the marginal distribution of idiosyncratic productivity is exogenous (and shown in figure 2.1).

A number of things can be learned from figure 2.7. The top-left panel shows that more than one third of the households are at, or very close to, the borrowing limit. Moreover, the wealth distribution has a fat right tail. Further, and likely not surprising, high-productive households tend to be wealthier on average, although also represented close to the borrowing

limit.

The bottom-right panel, with axes swapped from the norm to better match the bottom-left panel, shows a positively skewed distribution of hourly wages, with more productive households on average having a higher wage — which is also not surprising. It also reveals bunching around certain wage levels, which should be seen as a numerical issue rather than the distribution actually being multi-modal. Each bunching corresponds to a feasible productivity level, so increasing the number of possible productivity levels gives a smoother-looking distribution, but comes at a computational cost.

Finally, although perhaps cumbersome to read, the bottom-left panel shows how all three dimensions of the distribution interact. There is a clear tendency that wealthier households also have higher wages. The three state variables are positively correlated with each other, which it is non-controversial to state is in line with what we expect to find in the data, although the idiosyncratic productivity is difficult to observe.

Idiosyncratic productivity The source of heterogeneity in this economy is the exogenous stochastic process governing the idiosyncratic productivity, see equation (2.1), together with the wage-updating shocks. Hence, an interesting starting point for the analysis is the interaction of the idiosyncratic productivity (e), with the endogenous state variables wage (w) and wealth (b), respectively, which is shown in figure 2.8.

The top panel shows a clear positive relationship between the idiosyncratic productivity and the (hourly) wage.¹⁸ The relationship is distorted by two features: the fact that the optimal wage is also affected by the wealth of a household, which is barely visible in the dispersion of the red dots ($f = 1$); and the wage friction preventing updates in response to a change in productivity, illustrated by the dispersion of the blue dots ($f = 0$).¹⁹ The latter has by far the largest impact in this figure. Loosely speaking, the red dots in the figure hint at a line along which the wage perfectly matches the productivity.²⁰ The further above this line a household is, the more is labor demand suppressed by the mismatch, leading

¹⁸Productivity and the *effective* wage are *negatively* correlated, which is not shown in the figure.

¹⁹The latter is not present if the wages are fully flexible, see appendix figure 2.25.

²⁰Conditional on wealth, which tends to have a very small impact on the wage choice.




Figure 2.8: Joint distribution of idiosyncratic productivity and wage (top panel), and idiosyncratic productivity and wealth (bottom panel). Dot size represents the concentration of households. Colors represent households who **cannot update the wage** ($f = 0$), and households who **can update the wage** ($f = 1$). Axes limits have been chosen to exclude at most 0.5 percent of total households at each tail of each dimension.

to involuntary underemployment. The potential involuntary underemployment creates an asymmetry, and we find more households below the perfect-match line than above. The discretization of idiosyncratic productivity is visible as levels along which all households align as vertical lines. The figure also reveals how these potential productivity levels transmit into bunching in the wage dimension as well, which can be seen in the

lower right panel of figure 2.7.

The bottom panel of figure 2.8 shows a positive relationship between the productivity and the wealth, which is hardly surprising.²¹ The explanation is that a high productivity allows a high wage, without labor demand being suppressed. This, in turn, leads to high earnings, which are partly used to build up wealth, mainly for consumption-smoothing reasons. A wealthy and low-productive household will quickly eat out of her wealth while waiting for her productivity to hike. However, there are some, although few, high-productive households with very little wealth. These are mainly households whose productivity has risen since the last opportunity to reset the wage. They are stuck with a low wage, at which they will not bother to work much. Hence, their earnings remain low, leaving little room to build financial wealth. However, their human capital is quite high, due to the persistence of the productivity process. The figure also reveals that the difference between updating and non-updating households is less pronounced for the wealth than for the wage. However, a larger share of the non-updating households are represented close to the borrowing limit, which is mainly visible for the higher productivity levels in the lower panel of the figure, but true for all productivity levels. These are mainly households that are involuntary underemployed due to a recent decrease in productivity, and are hence represented above the invisible perfect-match line in the top panel.

2.5.2 Aggregate outcomes

Before digging deeper into the mechanisms at the household level, it might be helpful to look at the equilibrium outcomes for some aggregate variables. Table 2.2 provides results. The baseline model, with incomplete markets and sticky wages, is compared with two natural candidates: an identical model except for wages being flexible ($f = 1$ with certainty for all households at all times), and an identical model with complete markets for Arrow and Debreu (1954) securities, fully insuring against idiosyncratic shocks. The latter is equivalent to a model with a representative household with a fixed productivity (RANK).

The equilibrium interest rate differs substantially between the three

²¹Figure 2.26 in the appendix shows that qualitatively, the relationship between the idiosyncratic productivity and the wealth is similar when the wages are flexible.

Table 2.2: Aggregate outcomes in steady state

	Incomplete markets		Complete markets
	Sticky wages	Flexible wages	Sticky wages
Real rate ($R - 1$)	-.88%	-.07%	1.01%
Wage (W)	.57	.56	.57
Output (C, Y)	.88	.91	.86
Hours worked (L, N)	.82	.87	.80
Firm profits (D)	.41	.42	.40
Underemployment	8.5%	—	—
Involuntary	10.6%	—	—
Voluntary	-2.2%	—	—
On labor demand	24.6%	100.0%	100.0%
On borrowing limit	25.5%	50.4%	—
<i>Gini coefficient</i>			
Wealth	.67	.79	—
Wage	.23	.23	—
Productivity	.27	.27	—
Consumption	.12	.11	—
Labor supply	.15	.02	—
Earnings	.31	.24	—
Income	.16	.13	—

models.²² In the RANK model, the household is aware that there is a zero probability that the borrowing limit will be binding and hence, the precautionary motive for savings is not present. The only remaining reason to save is to smooth consumption and thus, the real return on savings completely offsets the myopic nature caused by discounting of the future utility. Hence, the real rate is such that $\beta R = 1$. With incomplete markets, however, the uninsurable idiosyncratic shock creates reasons to save in order to avoid being bound by the borrowing limit in the event of bad future shocks. This motive pushes up the willingness to save, especially for households close to the borrowing limit, compared to what is motivated by pure consumption-smoothing reasons. For the bond market to clear, the equilibrium real interest rate is lower. The precautionary motive is

²²Table 2.2 shows the net quarterly rate. Calculated in annualized terms, the differences are more extreme.

even stronger when wages are sticky rather than flexible, because of the risk of losing almost all earnings due to involuntary underemployment in response to bad idiosyncratic shocks. As a result, the equilibrium rate is lower with sticky wages than with flexible wages.

For the rest of the outcomes, the differences are overall rather small.²³ A more dispersed labor supply, in combination with non-forcing labor demand, causes aggregate labor input, and hence aggregate output, to be slightly lower when the wages are sticky than when they are flexible. Underemployment, defined as $1 - \textit{Employment status}$, see equation (2.7), is only available in the baseline model. Involuntary underemployment can be viewed as a proxy for unemployment, and is close to 10% in the aggregate, which is not very far from the average unemployment rate in the US. Voluntary underemployment is actually negative, but small, meaning that the tendency to work more than desired is stronger than the tendency to work less than desired due to the wage friction among households with non-binding labor demand. Note also that around three quarters of the households choose to supply labor strictly below the demand, hinting that a non-forcing labor demand is important to incorporate in this type of model.

In the baseline model, one quarter of the households are bound by the borrowing limit, to be compared to half of the households in the model with flexible wages. Potential involuntary underemployment makes it more harmful to be at the borrowing limit when the wages are sticky than when they are flexible. The perhaps surprisingly high share of households at the borrowing limit is explained by earnings from labor constituting a relatively small share of the income for low-productive households. Lump-sum dividends from the non-traded ownership of firms work as a sort of insurance, guaranteeing a certain level of consumption also with zero wealth and very bad idiosyncratic shocks. This is true both with sticky and flexible wages, as the firm dividends are almost 90% as large as the economy-wide average earnings.

Inequality, measured by the Gini coefficient, is shown in the bottom part of table 2.2 for a number of variables. It is irrelevant with complete

²³One needs to take into account that with incomplete markets, average idiosyncratic productivity is 1.13, while the productivity of the representative household is 1 with complete markets. This mechanically makes aggregate effective labor lower, and hence also output and profits.

markets, but comparing sticky with flexible wages is possible. Wealth inequality is lower with sticky wages, due to the stronger precautionary motives keeping more households away from the borrowing limit. However, the wage stickiness creates mismatches between productivity and wages, and hence makes hours worked vary a great deal more, e.g., involuntary underemployment, between households than what is the case with flexible wages. This also makes the earnings inequality higher with sticky wages.

In summary, the precautionary motives to save differ considerably between the three models that are compared, causing the equilibrium real interest rates to be quite different. Other aggregate outcomes are similar. The arguably most interesting analysis is at the households micro level, and yet remains to be discussed.

2.5.3 The labor market

One of the main motivations for considering wage stickiness in a macroeconomic model is that we believe that mismatching wages explain unemployment, at least in part. If a worker has a wage that is high relative to her productivity, her labor services are less attractive to an employer, and it will be difficult for her to find employment. Employment and unemployment are typically considered as occurring on the extensive margin, which is not present in this analysis, and hence I avoid that phrasing. However, also in this economy, people are working less than they would if the wage could be adjusted, which is what I refer to as *underemployment*, and can be viewed as a proxy for unemployment. The underemployment of a household is defined as $1 - \textit{Employment status}$, where the employment status is defined in equation (2.7). In words, underemployment is the fraction of lost hours worked, due to the wage deviating from the desired level. The economy-wide aggregate underemployment is the average of all households' individual underemployment.

Underemployment heterogeneity A model with rich household heterogeneity cannot only address questions of not only aggregate underemployment, but also questions of what characterizes those who suffer from it. Figure 2.9 shows how the employment status varies over the state space for non-updating households ($f = 0$). It also indicates which regions of the state space are most relevant in terms of household density. From




Figure 2.9: Employment status of non-updating households ($f = 0$). The leftmost panel shows the case of lowest possible idiosyncratic productivity, the middle panel shows median productivity, and the rightmost panel shows the highest possible productivity. Grey scale indicates employment status, as defined by equation (2.7), in percent. Blue dots show the distribution of households, scaled up to compensate for the fact that the extreme productivity levels contain much fewer households than the median-productivity level. Dot size represents the concentration of households. Axes limits have been chosen to exclude at most 0.5 percent of total households at each tail of each dimension.

the left-hand panel, we learn that a low-productive household must have a very low wage in order to have an employment status substantially higher than zero. The wealth dimension is of higher-order importance. Most households in this group have very little wealth and a low-enough wage to be employed. However, although difficult to read from the figure, a substantial share of this group suffers from involuntary underemployment. Figure 2.10 shows underemployment, voluntary and involuntary, for the different levels of idiosyncratic productivity. It reveals that involuntary underemployment is around $1/3$ in the group of least-productive households, while voluntary underemployment practically does not exist in this group.²⁴ Bear in mind that this group is very small, containing only 0.01% of the households. But the pattern prevails for other low levels of productivity.

The middle panel of figure 2.9 shows that a median-productive household is employable at higher wages, but only up to a certain level. Wages

²⁴To be clear, this does not mean that $1/3$ of the households in this group are more or less underemployed. Rather, it means that probably more than $1/3$ of the households in this group partly suffer from involuntary underemployment, such that the average hours worked in this group are $2/3$ of what they would be if everyone could reset their wage.

only little above this level lead to a substantial decrease of the employment status. Too low a wage, on the other hand, leads to a suppressed labor supply on a voluntary basis. We can also see that most households in this group are bunched around the wage level leading to full employment status, although deviations above and below occur, which is confirmed by the middle bars of figure 2.10. The group of median-productive households is fairly large, containing 21% of the households.

Finally, the right-hand panel of figure 2.9 shows that only a very high wage renders involuntary underemployment for a top-productive household. However, being stuck with a wage below the highest region leads to voluntary underemployment. Close to none of the top-productive households suffer from involuntary underemployment. Instead, it is among the high-productive households that we find the voluntary underemployed. As a share of total households, the very top-productive households are few, only 0.01% of the total population.




Figure 2.10: Underemployment in percent, divided into **involuntary** and **voluntary**, for different productivity levels. Underemployment is defined as $1 - \textit{Employment status}$, see equation (2.7).

An overall take-away from figure 2.10 is that involuntary underemployment co-varies negatively with productivity, while voluntary underem-

ployment does show, if anything, the opposite pattern.²⁵ It also reveals that a negative voluntary underemployment, i.e., overemployment, an employment status > 1 , is common in the group of households with a productivity slightly below the median. Voluntary underemployment is not the main focus of this paper. However, one view is to think of this as suppressed labor-force participation.

Figure 2.9 illustrates that the wage is an important determinant of the employment status, at least conditional on idiosyncratic productivity. Figure 2.11 further investigates the interaction between wages and employment status at the household level. The top panel shows the hourly wage on the horizontal axis. The colors show whether labor demand is binding or not, i.e., whether a deviation from full employment status is voluntary or not. It is not clear from the chart that involuntary underemployment is caused by “too high” a wage. There are fully employed households with a high wage, while a large share of the involuntary underemployment is found in households with a relatively low wage. However, it is clear that all households with an employment status below 50% are involuntary underemployed, while almost all slightly underemployed households are so voluntarily.

Recall that labor demand is determined by the *effective* wage (w/e), rather than the hourly wage.²⁶ Hence, the lower panel of figure 2.11, replacing the hourly wage with the effective wage on the horizontal axis, illustrates the labor-market dynamics better. Note that here colors represent idiosyncratic productivity and not employment categories. It is clear that although involuntary underemployed households have diverse hourly wages, all of them have an above-average effective wage, suppressing labor demand, and they are often low-productive. This means that a fully employed household with a high wage must also be highly productive. And an involuntary underemployed household with a low wage has a very

²⁵If the wage friction applies to the effective wage instead of the hourly wage, the involuntary underemployment co-varies positively with productivity. See Åhl (2020) for some details.

²⁶In fact, the demand for effective labor (el) is completely determined by the effective wage and aggregate factors. It is best realized by reshuffling equation (2.10) as

$$e(j)n^*(j) = \left(\frac{w(j)/e(j)}{W_d} \right)^{-\delta} N_d.$$



Figure 2.11: *Top panel:* interaction of hourly wage and employment status. Colors represent **non-updating** ($f = 0$) and **wage-updating** ($f = 1$) households. *Bottom panel:* interaction of effective wage and employment status. Colors represent idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$). Dot size represents the concentration of households. Axes limits are set to exclude at most 0.5 percent of households at each tail of each dimension.

low productivity. From the left part of the chart, we also learn that the typical voluntary underemployed household has a low effective wage and an above-average productivity — in line with figures 2.9 and 2.10. In

a sense, the left part of the bottom panel of figure 2.11 describes the aggregate labor-supply curve, while the right part describes the aggregate demand curve.

The figure also shows bunching of employment status around certain levels, e.g., 10%, 35% and 100%. This is an artifact of the discrete process for idiosyncratic productivity. The 35% level roughly corresponds to a one-level drop in idiosyncratic productivity, while the 10% level roughly corresponds to a two-level drop, within a spell of fixed wage. A denser grid for the idiosyncratic productivity results in less bunching of the employment status, but comes at a computational cost.

One last take-away from figure 2.11 is that the distribution of effective wages is much more concentrated than the distribution of hourly wages. This means that high-productive households tend to have higher hourly wages and vice versa. Although effective wages and idiosyncratic productivity are difficult to measure and observe, this is in line with what intuition tells us we would find in the data.




Figure 2.12: Actual labor supply as a share of demand, for households that cannot reset their wage ($f = 0$). The leftmost panel shows the case of lowest possible idiosyncratic productivity, the middle panel shows median productivity and the rightmost panel shows the highest possible productivity. Grey scale shows the ratio $\frac{g^l(b, w_{-1}, e, 0)}{n^*(b, w_{-1}, e, 0)}$, where $n^*(b, w_{-1}, e, 0)$ is given by equation (2.30). Blue dots show the distribution of households, scaled up to compensate for the fact that the extreme productivity levels contain much less households than the average-productivity level. Dot size represents the concentration of households. Axes limits have been chosen to exclude at most 0.5 percent of total households at each tail of each dimension.

Deviation from labor demand Leaving the analysis of underemployment, one novelty in this paper is the introduction of non-forcing labor demand. Natural questions arise: to what extent do households deviate from demand in their labor-supply decisions? And what characterizes the deviating households? Figure 2.12 attempts to answer this, focusing on non-updating households ($f = 0$), and three levels of idiosyncratic productivity: lowest, median, and highest.²⁷ White denotes areas where labor demand is binding, and scales of grey indicate to what extent supply deviates from demand.

The figure shows that a low-productive household must be stuck with a very low wage for labor demand not to bind. With higher productivity, there is still a “threshold wage” beneath which labor demand is no longer binding. The threshold is higher the higher is idiosyncratic productivity, which is natural since a higher idiosyncratic productivity attracts higher labor demand at any given wage. Below the threshold wage, there is a rapid increase in demand and hence, the supply-demand ratio drops to low regions. This indicates the importance of labor demand not being forcing for households. Although only 25% of the households supply exactly the demanded labor, the average labor supply-demand ratio is 78%, and shows a weak negative co-variance with idiosyncratic productivity.

The labor supply-demand ratio is, arguably, less interesting for the case of wage-updating households ($f = 1$). Nevertheless, it is shown in appendix figure 2.37 (note that the grey scale starts at 70). There is no clear dependency on either idiosyncratic productivity or wealth. Instead, there is a constant pattern that a wage is chosen such that labor demand is slightly above the desired supply. This gap renders lower-than-necessary current earnings, but provides a sort of “insurance” against severe involuntary underemployment in the case of future bad productivity shocks.

To summarize the labor market in a general-equilibrium steady state, each idiosyncratic productivity level has a tight matching region for the wage, in the sense that demand approximately equals the desired supply at this productivity-wage combination. A wage-updating household chooses

²⁷The labor supply-demand ratio is defined as $\frac{g^l(b, w_{-1}, e; f)}{n^*(b, w_{-1}, e; f)}$. Note that figure 2.12 shows how labor supply relates to *demand*, while figure 2.9 shows how labor supply relates to *desired supply*, which is a different thing.

a wage within this region. Exactly where is also determined by her wealth, which is less important for the wage decision. A household stuck with a mismatching high wage will suffer from suppressed labor demand and involuntary underemployment. A household stuck with a mismatching low wage will meet a high labor demand, but deviate from it downwards in actual labor supply. She might also work less than would be the case if she could reset her wage, but this voluntary underemployment is quantitatively less important than involuntary underemployment.

2.5.4 Consumption and savings

Although wage stickiness is mainly interesting for labor-market dynamics, the consumption-savings decision is at the core of all macroeconomic models. The interaction of consumption and net savings (i.e., the difference between the amount of bonds a household brings to the next period, and the amount brought from the last period) is shown in figure 2.13. A weak positive relationship is revealed. There are no households consuming little while also building up wealth — in this economy you do not become rich by a cheap lifestyle. However, there are clearly households keeping up consumption by dis-saving; these are mainly low-productive and found in the lower half of the figure. These households are typically involuntary underemployed (not shown in the figure), eating out of their wealth while waiting to end the underemployment spell by either an opportunity to update the wage, or a positive productivity shock.

Appendix figure 2.27 reveals a similar pattern when wages are flexible, with one difference being that low-productive households dis-save to a lesser extent. The reason is that if wages are flexible, there is no underemployment and hence, all low-productive households have earnings significantly above zero. Another notable difference is that if the wages are flexible, high-productive households always consume much, although their net savings differ. This is not the case with sticky wages, where also a high-productive household might suffer from both poverty and underemployment, and hence be forced to consume little.

It is common for households with a consumption below about 0.9 to have a zero-net-saving behavior. However, when the income is high enough to allow a higher consumption, the optimal trade-off is to also build up wealth for the future, mainly for consumption-smoothing reasons. When




Figure 2.13: Interaction of consumption and net savings. Colors represent idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$), and dot size represents the concentration of households. Axes limits are set to exclude at most 0.5 percent of households at each tail of each dimension.

the income is very high, the idiosyncratic productivity must also be high, and will hence likely decrease in the future.

Marginal propensity to consume and work In the past years, there has been a focus in the literature on the distribution of households' marginal propensity to consume (MPC), not the least because this seems to give a better understanding of the potency and transmission mechanism of monetary and fiscal policy.²⁸ The term “wealthy hand-to-mouth” was coined by Kaplan and Violante (2014), and seems to be well in line with the data, although MPC is difficult to measure empirically. How is the MPC affected by how we model household-level wage stickiness, and how does it relate to wealth? One answer is given in the top panel of figure 2.14, revealing that only households very close to the borrowing limit have a MPC out of wealth higher than 0.3. The economy-wide average is 0.14 in

²⁸By my definition, MPC denotes the share of a marginal unexpected increase in wealth (government bonds, b) that is consumed within the same period. E.g., $\text{MPC}=0.5$ means that consumption increases (decreases) by 50% of a marginal increase (decrease) in wealth.

the model, which is low as compared to the empirical estimates, see, e.g., the references listed in Kaplan and Violante (2014). Although my model does not produce wealthy hand-to-mouth households, it does produce some high-productive households with a high MPC, which is rare.²⁹ These are mainly poor households stuck with a low wage, preferring to give up consumption rather than working more at such a low wage, i.e., voluntary underemployed.

A concept closely related to MPC is the marginal propensity to supply labor (MPL), measuring the response of labor supply to a marginal change in wealth. The MPL is shown in the lower panel of figure 2.14, and the immediate impression is that it mirrors the MPC chart quite well. Households that are bound by labor demand mechanically have a $MPL=0$, since labor demand is the only thing shifting their hours worked, and labor demand is, of course, independent of a household's wealth.³⁰ However, for households not bound by labor demand, a natural response to a marginally increasing wealth is to decrease labor supply, trading off increases in consumption and leisure. I.e., the MPL is expected to be zero or negative, which is confirmed by figure 2.14. If the household is at the borrowing limit with a low idiosyncratic productivity, her labor supply might be high in order to keep consumption at an acceptable level, and the MPL might hence be far below zero.

So, should we expect to find a household given in mirroring positions in figure 2.14? Figure 2.15 shows that the answer is both yes and no, by plotting the interaction of MPL and MPC. There are two clear patterns; one downward sloping, and one vertical at $MPL=0$. The very highest MPC is found among involuntary underemployed households. These households are bound by labor demand, and hence their MPL must be zero, and they are found in the vertical pattern in the figure. The downward-sloping pattern consists of households not bound by labor demand. These households trade off consumption and leisure in response to marginal

²⁹In an economy without a precautionary motive to save, e.g., if markets are complete, the MPC typically equals the net interest rate, in accordance with the permanent-income hypothesis. In a model with incomplete markets and flexible wages, the MPC is typically significantly higher only for low-productive households; see appendix figure 2.28.

³⁰There could be an indirect effect on labor demand of a wealth change, if it also renders a change of the wage. However, this only applies to wage updating households, $f = 1$.




Figure 2.14: *Top panel:* interaction of wealth and MPC. *Bottom panel:* interaction of wealth and MPL.

Colors represent idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$), and dot size represents the concentration of households. Axes limits are set to exclude at most 0.5 percent of households at each tail of each dimension.

changes in their wealth.

2.5.5 Welfare heterogeneity

In a model with rich household heterogeneity, a natural question of how well off different households are arises. One way of addressing this question is to study the value function capturing both instantaneous and expected




Figure 2.15: Interaction of marginal propensities to work and consume. Colors represent idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$), and dot size represents the concentration of households. Axes limits are set to exclude at most 0.5 percent of households at each tail of each dimension.

future utility, see the Bellman equation (2.2). However, the value of a household is an abstract concept which is not observed, and lacks a measurement unit. Hence, it is often helpful to translate the value into something more concrete and easy to measure, such as wealth.

Given that the novelty of this paper is household-level wage friction, this section is focused on the lost value caused hereby. Shades of gray in figure 2.16 show the compensation (in wealth, i.e., government bonds) needed to make a non-updating household ($f = 0$) indifferent between receiving the compensation, and being able to reset the wage in the current period.³¹ This is a hypothetical analysis made without general-equilibrium feedback, i.e., all prices and aggregate outcomes are being held constant. The figure breaks this concept down to different regions of the state space, by wealth (b , horizontal axes), wage (w , vertical axes), and idiosyncratic productivity (e , panels). On top, the blue dots show how households are distributed over the state space.

³¹In consequence, for every wealth-productivity combination, there exists a beginning-of-period wage such that the compensation needed is zero, namely the (optimal) wage chosen by wage-updating households with the same wealth and productivity.




Figure 2.16: The welfare loss of non-updating households ($f = 0$), measured as the extra wealth that would compensate for the inability to update the wage. The leftmost panel shows the case of lowest possible idiosyncratic productivity, the middle panel shows median productivity and the rightmost panel shows the highest possible productivity. Grey scale indicates the additional wealth needed to be indifferent between receiving the wealth or an immediate opportunity to update the wage. Blue dots show the distribution of households, scaled up to compensate for the fact that the extreme productivity levels contain much fewer households than the median-productivity level. Dot size represents the concentration of households. Axes limits have been chosen to exclude at most 0.5 percent of total households at each tail of each dimension.

First, note that different scales for the gray shades are used for different levels of idiosyncratic productivity. The leftmost panel shows the group of very lowest-productive households. The optimal wage is low, and we learned from figure 2.9 that a wage more than slightly above it renders very low labor demand, and hence involuntary underemployment. Among the involuntary underemployed, there is a positive compensation making a household indifferent to being able to reset the wage and hence, end the underemployment spell. Except for a region close to the optimal wage, how far off the wage is plays a second-order role. This is because earnings will be very low, due to close to zero hours worked, regardless of how far away from the optimum the wage is. More significantly, the compensation shrinks as the wealth increases, since a wealthy household can keep up consumption by dis-saving while waiting for either an opportunity to update the wage, or a positive productivity shock, to end the underemployment spell. However, all households at this productivity level are wealth poor, and most have a very low wage.

The middle panel shows the relatively large group of median-productive households. Here, we find both households with a positive compensation

making them indifferent to lowering the wage to get out of involuntary underemployment, and households with a positive compensation making them indifferent to increasing the wage to maximize their earnings. Some households in this group suffer more, in terms of compensating wealth, than the worst off households in the low-productive group. However, most households have a wage that matches their productivity quite well, and hence suffer only marginally from the wage friction.

We now turn to the group of highest-productive households, in the rightmost panel of figure 2.16. Here the optimal wage is, of course, higher. Being this productive provides a rare but great opportunity to work hard, earn a great deal, and build a substantial wealth for the future — that is, if the wage matches the productivity, that is. Hence, we see compensations making low-wage households indifferent to increasing their wage and hence their earnings. The compensation depends very little on the current level of wealth, but increases considerably with a falling wage. The reason is that earnings are, to a considerable extent, explained by the wage in this group. However, at this productivity, it is almost impossible to be involuntary underemployed, so very little is needed to compensate for not being able to lower the wage.

From comparing the scales of the different panels, we see that it takes a much higher amount to compensate a high-productive low-wage household to indifference, than a low-productive underemployed household. This might be surprising, considering that involuntary underemployment is a very undesirable situation in this economy, as it almost completely takes away earnings. The reason behind the different compensating levels is that the stationary property of the process for idiosyncratic productivity makes it very likely that the productivity will rise for the least-productive households, and thus end the underemployment spell. Should this not be the case, it is quite likely that a wage-setting opportunity ends the underemployment within a few periods. However, the opposite is true for the top-productive households. Their productivity is expected to decrease drastically in the future, so if stuck with a low wage, they forego a very rare opportunity to earn spectacular amounts, and hence to build wealth very quickly.

2.5.6 Marginal distributions

Ideally, a model with rich household heterogeneity matches a number of cross-sectional data distributions. Section 2.4 provides a number of such distributions for the US economy. This section compares the model outcomes with real-world data. However, before proceeding, a few clarifying comments are in place.

There are several reasons why we should not expect the model to match the data distributions. For one, the model is not calibrated to match any data, see section 2.3. Further, this analysis concerns the steady-state behavior of the model, not only assuming that no aggregate shocks have hit the economy, but also that all agents in the economy are convinced that no aggregate shocks will ever occur. There are strong reasons to believe that this was not the case when the US data was collected. Yet another thing to bear in mind is that in the model, all agents are assumed to be fully informed and act completely rationally. Even if rationality is a good assumption for the aggregate behavior in the real-world economy, which is far from certain, this is certainly not true at the individual level. Given these caveats, this section investigates the distributions of household decisions for observable variables, one at a time. The exercise should be viewed as qualitative rather than quantitative.

Figure 2.17 (partly the same as the top-left panel of figure 2.7) shows the wealth distribution in the model, and has its data counterpart in figure 2.2.³² Both the model and the data show a sharp spike of households at zero wealth. In the data, there are households with negative net wealth, which is prevented by the exogenous borrowing limit in the model. However, the model is in line with the data regarding a fat right tail, although not pronounced enough to replicate the high level of wealth inequality found in the US data. This is a common issue in the class of heterogeneous-agents models relying on incomplete markets and shocks to idiosyncratic productivity; see, for example Hubmer et al. (2020). In the model, wealth co-varies positively with idiosyncratic productivity, in line with figure 2.8.

Figure 2.18 (partly the same as the bottom-right panel of figure 2.7) shows that the model produces a positively skewed wage distribution,

³²For a comparison with the model with fully flexible wages, see appendix figure 2.31.




Figure 2.17: Model distribution of wealth. Left is the probability mass function approximated by a histogram. Right is the corresponding cumulative distribution function. Bar color represents idiosyncratic productivity, from **lowest** to **highest**. The limits of the horizontal axes are set to exclude at most 0.5 percent of households at each tail of the distribution.




Figure 2.18: Model distribution of hourly wages. Left is the probability mass function approximated by a histogram. Right is the corresponding cumulative distribution function. Bar color represents idiosyncratic productivity, from **lowest** to **highest**. The limits of the horizontal axis are set to exclude at most 0.5 percent of households at each tail of the distribution.

which is also the case in the data, see figure 2.3.^{33,34} In the model, there is no legislated minimum wage causing a pronounced spike, as is the case in the data.³⁵ The model fails to produce the fat right tail of the data distribution. The multi-modal appearance of the wage distribution is an artifact of the discrete productivity distribution. Each spike represents an optimal-wage region, corresponding to a feasible productivity level, for a household that is allowed to reset the wage. But some households will stay in that region due to the friction, despite changing productivity. A denser grid for the idiosyncratic productivity would yield a smoother distribution, but comes at a computational cost, and would add very little to the overall model properties.




Figure 2.19: Model distribution of labor supply. Left is the probability mass function approximated by a histogram. Right is the corresponding cumulative distribution function. Bar color represents idiosyncratic productivity, from **lowest** to **highest**. The limits of the horizontal axis are set to exclude at most 0.5 percent of households at each tail of the distribution.

Figure 2.19 shows a spike in labor supply at a level which can be

³³The skewness of the wage distribution follows the skewness of the idiosyncratic-productivity distribution, see figure 2.1.

³⁴For a comparison with the model with fully flexible wages, see appendix figure 2.32.

³⁵The US federal minimum wage was USD 7.25 in 2017, but may differ at the state level.

thought of as “full-time” work, with a bunching of households around it.³⁶ Below this level, we see clusters of households at quite low levels of labor supply. These are involuntary underemployed households, and the clusters correspond to level drops in productivity, see section 2.5.3. In line with figures 2.9 and 2.10, low-productive households are over-represented here. The data counterpart in figure 2.4 indeed shows a spike at full time, with some households working less.³⁷ However, the real world data also shows households working far more than full time, contrary to the model.




Figure 2.20: Model distribution of earnings. Left is the probability mass function approximated by a histogram. Right is the corresponding cumulative distribution function. Bar color represents idiosyncratic productivity, from **lowest** to **highest**. The limits of the horizontal axis are set to exclude at most 0.5 percent of households at each tail of the distribution.

The earnings distribution generated by the model is shown in figure 2.20, with the data counterpart in figure 2.5.³⁸ Despite missing an extensive labor margin, the model replicates quite well the share of households with zero or very low earnings. Earnings are positively skewed in the model, although perhaps not as much as in the data. Not surprisingly, the

³⁶For a comparison with the model with fully flexible wages, see appendix figure 2.33.

³⁷Note that households without a wage are excluded in figure 2.4. This probably excludes most unemployed households.

³⁸For a comparison with the model with fully flexible wages, see appendix figure 2.34.

earnings tend to be higher for high-productive households in the model.




Figure 2.21: Model distribution of consumption. Left is the probability mass function approximated by a histogram. Right is the corresponding cumulative distribution function. Bar color represents idiosyncratic productivity, from **lowest** to **highest**. The limits of the horizontal axis are set to exclude at most 0.5 percent of households at each tail of the distribution.

The distribution of consumption in the model economy is shown in figure 2.21, and is less positively skewed than the data counterpart in figure 2.6.³⁹ Moreover, the data shows households consuming less than 1/3 of the median. This is not the case in the model economy, where all households receive firm dividends, which can be considered as capital income, lump-sum. That is clearly not the case in the US economy, and is further discussed in section 2.5.7. Figure 2.21 also reveals that high-productive households tend to consume more than others.

How are these distributional model features affected by the main contribution of this paper, i.e., idiosyncratic wage stickiness? One answer to this question is to make a comparison with the corresponding features of the same model with fully flexible wages. The counterparts of figures 2.17-2.21 are found in appendix section 2.A.4. Most distributions are qualitatively alike in the two models. One notable difference is the distribution of hours worked, which is much more concentrated around

³⁹For a comparison with the model with fully flexible wages, see appendix figure 2.35.

the full-time level with flexible wages, see appendix figure 2.33. The main reason is that there are no underemployment in the model with flexible wages. For the same reason, there are no households with close to zero earnings if wages are flexible, see appendix figure 2.34.

To summarize this section on marginal distributions generated by the model, one conclusion is that the model captures many features of the US data qualitatively well. There are natural explanations for the features that are captured less well. A serious calibration, aiming at mimicking some features of the data, could be interesting, but is beyond the scope of this paper.

2.5.7 Sensitivity analysis

An important purpose of this chapter is to investigate the impact of household-level wage stickiness has on the aggregate economy with incomplete markets. As is clear from section 2.3, the model is not calibrated to match any particular features of the data. Questions may arise about to what extent the results and mechanisms I have presented hinge on the parameter values I use. In this section, I alter the values of some of the key parameters of the model, one at the time. This yields a possibility to study the impact of each parameter. It also yields an opportunity to delve deeper into the mechanisms present in this model.

I conduct five different sensitivity exercises: higher wage stickiness, less substitutable skill types, heterogeneous returns from firm ownership, and a lower volatility or persistence of the idiosyncratic productivity process. Each is explained in detail and commented below. A comparison of the effects on aggregate outcomes is shown in table 2.3.

Higher wage stickiness

The key difference between this economy and a standard HANK model is the Calvo (1983) friction applying to individual wages. Thus, a natural first candidate for sensitivity analysis is the degree of wage stickiness.⁴⁰ In this exercise, I consider a substantially higher degree of stickiness, namely a $\theta_w = 9/10$, instead of $\theta_w = 3/4$ as in the baseline. I.e., the chance of

⁴⁰Also the empirical literature arrives at different conclusions about the degree of wage stickiness. See, e.g., Grigsby et al. (2019) and references therein.

Table 2.3: Sensitivity of aggregate outcomes

	Baseline	Higher stickiness	Less substitut.	Heterog. returns	Smaller volatility	Less persist.
	$\theta_w = 9/10$	$\delta = 2$	$D \propto e$	$\sigma_e = .13/2$	$\rho_e = .7$	
Real rate ($R - 1$)	-.88%	-.95%	.21%	-2.81%	.49%	.40%
Wage (W)	.57	.57	.60	.57	.57	.58
Output (C, Y)	.88	.87	.79	.87	.87	.84
Hours worked (L, N)	.82	.81	.71	.82	.81	.77
Firm profits (D)	.41	.41	.37	.41*	.40	.39
Underemployment	8.5%	17.9%	-0.2%	11.2%	7.0%	12.2%
Involuntary	10.6%	17.1%	0.9%	11.6%	7.5%	11.8%
Voluntary	-2.2%	0.8%	-1.1%	-0.4%	-0.5%	0.4%
On labor demand	24.6%	29.0%	95.8%	17.5%	39.7%	43.7%
On borrowing limit	25.5%	20.7%	19.4%	20.2%	11.7%	3.9%
<i>Gini coefficient</i>						
Wealth	.67	.64	.64	.63	.59	.37
Wage	.23	.20	.13	.25	.11	.06
Productivity	.27	.27	.27	.27	.14	.10
Consumption	.12	.13	.07	.20	.07	.03
Labor supply	.15	.18	.09	.17	.10	.15
Earnings	.31	.31	.16	.26	.17	.18
Income	.16	.16	.09	.25	.09	.09

*Average value. Individual outcomes lie in the interval [0.05, 2.36].

updating the wage is 10% instead of 25% each quarter. Note that the case $\theta_w = 0$ corresponds to fully flexible wages and the standard HANK model.⁴¹

Making a comparison with the baseline in table 2.3, we note several things. First, longer spells with a fixed wage result in a higher risk that idiosyncratic productivity drops during the spell, leading to involuntary underemployment. The aggregate underemployment rate is higher in this exercise, mainly due to involuntary underemployment being higher. One additional effect is that, as the risk for underemployment increases, so does the precautionary motive to save. Thus, the equilibrium interest rate is pushed even further down, although not by very much.

Figure 2.22, to be compared with figure 2.10 for the baseline, illustrates a much higher involuntary underemployment among low productive

⁴¹Non-forcing labor demand is irrelevant if wages are flexible, since labor demand is determined, in each period by all households via the wage choice. Hence, there is never any incentive to deviate from labor demand.

households. In addition, high productive households are voluntary underemployed to a greater extent, by a similar reasoning.




Figure 2.22: Underemployment, divided into **involuntary** and **voluntary**, for different productivity levels. Underemployment is defined as $1 - \text{Employment status}$, see equation (2.7).

Less substitutable skill types

As pointed out in section 2.5.3, involuntary underemployment arises when the wage is high relative to idiosyncratic productivity, because a firm can substitute that worker for someone else, with a different skill type. From this reasoning, it is intuitive that how easy it is for a firm to substitute between skill types is important for the functioning of the labor market. This is captured by the model parameter for elasticity of substitution between skill types, δ , see equation (2.8). A lower elasticity, which is analyzed here, makes it more difficult for firms to substitute an expensive skill type for a cheaper one, without suffering any severe output damage. Hence, labor demand is less affected by a wage-productivity mismatch. This is illustrated in a stylized way in figure 2.23, where the labor demand for two households, a and b , is shown on the different axes. They have the same idiosyncratic productivity, but a 's wage is 20% higher than b 's. The blue and red curves illustrate different possible combinations of labor input

from a and b to end up with a certain level of aggregate labor input, for $\delta \in \{2, 5\}$. The circles show the firm's optimal labor demand, respectively, taking the relative wage of a and b into account. With the high elasticity of substitution (baseline case, red), the demand for a 's labor is 67% lower than that for b . However, with the lower elasticity (blue), the difference is only 31%. I.e., a is hit much harder by the wage-productivity mismatch if the elasticity is high, as in the baseline.




Figure 2.23: A stylized illustration of different elasticities of substitution between skill types, δ . Axes show demand for hours worked by the two workers, respectively. Curves show labor combinations resulting in the same aggregate effective labor for the two cases: **low substitutability** ($\delta = 2$), and **high substitutability** ($\delta = 5$, baseline). Circles indicate cost-minimizing compositions when worker a 's wage is 20% higher than that of worker b , but their idiosyncratic productivities coincide.

As expected, the problem of involuntary underemployment decreases, as can be seen from table 2.3, with a lower elasticity of substitution between skill types. In fact, it disappears almost completely. Also note that, as the risk of underemployment is so much lower, almost all households obey labor demand. Another consequence of less underemployment risk is that the precautionary savings motive is mitigated, and the equilibrium rate is higher. Moreover, a more equally distributed labor supply results in more equally distributed earnings and consumption, which can be read from the Gini coefficients.

One more notable feature in table 2.3 is that aggregate labor supply is substantially lower, and consequently so is output. This is not necessarily because households work less hours. The explanation, more technical than economic in its nature, is mainly that the firms' technology to aggregate labor has changed.

Heterogeneous dividends

In the baseline analysis, I assume that firms are owned by households, but there is no possibility to trade the shares, and the profits are paid lump-sum as dividends to the households. This is a simplification to avoid the problems that tend to arise when multiple assets are introduced, and hence a portfolio choice.⁴² The baseline assumption is that the shares are uniformly distributed among all households, so that the dividend income is equally large for all. Under this assumption, the dividend income is almost as large as the earnings for the median household. For a low-productive, or severely underemployed household, the dividend income is much larger than the earnings, and hence constitutes a large share of total income. In a sense, this works as an insurance against underemployment. However, it is clear from the data that poor, low-earning households rarely have any dividend income from equity shares. An alternative assumption, which does not impose any complications by distorting households' decisions, is to let the shares in the firms be unequally distributed as a function of something exogenous to the households. A suitable candidate is idiosyncratic productivity, since we know that it correlates well with earnings and wealth in the model economy. In this sensitivity analysis, I replace the lump-sum dividends D in the households' budget constraint (2.3) by a function $D(e)$. More specifically, I let $D(e) = e\bar{D}$, so that dividend income is proportional to idiosyncratic productivity. \bar{D} is a constant chosen so that aggregate dividends are consistent with the equilibrium condition (2.20). The difference between the assumptions is illustrated in figure 2.24.⁴³

Without the income insurance provided by high dividends, the fear of

⁴²See, e.g., Foltyn (2020).

⁴³In a similar way, one could consider making the lump-sum tax/transfer T progressive, and hence taking on a role as an insurance against low earnings. However, taxes are very small compared to dividends in this model, and hence the impact of such an exercise is very limited.




Figure 2.24: Distribution of firm profits as a function of idiosyncratic productivity: **proportional**, and **uniform** (baseline).

becoming underemployed due to decreasing productivity becomes more severe. This enhances the precautionary motive to save, and decreases the equilibrium rate dramatically, see table 2.3. Despite the extremely low interest rate, fewer households are found at the borrowing limit as compared to the baseline.

Although the earnings inequality is slightly smaller in this case, the positive co-variance between productivity and earnings (and hence also between dividend income and earnings) makes the income inequality, where also dividends are included, much larger than in the baseline. In response, consumption inequality is also higher when dividends are distributed in this way.

Alternative process for idiosyncratic productivity

The idiosyncrasy of productivity is the source of household heterogeneity in this model. It is also the main driver, in combination with the wage stickiness, of underemployment. However, productivity is not observable on the household level. Estimates of the process for idiosyncratic productivity often infer it from labor-market outcomes, with some underlying structural model. However, these models typically handle wage setting very

differently than I do in this economy and hence, an off-the-shelf process for idiosyncratic productivity might be inappropriate in this model.

I assume that idiosyncratic productivity follows an $AR(1)$ process with i.i.d. stochastic innovations, see equation (2.1). The properties of the productivity process are determined by two parameters: ρ_e , determining how persistent the process is; and σ_e , determining how large the innovations, or jumps, in productivity are. Here, I analyze the effect of smaller innovations, and of less persistence in the process.

First, consider half as large a standard deviation of the innovations to idiosyncratic productivity as in the baseline. With smaller jumps, the risk of a large wage-productivity mismatch, and hence severe underemployment, decreases. Table 2.3 shows that underemployment is lower with this assumption, and also that the equilibrium rate is higher, due to less precautionary motives to save. The lower inequality in productivity is reflected in lower overall inequality. As an example, it is less likely to draw enough negative shocks to end up at the borrowing limit.

On the other hand, if the innovation size is the same as in the baseline, but the deviations from average productivity are assumed to be less persistent, a household runs the same risk of being hit by underemployment, *ceteris paribus*, as in the baseline. However, the low productivity tends to return upwards more quickly, making the underemployment spells shorter, and thus less severe. Underemployment is more common than in the baseline, due to different wage-setting behavior, but not as deterrent. Hence, the precautionary motive is weaker, and the equilibrium rate higher. Even though the variance of the productivity innovations is the same as in the baseline, the lower persistence makes the unconditional distribution of productivity less disperse, reflected in a lower Gini coefficient for productivity. In general, this results in lower inequality, with labor supply as an exception.

Summary of sensitivity analysis A general conclusion from the sensitivity exercises in this section is that the overall qualitative results are robust to alternative values of crucial parameters. The risk of involuntary underemployment is a key driver of the precautionary motive to save in the baseline model. How this is affected by alternative parameter values — either the risk of becoming underemployed, or the severity of being

so — determines the effect on the equilibrium real interest rate, which reflects the precautionary savings motive. Other aggregate variables, such as the wage and production, are affected to a smaller extent. Inequality is somewhat affected by the parameter assumptions, and these variations are in line with intuitive reasoning.

2.6 Conclusions

In this paper, I have developed a general-equilibrium model featuring household heterogeneity, individual-level wage stickiness, and standard price stickiness on the firm side. The model also features non-forcing labor demand, which is non-standard in the literature. I used the model to analyze a steady state without aggregate uncertainty. The main focus was on how the wage friction affects aggregate outcomes, in particular in the labor market.

The combination of idiosyncratic shocks to productivity, and household-level wage stickiness, creates occasional wage-productivity mismatches for individual households. These mismatches affect labor demand, and may consequently lead to spells of severe underemployment at the household level. My main findings are that this mechanism has a large impact on the precautionary motive to save and hence, on the equilibrium interest rate. The underemployment also distorts labor supply, leading to suppressed production compared to an economy with fully flexible wages.

The model produces cross-sectional distributions of observable household characteristics that match the US data qualitatively. I have also shown that the main results are robust to variations in crucial parameter assumptions.

Earlier literature has proved wage frictions to make the transmission mechanisms of aggregate shocks more realistic. In this paper, I have only studied an economy in steady state, without aggregate fluctuations. A natural next step is to investigate how the type of wage friction I have studied here affects the impact and transmission mechanisms of aggregate shocks.

References

- Åhl, M. (2020). Wage stickiness and household heterogeneity. Thesis chapter 1, Institute for International Economic Studies.
- Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, 109(3):659–684.
- Arrow, K. J. and Debreu, G. (1954). Existence of an Equilibrium for a Competitive Economy. *Econometrica*, 22(3):265–290.
- Barattieri, A., Basu, S., and Gottschalk, P. (2014). Some Evidence on the Importance of Sticky Wages. *American Economic Journal: Macroeconomics*, 6(1):70–101.
- Bellman, R. (1957). A Markovian decision process. *Journal of mathematics and mechanics*, pages 679–684.
- Bewley, T. (1986). Stationary monetary equilibrium with a continuum of independently fluctuating consumers. *Contributions to mathematical economics in honor of Gérard Debreu*, 79.
- Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383 – 398.
- Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312 – 320.
- Dixit, A. K. and Stiglitz, J. E. (1977). Monopolistic Competition and Optimum Product Diversity. *The American Economic Review*, 67(3):297–308.
- Erceg, C. J., Henderson, D. W., and Levin, A. T. (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics*, 46(2):281 – 313.
- Fisher, I. (1930). *Theory of interest: as determined by impatience to spend income and opportunity to invest it*. Augustusm Kelly Publishers, Clifton.

- Foltyn, R. (2020). Experience-based Learning, Stock Market Participation and Portfolio Choice. *Available at SSRN 3543442*.
- Galí, J. (2008). Introduction to Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework. In *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Introductory Chapters. Princeton University Press.
- Grigsby, J., Hurst, E., and Yildirmaz, A. (2019). Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data. Working Paper 25628, National Bureau of Economic Research.
- Hubmer, J., Krusell, P., and Smith, A. (2020). Sources of U.S. Wealth Inequality: Past, Present, and Future. In *NBER Macroeconomics Annual 2020, volume 35*. National Bureau of Economic Research, Inc.
- Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5):953 – 969.
- Imrohoroglu, A. (1989). Cost of Business Cycles with Indivisibilities and Liquidity Constraints. *Journal of Political Economy*, 97(6):1364–1383.
- Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary Policy According to HANK. *American Economic Review*, 108(3):697–743.
- Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.
- Kaplan, G. and Violante, G. L. (2018). Microeconomic Heterogeneity and Macroeconomic Shocks. *Journal of Economic Perspectives*, 32(3):167–94.
- Krusell, P., Mukoyama, T., and Smith, A. A. (2011). Asset prices in a Huggett economy. *Journal of Economic Theory*, 146(3):812 – 844. Incompleteness and Uncertainty in Economics.
- MaCurdy, T. E. (1981). An Empirical Model of Labor Supply in a Life-Cycle Setting. *Journal of Political Economy*, 89(6):1059–1085.

- McKay, A., Nakamura, E., and Steinsson, J. (2016). The Power of Forward Guidance Revisited. *American Economic Review*, 106(10):3133–58.
- Nakamura, E., Steinsson, J., Sun, P., and Villar, D. (2018). The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation*. *The Quarterly Journal of Economics*, 133(4):1933–1980.
- Ravn, M. O. and Sterk, V. (2017). Job uncertainty and deep recessions. *Journal of Monetary Economics*, 90:125 – 141.
- Rouwenhorst, G. (1995). ‘Asset Pricing Implications of Equilibrium Business Cycle Models’ Chapter 10, *Frontiers of Business Cycle Research*, T. Cooley.
- Schmitt-Grohé, S. and Uribe, M. (2005). Optimal Fiscal and Monetary Policy in a Medium-Scale Macroeconomic Model. *NBER Macroeconomics Annual*, 20:383–425.
- Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, 39:195 – 214.

2.A Appendix

2.A.1 Derivation of firms' intertemporal equilibrium conditions

Here I lay out the details of the derivation of equilibrium conditions (2.17), (2.18), and (2.19) from the firms' dynamic problem (2.13). Without loss of generality, fix the current period to 0, which simplifies the notation. Note that

$$\frac{\partial \tilde{d}_{t|0}(i)}{\partial p_0(i)} = (1 - \epsilon) \left(\frac{p_0(i)}{P_t} \right)^{-\epsilon} C_t + \frac{\epsilon}{1 - \alpha} \left(\frac{p_0(i)}{P_t} \right)^{-(1 + \frac{\epsilon}{1 - \alpha})} W_t C_t^{\frac{1}{1 - \alpha}}$$

The first-order condition to (2.13) is

$$\begin{aligned} \sum_{t=0}^{\infty} E_0 \left[(\beta \theta_p)^t \frac{\partial \tilde{d}_{t|0}(i)}{\partial p_0(i)} \right] &= 0 \quad \Longleftrightarrow \\ E_0 \left[\underbrace{\sum_{t=0}^{\infty} (\beta \theta)^t \tilde{p}_0^{*1 - \epsilon} P_t^{\epsilon} C_t}_{\equiv \tilde{J}_0^I} \right] &= \overbrace{\frac{\epsilon}{(\epsilon - 1)(1 - \alpha)}}^{\equiv \mathcal{M}} E_0 \left[\underbrace{\sum_{t=0}^{\infty} (\beta \theta)^t \tilde{p}_0^{* - \frac{\epsilon}{1 - \alpha}} P_t^{1 + \frac{\epsilon}{1 - \alpha}} W_t C_t^{\frac{1}{1 - \alpha}}}_{\equiv \tilde{J}_0^{II}} \right], \end{aligned}$$

where \tilde{p}^* is the optimal (absolute) price for firm i .⁴⁴ Further, the nominal auxiliary variable \tilde{J}_0^I can be recursively expressed as

$$\begin{aligned} \tilde{J}_0^I &= \tilde{p}_0^{*1 - \epsilon} P_0^{\epsilon} C_0 + \sum_{t=1}^{\infty} (\beta \theta_p)^t \tilde{p}_0^{*1 - \epsilon} P_t^{\epsilon} C_t \\ &= \tilde{p}_0^{*1 - \epsilon} P_0^{\epsilon} C_0 + \beta \theta_p \left(\frac{\tilde{p}_0^*}{\tilde{p}_1^*} \right)^{1 - \epsilon} \tilde{J}_1^I. \end{aligned}$$

However, \tilde{J}_0^I is a nominal variable, and needs to be scaled by the price level to be stationary. Defining $J_t^I \equiv \frac{\tilde{J}_t^I}{P_t}$, and dividing both sides by P_0 , we end up with (2.18), where $p_t^* \equiv \frac{\tilde{p}_t^*}{P_t}$ is the optimal relative price.

⁴⁴I leave out the firm-specific index i , because by symmetry all re-optimizing firms choose the same price \tilde{p}^* .

Similarly, \tilde{J}_0^{II} can be expressed recursively as

$$\tilde{J}_0^{II} = \tilde{p}_0^*{}^{-\frac{\epsilon}{1-\alpha}} P_0^{1+\frac{\epsilon}{1-\alpha}} W_0 C_0^{\frac{1}{1-\alpha}} + \beta \theta_p \left(\frac{\tilde{p}_0^*}{\tilde{p}_1^*} \right)^{-\frac{\epsilon}{1-\alpha}} \tilde{J}_1^{II},$$

and scaling yields (2.19).

2.A.2 Missing equilibrium conditions

Evolution of the price index

A standard result in the new-Keynesian literature with Calvo (1983)-type pricing friction is that the price index (CPI), defined in equation (2.15), can be expressed as $P^{1-\epsilon} = \int_0^1 p(i)^{1-\epsilon} di$. Hence, it evolves as

$$\begin{aligned} P^{1-\epsilon} &= \int_0^1 p(i)^{1-\epsilon} di \\ &= \theta \int_0^1 p_{-1}(i)^{1-\epsilon} di + (1-\theta) \int_0^1 p^*(i)^{1-\epsilon} di \\ &= \theta P_{-1}^{1-\epsilon} + (1-\theta) p^{*1-\epsilon} \\ 1 &= \theta \Pi^{\epsilon-1} + (1-\theta) p^{*1-\epsilon}. \end{aligned} \quad \Longleftrightarrow \quad (2.35)$$

Aggregate resource constraint

Integrating individual households' budget constraints gives in the aggregate resource constraint:⁴⁵

⁴⁵For compactness, I let $g^x(b, w_{-1}, e, f)$ be the shorthand notation for $\theta_w g^x(b, w_{-1}, e, 0) + (1-\theta_w) g^x(b, w_{-1}, e, 1)$ here, where $x \in \{c, b, w, l\}$ represents any of the choice variables.

$$\int_{\Omega} \left(g^c(b, w_{-1}, e, f) + \frac{g^b(b, w_{-1}, e, f)}{R} \right) = \int_{\Omega} \left(g^w(b, w_{-1}, e, f) g^l(b, w_{-1}, e, f) \right. \\ \left. + b + D - T \right) \quad \Longleftrightarrow$$

$$\int_{\Omega} g^c(b, w_{-1}, e, f) + \frac{\int_{\Omega} g^b(b, w_{-1}, e, f)}{R} = \int_{\Omega} \left(g^w(b, w_{-1}, e, f) g^l(b, w_{-1}, e, f) \right) \\ + \int_{\Omega} b + D - T \quad \begin{array}{l} (2.24) \\ (2.25) \\ (2.27) \\ \Longleftrightarrow \end{array}$$

$$C + \frac{B'}{R} = WN + B + D - T \quad \begin{array}{l} (2.23) \\ (2.20) \\ \Longleftrightarrow \end{array}$$

$$C + \frac{B'}{R} = WN + B + (Y - WN) - \left(B - \frac{B'}{R} \right) \quad \Longleftrightarrow \\ C = Y. \quad (2.36)$$

2.A.3 Complete markets

In the baseline model, markets are assumed to be incomplete. If instead complete markets are assumed, this is equivalent to assuming a representative household.⁴⁶ In this case, the distribution of households over the state space collapses to one point, and the numerical approach to the households' problem (2.2) can be replaced by the following analytical

⁴⁶Under some circumstances, which are fulfilled here; see, e.g., Schmitt-Grohé and Uribe (2005).

equilibrium conditions:

$$C^{-\sigma} = \beta RE \left[C'^{-\sigma} \right] \quad \text{Euler equation}$$

$$J^{III} = J^{IV} \quad \text{Wage-setting condition}$$

$$J^{III} = C^{-\sigma} w^* \left(\frac{w^*}{W} \right)^{-\delta} N + \beta \theta_w E \left[\left(\frac{w^*}{w'^*} \right)^{1-\delta} \Pi'^{\delta-1} J'^{III} \right] \quad \text{Recursive formulation}$$

$$J^{IV} = \frac{\delta}{\delta-1} \left(\frac{w^*}{W} \right)^{-\delta(1+\varphi)} N^{1+\varphi} + \beta \theta_w E \left[\left(\frac{w^*}{w'^*} \right)^{-\delta(1+\varphi)} \Pi'^{\delta(1+\varphi)} J'^{IV} \right] \quad \text{Recursive formulation}$$

$$W^{1-\delta} = \theta_w W_{-1}^{1-\delta} + (1 - \theta_w) w^{*1-\delta}. \quad \text{Wage evolution}$$

2.A.4 Flexible wages

If prices are assumed to be fully flexible, the households' problem (2.2) is changed such that the constraint (2.5) is removed, which is also equivalent to letting $\theta_w = 0$. In all other aspects, the model is identical. It is worth noting that with flexible wages, households' wage choice is no longer dynamic. In consequence, all incentives to ever deviate from labor demand disappear. Hence, the constraint (2.4) always holds with equality. The point of allowing labor supply to deviate from demand is not there with flexible wages.

Figures: flexible wages

Here I present some flex-wage results, captured in figures corresponding to the section 2.5 figures of the baseline model.




Figure 2.25: Flexible wages: joint distribution of idiosyncratic productivity and wage. Dot size represents the concentration of households. Axes limits have been chosen to exclude at most 0.5 percent of total households at each tail of each dimension.




Figure 2.26: Flexible wages: joint distribution of idiosyncratic productivity and wealth. Dot size represents the concentration of households. Axes limits have been chosen to exclude at most 0.5 percent of total households at each tail of each dimension.




Figure 2.27: Flexible wages: joint distribution of consumption and net savings. Colors represent idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$), and dot size represents the concentration of households. Axes limits have been chosen to exclude at most 0.5 percent of total households at each tail of each dimension.




Figure 2.28: Flexible wages: joint distribution of wealth and marginal propensity to consume (MPC). Colors represent idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$), and dot size represents the concentration of households. Axes limits have been chosen to exclude at most 0.5 percent of total households at each tail of each dimension.




Figure 2.29: Flexible wages: joint distribution of wealth and marginal propensity to supply labor (MPL). Colors represent idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$), and dot size represents the concentration of households. Axes limits have been chosen to exclude at most 0.5 percent of total households at each tail of each dimension.




Figure 2.30: Flexible wages: joint distribution of marginal propensities to consume (MPC), and to supply labor (MPL). Colors represent idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$), and dot size represents the concentration of households. Axes limits have been chosen to exclude at most 0.5 percent of total households at each tail of each dimension.




Figure 2.31: Flexible wages: model distribution of wealth. Bar color represents idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$).




Figure 2.32: Flexible wages: model distribution of wages. Bar color represents idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$).




Figure 2.33: Flexible wages: model distribution of hours worked. Bar color represents idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$).




Figure 2.34: Flexible wages: model distribution of earnings. Bar color represents idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$).




Figure 2.35: Flexible wages: model distribution of consumption. Bar color represents idiosyncratic productivity, from **lowest** ($e = 0.15$) to **highest** ($e = 6.6$).

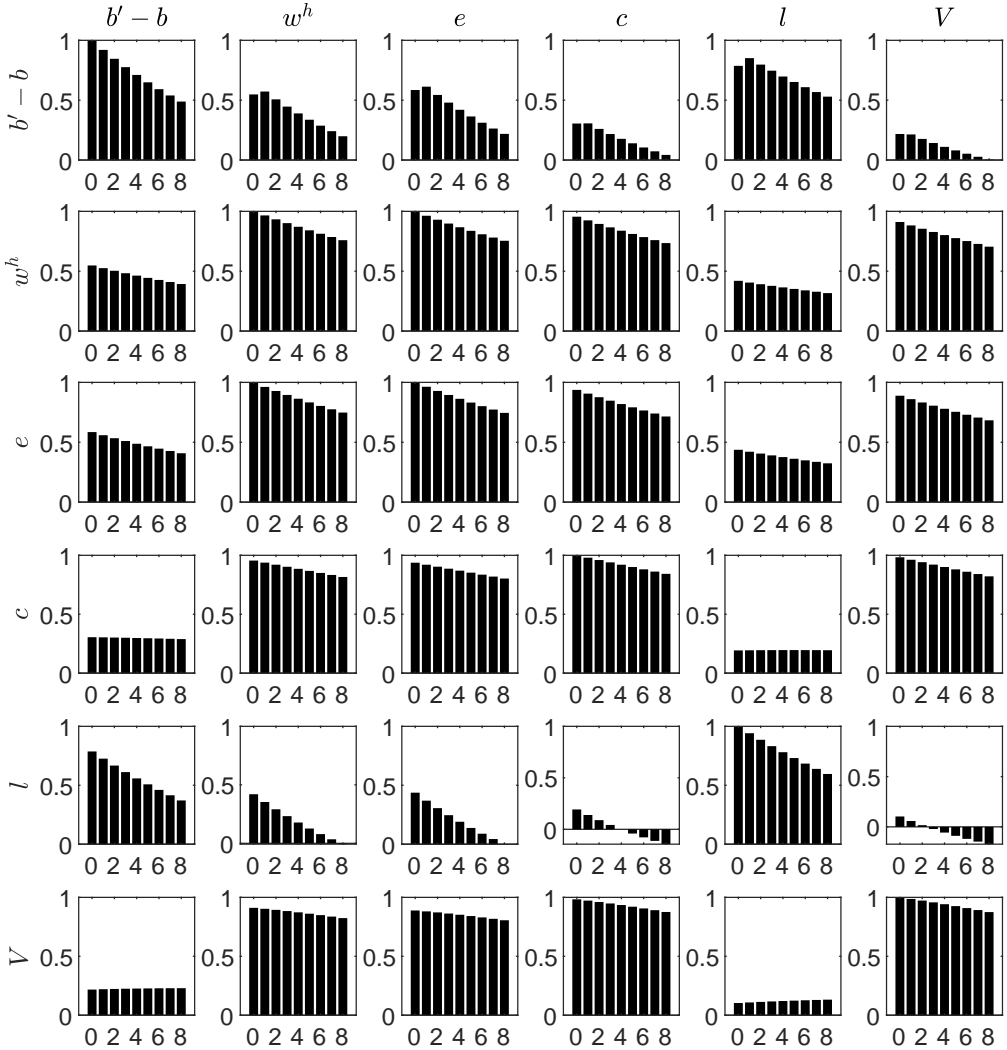


Figure 2.36: Flexible wages: time correlations of individual variables in version 2. Panel rows and columns correspond to labeled household-level variables. The horizontal axis in each panel shows lag/lead periods, while the vertical axis shows the correlation coefficient: $\text{corr}(\text{column}_t, \text{row}_{t+k})$, where k is the value on the horizontal axis. As an example, the panel in row 5 (l), column 4 (c) shows that $\text{corr}(c_t, l_{t+8}) < 0$, while the panel in row 4 (c), column 5 (l) shows that $\text{corr}(l_t, c_{t+8}) > 0$.

2.A.5 Additional figures

This section presents some additional figures for the baseline model.




Figure 2.37: Actual labor supply as share of demand, for wage-updating households ($f = 1$). Grey scale shows the ratio $\frac{g^l(b, w_{-1}, e, 1)}{n^*(b, w_{-1}, e, 1)}$, where $n^*(b, w_{-1}, e, 1)$ is given by equation (2.30). Blue dots show the distribution of households, and dot size represents the concentration of households. Axes limits have been chosen to exclude at most 0.5 percent of total households at each tail of each dimension.

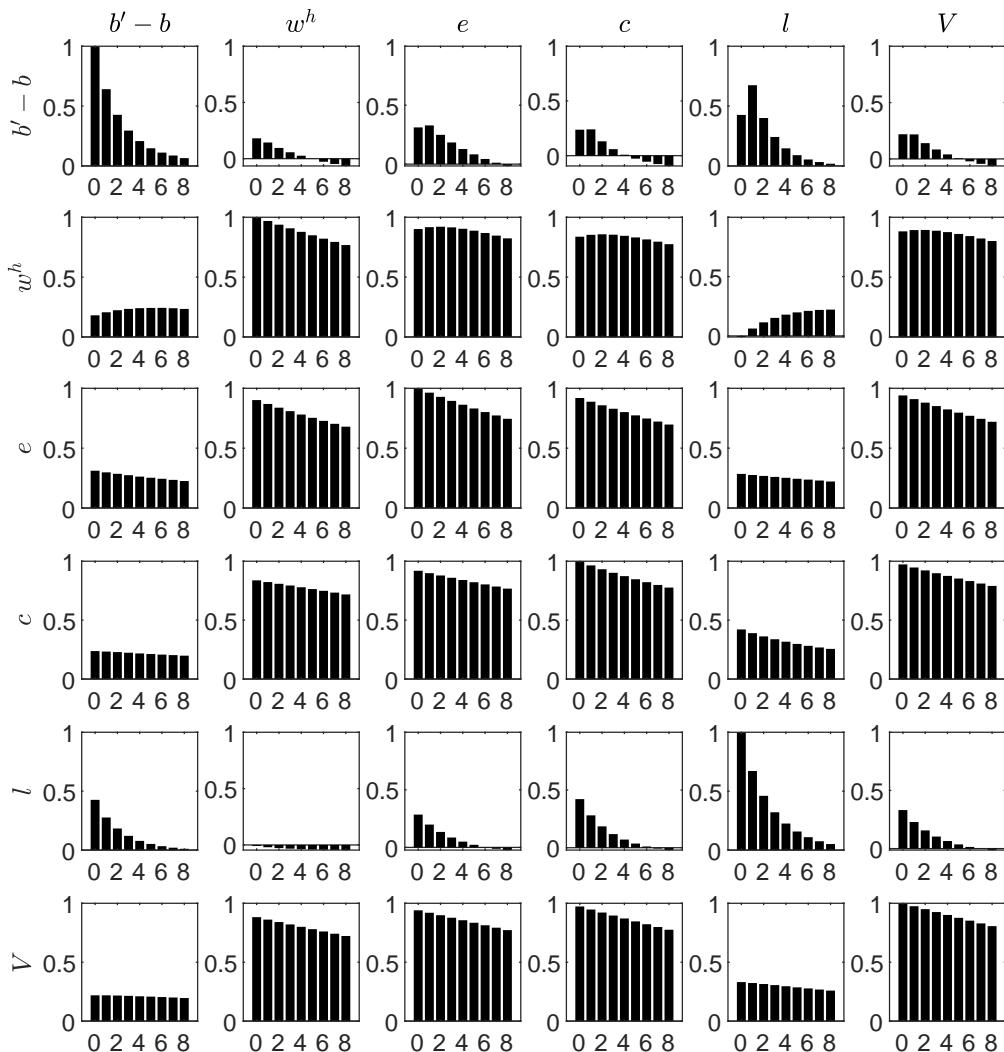


Figure 2.38: Time correlations of individual variables in version 2. Panel rows and columns correspond to labeled household-level variables. The horizontal axis in each panel shows lag/lead periods, while the vertical axis shows the correlation coefficient: $\text{corr}(\text{column}_t, \text{row}_{t+k})$, where k is the value on the horizontal axis. As an example, the panel in row 2 (w^h), column 5 (l) shows that $\text{corr}(l_t, w_{t+6}^h) > 0$, while the panel in row 5 (l), column 2 (w^h) shows that $\text{corr}(w_t^h, l_{t+6}) < 0$.

2.A.6 PSID (2017) data

This section lists the variables from PSID (2017) that have been used in section 2.4.

Table 2.4: PSID 2017 variables used

Variable	PSID code	Comment
Wealth, b	ER71483	Wealth excluding equity
Wage, w	<i>Added</i> ER66217 ER66211/(ER66170×ER66172)	Hourly regular wage rate <i>Yearly salary</i> divided by <i>hours worked</i> (see below)
Labor supply, l	ER66170×ER66172	Product of <i>weeks employed</i> and <i>average hours per week</i>
Earnings, wl	ER67046	Earnings from wages or salaries
Consumption, c	<i>Added</i> ER71487 ER71491 ER71497 ER71503 ER71515 ER71516 ER71517 ER71522 ER71523 ER71524 ER71525 ER71526 ER71527 <i>Subtracted</i> ER71492	Food Housing Utility Transportation Education Childcare Health care Computing Household repairs Household furnishing Clothing Trips Other recreation Mortgage

Chapter 3

How big is the toolbox of a central banker?

Managing expectations with policy-rate forecasts: Evidence from Sweden*

*An earlier version of this paper is number 339 in the Riksbank Working Paper Series. I would like to thank Jan Alsterlind, Jens Iversen and Ulf Söderström for detailed comments on an earlier draft, as well as valuable discussions. I would also like to thank Paolo Bonomolo, Henrik Eriksson, David Kjellberg, Per Krusell, Jesper Lindé, Jonna Olsson, Lars EO Svensson, David Vestin, Karl Walentin, and seminar participants at Universitat Autònoma de Barcelona, IIES and the Riksbank for valuable input and discussions, and Lina Fransson for great help with data collection.

3.1 Introduction

Monetary policy conducted by central banks affect the financial conditions, and hence potentially the lives, of most people throughout the world. The ultimate goal of any modern central bank is to create price stability; may it be prices of commodities, currencies or something else. Although the textbook means of achieving the target are either via money supply, or via controlling the short-term nominal interest rate, the real-world toolbox of central banks is somewhat more complex. As an example, the Federal Reserve has conducted several programs of *Quantitative Easing* (QE), i.e., large-scale asset purchases, since the onset of the financial crisis in 2008. This paper addresses another tool that might potentially be useful for achieving the target of a central bank.

Over the past decades, there has been a trend towards more transparency of both the decisions of central banks, and the motivating analysis behind the decisions; see, e.g., Dincer and Eichengreen (2007). One such step towards transparency is that a few central banks in developed countries have begun to not only announce the level of the policy rate but also the intended future development of the policy rate, a policy-rate path, beginning with the *Reserve Bank of New Zealand* in 1997.¹

One of the foundations of economic research is that expectations of the future matters for decisions taken today. Hence, using communication to affect the agents' beliefs about the future is one way for a central bank to steer the economy and achieve its target. One specific type, where the central bank communicates its own intended future actions, has broadly been labeled *forward guidance*. Besides providing transparency, the publication of a policy-rate path is also one way for a central bank to conduct forward guidance; see for example Archer (2005) and Ingves (2007). The main question I address in this paper is if forward guidance through the publication of a policy-rate path is an effective tool for a central bank. It is only lately that enough data is available to approach this question quantitatively

I perform an event study of the impact on Swedish market expectations of surprises in the Riksbank's announcements. An announcement consists of at least three parts: a policy-rate decision, a policy-rate path for the

¹The other central banks announcing policy-rate paths are *Norges Bank* (Norway, 2005), *Sveriges Riksbank* (Sweden, 2007) and the *Czech National Bank* (2010).

future, and a report analyzing the current economic situation, including forecasts of other macroeconomic variables. I try to distinguish the effect of surprises in the repo-rate path. The identification relies in part on high-frequency data around announcements, and I use movements of forward rates on different horizons to measure the impact on expectations.

My main finding is that a surprise in the policy-rate path published by the Riksbank does move market the expectations of the future policy rate. However, the effect is not one-for-one, and only significant up to around a year and a half, which is shorter than the forecast horizon. The results are robust to a number of variations in my measurements, which is particularly important because the main variables of interest cannot be directly observed, and might hence suffer from measurement errors.

The analysis relates closely to other studies, mainly using New-Zealand data. Moessner and Nelson (2008) conclude that the policy-rate paths published by the Reserve Bank of New Zealand impact market prices, and that they do not impair the functioning of the market. I mainly deviate in the definition of the market expectation of a published policy-rate path, just prior to the announcement. Ferrero and Secchi (2009) also use New-Zealand data to address the same question as I do in this paper. My approach is similar to theirs, although I use data from a different country, but it differs in how to proxy the main explanatory variables, the set of controls included, and how to check for robustness. Finally, Detmers and Nautz (2012) analyze how the market impact in New Zealand has changed since the financial crisis. They follow Moessner and Nelson (2008) in the definition of the market expectation of a published policy-rate path prior to the announcement, and hence this is where I mainly deviate. A study addressing similar questions as I do, with similar methods, using Swedish data, is Iversen and Tysklind (2017), which studies the impact of policy-rate path surprises on different financial prices. In general, my results are in line with the main results of all these studies.

The structure of the rest of the paper is that section 3.2 presents the econometric approach I take, and discusses the data in detail. Section 3.3 presents the main results with some alternative views on the data, and also discusses the robustness exercises I perform. Finally, section 3.4 summarizes and draws conclusions from the results.

3.2 Method

This section describes the model used to analyze the question of interest. It also describes the data used in the estimations in detail, and some assumptions that have been made and in some cases relaxed. There is also a discussion of some potential problems with the analysis.

3.2.1 Econometric approach

In the baseline analysis, I use a regression approach in an attempt to quantify the impact of the Riksbank's announcement of a repo-rate path on the market expectation of the future repo rate. The regression equation is

$$Impact_{h,t} = \beta_h Surprise_{h,t} + \gamma_h X_{h,t} + \varepsilon_{h,t}, \quad (3.1)$$

where $Impact_{h,t}$ is the movement in market expectations of the repo rate h quarters into the future at an announcement of a new repo-rate path at time t , $Surprise_{h,t}$ is the surprise component of the announced repo-rate path, $X_{h,t}$ is a vector of controls (including a horizon-specific constant) and $\varepsilon_{h,t}$ is an error term.² In section 3.2.2, I discuss how to measure these variables. The coefficient of main interest is β_h , which measures how much the market expectations are affected by (the surprise element of) the announced repo-rate path.

The main identifying strategy is the use of high-frequency data, as is typically the case in event studies.³ However, there is still room for non-causal interpretation of β_h . One potential problem is that more information than just the policy-rate path is announced at the same time, e.g., the contemporaneous decision on the repo rate, and a monetary policy report containing the Riksbank's view on the current economic status, and future development. This motivates including the control variables $X_{h,t}$, which might improve the causal interpretation.

An alternative specification, following Detmers and Nautz (2012), is

$$Impact_{h,t} = \beta_h^S Surprise_{h,t} + \beta_h^A Anticipated_{h,t} + \gamma_h X_{h,t} + \varepsilon_{h,t}, \quad (3.2)$$

²Note that the coefficients in the equation are indexed by the horizon. There is one equation, and one regression, per horizon. Some control variables might be common for all horizons while others are horizon-specific.

³See Gürkaynak and Wright (2013) for an overview of the event-study methodology.

where $Anticipated_{h,t}$ is the expected change (by an announcement at t) of the repo-rate path (at horizon h quarters) since the last announcement by the Riksbank. The parameter β_h^A measures the effect of an adjustment of the repo-rate path that is fully expected by the market. Without any measurement errors, in accordance with the effective-markets hypothesis, one would expect this effect to be zero. Hence, one possible interpretation is that any deviation from $\beta_h^A = 0$ can be viewed as an indication that there are measurement errors present in the variables. It is important to include $Anticipated_{h,t}$ since it is difficult to measure the anticipated communication by the Riksbank, and this provides an indicator of the quality of the measure that I use. This is also in line with the findings of Kuttner (2001), although there may be other plausible interpretations as well.

It is far from obvious which is the best way to match the three variables $Impact_{h,t}$, $Surprise_{h,t}$ and $Anticipated_{h,t}$ to available data. Neither is it obvious what the relevant control variables to include in $X_{h,t}$ are. The main strategy is presented in section 3.2.2, together with a number of alternative specifications to ensure the robustness of the results.

Aside from the baseline analysis, I motivate the main question I try to answer by investigating whether the expectations of the future repo rate tend to change more on announcement days, i.e., days when the Riksbank announces a repo-rate decision, publishes a new forecast for the repo rate, and publishes a new monetary policy report.⁴ This is carried out by a simple regression of the kind

$$|Impact_{h,t}| = \alpha_h + \delta_h D_t^{Announcement} + \eta_h D_t^{Expiration} + \varepsilon_{h,t}, \quad (3.3)$$

where $|Impact_{h,t}|$ is the absolute value of the movement in market expectations of the repo rate h calendar quarters into the future at day t , $D_t^{Announcement} \in \{0, 1\}$ is a dummy variable indicating the days on which a new repo-rate path was announced, $D_t^{Expiration} \in \{0, 1\}$ is a dummy variable indicating the expiration dates of the contracts used to measure market expectations, and $\varepsilon_{h,t}$ is an error term. The nature of $D_t^{Expiration}$ is technical, and explained in more detail below.

The coefficient of interest is δ_h , which measures to what extent market

⁴On some occasions in my sample, a full report is not published, but rather a less extensive update.

expectations tend to move more, in any direction, on announcement days. If δ_h is significantly larger than zero, this is evidence that the market expectation of the future repo rate at horizon h is affected by the announcement. If $\delta_h = 0$ cannot be rejected, it indicates one of two things: either the surprise elements of announced repo-rate paths do not impact the market pricing, or the announced repo-rate paths in the sample are well in line with the pre-announcement expectations.

3.2.2 Data

Most of the variables introduced in section 3.2.1 are not directly observable. Hence, I need to proxy for them, which will result in potential measurement errors. This section presents and discusses the data used in the empirical analysis.

The Riksbank publishes a repo-rate forecast, or repo-rate path, six times per year, and has done so since the beginning of 2007. The path consists of quarterly averages for the forecast of the repo rate, and typically has a horizon of 12 quarters. It is announced together with a repo-rate decision and a monetary policy report or update, containing forecasts for a number of macroeconomic variables along with an analysis of the current economic situation. The announcement dates are known beforehand.⁵

Market data As a measure of the market expectations of the future repo rate, I use *Forward Rate Agreements* (FRAs) adjusted for a time-independent premium.⁶ These are futures contracts on an underlying

⁵Typically, the announcement dates are known far in advance. One exception is the announcement on December 4 2008. On December 1, a press release made clear that this announcement had been rescheduled from December 17. The press release does not mention the reason, but it should have been clear to everyone of interest at that point that the rescheduling was due to developments in the ongoing financial crisis.

⁶It is important to distinguish between expected *communication* and expected *action* by the central bank. The FRAs, compensated for premia, are used as measures of the expected action, but do not provide any information on which repo-rate path the market expects the Riksbank to communicate.

3-month interbank rate, STIBOR.^{7,8} The usage of such contracts as a measure of market expectations of the future policy rate is in line with the existing literature; see, for example Gürkaynak et al. (2007), Moessner and Nelson (2008) and Ferrero and Secchi (2009). This is also how the Riksbank measures expectations of future monetary policy in its own analysis; see Sveriges Riksbank (2013). However, there is need for some caution here. It may well be that the FRAs are subject to a time-varying premium and hence do not directly reflect the expected repo rate.⁹ There are methods for estimating such time-varying premiums, but different methods tend to give substantially different and uncertain results, so in the main analysis I keep the assumption that the premium is fixed. This assumption is relaxed in section 3.3.2.

The FRA contracts expire two bank days prior to the third Wednesday of the last month in each quarter, i.e., approximately two weeks before the beginning of a new calendar quarter. Hence, if compensated for premia, the FRAs are good measures of the expectations of the average overnight interbank rate in a calendar quarter by the expectations hypothesis.¹⁰ Furthermore, the overnight interbank rate is very well correlated with the repo rate.

Figure 3.1 shows the outcome of the repo rate, together with one forecast per year by the Riksbank and corresponding expectations according to the FRAs, for the period of interest.¹¹ Note that the Riksbank and the market have agreed about the future development in some periods, and dis-

⁷The difference between the 3-month STIBOR and the repo rate has been rather constant and on average 0.3 percent over the period of interest. Hence, the FRA quotes are adjusted down by 0.3 percentage points in order to better reflect the expected repo rate.

⁸Also RIBA futures, similar to the FRAs but with the repo rate as the underlying rate, are traded. These are available from 2009, not for as many horizons, and they are traded in smaller volumes than the FRAs, and are therefore not used in the main analysis. However, the main results are robust to replacing the FRAs with RIBA when possible.

⁹Such a premium could reflect the pricing of many things, including liquidity risk, and credit risk.

¹⁰To get an even better match with calendar quarters, I assign weights of $\frac{5}{6}$ and $\frac{1}{6}$, respectively, to two consecutive FRA contracts, following Detmers and Nautz (2012). An alternative would be to use the method suggested by Nelson and Siegel (1987) or the extended version in Svensson (1994).

¹¹When all forecasts are included, the figure becomes difficult to comprehend; see figure 3.4 in appendix 3.A.1.




Figure 3.1: Outcome of the **repo rate**, and selected forecasts by the Riksbank (black dashed) and the market (grey dashed) as quarterly averages at announcement dates. The full set of forecasts is available in figure 3.4 in appendix 3.A.1.

agreed in others. There are several plausible explanations for the periods of disagreement; the information available to the market might differ from that available to the Riksbank, the premia of the FRAs might change, the view of a steady-state level of interest rate might differ, different models for the economy might also be used, and the Riksbank's communication might be viewed as non-credible by the market. Probably all of the above are true to some extent, and there might also be other explanations. The reasons for the historical disagreement are both important and interesting per se, but it is not the aim of this paper to explain why it has arisen. For an analysis of the consequences of differences between market rates and communicated policy-rate paths, see De Graeve and Iversen (2015).

The most striking period of disagreement is perhaps in 2011, when the Riksbank projected the repo rate to continue to increase at a rapid pace, while the market expected the repo rate to increase at a much slower pace

or even decrease. As can be seen in the figure, the market turned out to be right ex post. This episode is discussed in more detail in Svensson (2015).

The FRA quotes are observed for horizons 1 to 12 quarters. More formally, we have the following relationship between the FRAs and the expected future repo rate:

$$FRA_{h,t} = E_t \left[i_{t+h}^{repo} \right] + \zeta_{h,t}, \quad (3.4)$$

where $FRA_{h,t}$ is the observed futures rate for horizon h at time t , $E_t \left[i_{t+h}^{repo} \right]$ is the market expectation of the repo rate h calendar quarters into the future, and $\zeta_{h,t}$ is the premium for horizon h .¹² In principle, t applies to any point in time. However, I consider t as occurring just after an announcement by the Riksbank. Under the assumption that the premium is not affected by the announcement, i.e., $\zeta_{h,t} = \zeta_{h,t-\epsilon}$ for all $h \in \{1, 2, \dots, 12\}$, it is straightforward to define the data version of the dependent variable as

$$\begin{aligned} Impact_{h,t} &= FRA_{h,t} - FRA_{h,t-\epsilon} \\ &= E_t \left[i_{t+h}^{repo} \right] - E_{t-\epsilon} \left[i_{t+h}^{repo} \right], \end{aligned} \quad (3.5)$$

where $t - \epsilon$ refers to a point in time just prior to the announcement. The assumption that the premium is unaffected by the announcement is possibly strong, but difficult to overcome.¹³ If this assumption is too strong, $Impact_{h,t}$ cannot be interpreted as reflecting the market expectation of the future policy rate, but rather market rates in general. Although this makes the analysis less specific, it is still of interest as a monetary-policy tool.

What is meant by “just after” and “just prior to” an announcement?

¹²Björk (2004) shows that even in a risk neutral setting, the expectations hypothesis need not hold. However, most central banks, including the Riksbank, rely on the expectations hypothesis adjusted for premia in this type of analysis, so I follow their example.

¹³In the spring of 2015, towards the end of my sample, the Riksbank started a program of quantitative easing (purchase of government bonds). The announcements of the measures coincided with regular Riksbank announcements, and were partly aimed at affecting the market risk premia; see Sveriges Riksbank (2015) and the references therein. This makes these observations extra prone to violate this assumption.

Is the difference one day, hour, minute, second or something else? In this study, I use end-of-day quotes, so ϵ corresponds to one day. This is common in the literature, see, e.g., Ehrmann and Fratzscher (2004) and Moessner and Nelson (2008), and has the advantage that the market has time to fully incorporate the new information announced by the Riksbank in the prices used. However, a drawback is that the prices will also be influenced by other news and information arriving within the same day.¹⁴ An alternative would be to use intra-day quotes, as in Gürkaynak et al. (2005) and advocated by Winkelmann (2010). Choosing this approach instead does not seem to affect the results to any considerable extent.¹⁵ In section 3.3.2, I also apply a method aimed at controlling for other news arriving within the same days.

Anticipated changes of the repo-rate path I now turn to the variable $Anticipated_{h,t}$ in equation (3.2). This variable is the market's expected change in the repo-rate path between two consecutive Riksbank announcements. Another way of putting it is that the repo-rate path that the market expects the Riksbank to announce, just prior to the announcement, is the sum of the last published repo-rate path and the variable $Anticipated_{h,t}$. The idea is that the market uses all available information — that which was previously announced by the Riksbank and the new information that has arrived since the last announcement — to predict the content of the new Riksbank announcement.¹⁶ Some alternative views on this variable are discussed in section 3.3.2.

In the baseline case, I assume that the market expects the Riksbank to update its view on the repo-rate path in the same way that the market itself updated its view since the last announcement. In this case, I define

$$Anticipated_{h,t} = FRA_{h,t-\epsilon} - FRA_{h,t_p}, \quad (3.6)$$

¹⁴The best case scenario is that the additional information is independent of the announcement by the Riksbank, in which case it only adds noise, and hence decreases the power of the estimates without inflicting any bias.

¹⁵I do not have access to intra-day quotes for the entire period of interest or all horizons, but combining daily data with what intra-day data I have only results in minor changes to the results.

¹⁶This is similar to what Archer (2005) does. Winkelmann (2010) takes another approach, using jumps in medium- to long-term rates on announcement days to identify anticipated and unanticipated surprises in the announced path.

where $FRA_{h,t-\epsilon}$ is the futures rate of horizon h just prior to the announcement at time t , as before, and FRA_{h,t_p} is the futures rate just after the previous announcement by the Riksbank.¹⁷ One implication of this definition is that I assume that the market expects any discrepancy between the market expectation and the forecast in the Riksbank's last announced path to remain unchanged in the coming announcement, given time-fixed premia.

With this definition of the anticipated change of the repo-rate path, the surprise, or unanticipated change, is defined as the difference between the actual and the anticipated change;

$$Surprise_{h,t} = \left(Path_{h,t}^{RB} - Path_{h,t_p}^{RB} \right) - Anticipated_{h,t}, \quad (3.7)$$

where $Path_{h,t}^{RB}$ is the repo-rate path for horizon h , announced by the Riksbank at time t and $Path_{h,t_p}^{RB}$ is the previously announced path for the corresponding calendar quarter.

As noted above, defining $Anticipated_{h,t}$ and $Surprise_{h,t}$ by equations (3.6) and (3.7) assumes that the market expects the Riksbank to update its views in the same way that the market has updated its views. This need of course not be the case. Alternatively, the anticipated and unanticipated changes in the repo-rate path can be defined as the explained parts and residuals of the following regressions (one per horizon), respectively:

$$Path_{h,t}^{RB} = \alpha_h + \mu_h^M FRA_{h,t-\epsilon} + \mu_h^{M_p} FRA_{h,t_p} + \mu_h^P Path_{h,t_p}^{RB} + Surprise_{h,t}. \quad (3.8)$$

After running these regressions, it is natural to define the anticipated change in the repo-rate path since the last announcement as

$$Anticipated_{h,t} = \alpha_h + \mu_h^M \cdot FRA_{h,t-\epsilon} + \mu_h^{M_p} \cdot FRA_{h,t_p} + \left(\mu_h^P - 1 \right) \cdot Path_{h,t_p}^{RB}. \quad (3.9)$$

The explained part of the right-hand side of equation (3.8) contains the level, and change since the last announcement, of the market rates as well as the previously announced path by the Riksbank. This way

¹⁷The Riksbank publishes a new repo-rate path six times per year, so on average the previous announcement was made two months earlier. However, the intervals between meetings differ over the year.

of defining the anticipated changes and surprises through regression is similar to what Moessner and Nelson (2008) suggest and to what Ferrero and Secchi (2009) do. Note that the simpler definition in equation (3.6) corresponds to the case $\alpha_h = 0$, $\mu_h^M = \mu_h^P = 1$ and $\mu_h^{M_p} = -1$ in equation (3.9).

Regardless of whether *Anticipated*_{*h,t*} and *Surprise*_{*h,t*} are defined by equations (3.6, 3.7), or by equations (3.8, 3.9), there is an obvious risk of correlation between the two. My variable of interest is *Surprise*_{*h,t*}, so if *Anticipated*_{*h,t*} is also correlated with the dependent variable *Impact*_{*h,t*}, it should be included on the right-hand side in the main analysis to avoid omitted-variable bias. I.e., if that is the case, I should use equation (3.2) rather than (3.1).

As should be clear from above, the measure of *Surprise*_{*h,t*} is uncertain and may well contain measurement errors. If that is the case, the regression equations (3.1) and (3.2) will suffer from regression dilution, also known as attenuation bias, and the estimates of β_h and β_h^S will be biased towards zero. I.e., the true coefficients may in fact be larger in size than suggested by the results in section 3.3.1.

Control variables Next I pay some attention to the potential vector of control variables, $X_{h,t}$ in equations (3.1) and (3.2). There might be two reasons to include control variables. The first, and most important, reason would be to prevent an omitted-variables bias. It is known that omitting any independent variable that is correlated with the dependent variable and the independent variable of interest will bias the coefficient of interest; see for example Angrist and Pischke (2008). The direction of the bias depends on the correlations in question and is in general not known. Hence, I include independent variables that I suspect can have an explanatory value for both the *Surprise*_{*h,t*} and *Impact*_{*h,t*} variables. The second reason to include more independent variables is that there might be variables that are not correlated with *Surprise*_{*h,t*}, but when interacted with *Surprise*_{*h,t*} explains the dependent variables. Including more independent variables comes at the cost of lower power of the results. This is especially notable when the sample size is small, as in this case.

Section 3.3 presents results with different specifications of the control vector. The following variables are included mainly to prevent omitted-

variables bias:

Surprise in decision: A measure of the surprise in the repo-rate decision.¹⁸ One can suspect that this correlates very well with surprises along the repo-rate path. Details on how this measure is constructed are found in appendix 3.A.3.

Surprises on other horizons: The average surprise for all horizons *except* the one the regression concerns.¹⁹ If the market pays no attention to the time precision of the repo-rate path, and only reacts to movements in the entire path for all horizons, it will be captured by this term rather than in β_h or β_h^S .

Dummy, effective lower bound: The Riksbank has, on some occasions, communicated that lowering the repo rate further might result in technical difficulties due to an effective lower bound. Such a lower bound might affect both the communication by the Riksbank and the interpretation by the market.²⁰

Disagreement: As can be seen in figure 3.4, there have been periods when the level of disagreement between the Riksbank's forecasts and the market expectations has been both high (with a positive and negative sign) and low. It might be that the level of disagreement affects the reasoning by the Riksbank as well as the market's reaction to the Riksbank communication. Therefore, I include a backward-looking one-year moving average of the disagreement (average for all horizons) between the Riksbank's forecast and the market pricing, at the time of the last announcement. This proxy measure of disagreement is de-meanned, and shown in appendix figure 3.5.

¹⁸A similar control variable is also used in Ferrero and Secchi (2009), although constructed slightly differently. None of the other covariates listed here seem to be present in the literature addressing this question.

¹⁹Since two consecutive announcements are often made in two different quarters, the repo-rate path from the previous announcement only covers the 11 first quarters of the new announced path. There are not enough data points where this is not the case to analyze the surprise in the 12 quarter horizon. Hence, this control variable is the average surprise in horizon 1–11 quarters, except the horizon that the regression concerns h .

²⁰The communication whether the interest rate is on the effective lower bound, or close enough to affect the monetary policy, is not always clear. I regard a lower bound to be effective for the period July 2009 – April 2010 and July 2014 – July 2015.

As mentioned above, an announcement by the Riksbank contains more than just a repo-rate decision and path. Aside from the list presented above, it would be desirable to also include controls for the market surprise in the remaining parts of the announcement, i.e., forecasts for other macroeconomic variables and an analysis of the current economic situation. However, it is very difficult to find measures of such expectations.

The following independent variables are included mainly because I am interested in the interaction effect with the surprise:

Dummy, surprise decreases the disagreement: Along the line of thought that the level of agreement between the Riksbank forecast and market expectations might affect the impact on market expectations, I include a dummy for whether the surprise works to increase or decrease the disagreement. A surprise that decreases the disagreement might be viewed as more credible by the market than a surprise that increases the disagreement further. When included, the dummy variable is de-measured and interacted with the surprise.

Announcement timing: The monetary policy meetings of the Riksbank are held at different times within the quarter. Consequently, at some meetings the one quarter ahead forecast refers to a quarter beginning only a few days later, while at other meetings the one quarter ahead forecast refers to a quarter beginning almost three months from the meeting. A reasonable hypothesis is that a repo-rate path announced closer to the beginning of a new quarter will be viewed as more credible, and hence a surprise in such a meeting could have a larger impact on the market expectations, particularly for short horizons. To capture this, the fraction of the quarter still remaining is included as a control, demeaned and interacted with the surprise.

Using the full set of control variables, equation (3.2) can be written as

$$\begin{aligned}
 Impact_{h,t} = & \beta_h^S Surprise_{h,t} + \beta_h^A Anticipated_{h,t} + \gamma_{0,h} + \gamma_{1,h} Surprise_{0,t} \\
 & + \gamma_{2,h} \frac{1}{10} \sum_{j \neq h} Surprise_{j,t} + \gamma_{3,h} D_t^{ELB} + \gamma_{4,h} \widetilde{Disagreement}_t \\
 & + \gamma_{5,h} \widetilde{D}_{h,t}^{Closing} Surprise_{h,t} + \gamma_{6,h} \widetilde{Fraction}_t Surprise_{h,t} + \varepsilon_{h,t},
 \end{aligned} \tag{3.10}$$

where \sim denotes the deviation from the horizon-specific mean. Note that under this specification, the effect of $Surprise_{h,t}$ on $Impact_{h,t}$ is not entirely captured by β_h^S , but rather we have

$$\frac{\partial Impact_{h,t}}{\partial Surprise_{h,t}} = \beta_h^S + \gamma_{5,h} \widetilde{D}_{h,t}^{Closing} + \gamma_{6,h} \widetilde{Fraction}_t. \tag{3.11}$$

In other words, β_h^S is a good approximation of the effect of $Surprise_{h,t}$ on $Impact_{h,t}$ if the remaining two terms in (3.11) are well approximated by zero, due to the coefficient being small, the independent variables being small, or both. In general, more than just the estimates of β_h^S must be considered. I discuss this further in section 3.3.1, together with the results.

3.3 Results

Before turning to the main analysis, I briefly motivate why it is worth digging into the questions I address in this paper. Table 3.1 reports the regression results of equation (3.3), where I have used end-of-day FRA quotes for all trading days between February 2005 and July 2015.²¹ Note that one regression is run per horizon h . The coefficient of interest is δ_h , which is interpreted as the extra movement of FRA quotes on days when a new repo-rate path is announced, in total 49 days in the sample. We note that δ_h is significantly larger than zero for all horizons, indicating that the repo-rate expectations tend to move more on announcement days than

²¹The standard errors reported in regression tables throughout the paper are heteroscedasticity-consistent; see, e.g., Angrist and Pischke (2008).

non-announcement days, still under the assumption that the premium is approximately unaffected by the announcement. Comparing the size of δ_h with the size of α_h , which captures the average movement of the FRA quote on non-announcement trading days, we see that the effect is not only statistically significant but also economically very significant, especially for shorter horizons.

The variable $D_t^{Expiration}$ in equation (3.3) is a dummy for the expiration dates, i.e., the dates when a FRA contract switches from referring to one calendar quarter to the next. This must, of course, be accounted for. The interpretation of η_h is hence the average difference of the FRA price between two consecutive horizon quarters, at the expiration dates. This might capture both the premia and the expectations of future short rates, but mainly captures how the expected future repo rate varies with the forecast horizon. There are in total 42 expiration days in the sample.

The overall conclusion from this introductory analysis is that there is excess volatility in the FRAs on announcement days, suggesting that the market is influenced by the information released by the Riksbank at these announcements. This motivates a further investigation of the announcements in general and, more specifically, the questions I address in this paper. I now proceed to presenting my main findings.

3.3.1 Main results

I begin by investigating the very simplest case, and thereafter add complexity in steps. The very simplest case is to run the regressions, one per horizon, in equation (3.1) without any control variables, i.e., $X_{h,t}$ is only a vector of ones so that γ_h is an intercept. I also use the simpler definition of $Surprise_{h,t}$, i.e., it is defined by equations (3.6) and (3.7). The results of these regressions are presented in table 3.2.

Table 3.1: Regression results of equation (3.3)

$h =$	1	2	3	4	5	6	7	8	9	10	11	12
α_h	0.02*** (0.00)	0.02*** (0.00)	0.03*** (0.00)	0.03*** (0.00)	0.04*** (0.00)	0.04*** (0.00)	0.04*** (0.00)	0.04*** (0.00)	0.03*** (0.00)	0.04*** (0.00)	0.04*** (0.00)	0.03*** (0.00)
δ_h	0.07*** (0.02)	0.05*** (0.01)	0.05*** (0.01)	0.04*** (0.01)	0.03*** (0.01)	0.02*** (0.01)	0.02*** (0.01)	0.02*** (0.01)	0.02*** (0.01)	0.02*** (0.01)	0.02*** (0.01)	0.02*** (0.01)
η_h	0.18*** (0.03)	0.14*** (0.02)	0.09*** (0.02)	0.10*** (0.02)	0.09*** (0.02)	0.08*** (0.02)	0.07*** (0.01)	0.06*** (0.01)	0.05*** (0.01)	0.04*** (0.01)	0.03*** (0.01)	0.02*** (0.01)
R^2	0.25	0.16	0.10	0.10	0.06	0.06	0.02	0.01	0.03	0.01	0.01	0.01
p_F	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Obs.	2604	2604	2604	2597	2600	2604	2597	2604	2604	2604	2604	2598

Equation: $|Impact_{h,t}| = \alpha_h + \delta_h D_t^{Announcement} + \eta_h D_t^{Expiration} + \varepsilon_{h,t}$.

Data sources: Bloomberg, Nasdaq OMX and the Riksbank.

Note: h refers to the horizon in quarters. ***, ** and * refer to significance at the 1%, 5% and 10% levels, respectively.

Table 3.2: Regression results of equation (3.1)

$h =$	1	2	3	4	5	6	7	8	9	10	11
β_h	0.55*** (0.16)	0.25*** (0.11)	0.14** (0.07)	0.10*** (0.05)	0.08*** (0.04)	0.06* (0.03)	0.06* (0.03)	0.03 (0.03)	0.01 (0.02)	-0.00 (0.02)	-0.01 (0.02)
γ_h	0.01 (0.01)	0.00 (0.02)	-0.01 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.00 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.00 (0.01)	0.01 (0.01)
R^2	0.53	0.16	0.13	0.13	0.13	0.11	0.09	0.03	0.01	0.00	0.00
p_F	0.00	0.01	0.01	0.01	0.01	0.02	0.05	0.24	0.56	0.98	0.68
Obs.	41	46	47	48	48	47	46	46	46	47	40

Equation: $Impact_{h,t} = \beta_h Surprise_{h,t} + \gamma_h + \varepsilon_{h,t}$.
Data sources: Bloomberg, Nasdaq OMX and the Riksbank.
Note: h refers to the horizon in quarters. ***, ** and * refer to significance at the 1%, 5% and 10% levels, respectively.

The estimates of β_h suggest that a surprise in the repo-rate path announced by the Riksbank might have a significant effect on market expectations up to a horizon of about 5–7 quarters. However, the suggested effect is quite small for horizons beyond 1 or perhaps 2 quarters.²² Note in table 3.2 that the coefficient of determination, R^2 , is low for horizons beyond 1 quarter, suggesting that this model does not perform well in explaining how market expectations are updated on announcement days.

The results of these first simple regressions suggest that the effect we are looking for, the ability of the Riksbank to affect market expectations with the repo-rate path, is present. However, the results should be interpreted with caution. There is reason to believe that the estimates of β_h may be biased, partly since the measure of the surprise in the announced repo-rate paths might be bad, and partly because there might be other explanatory variables that are correlated with both the impact on expectations and the surprise part of the repo-rate paths. Next, I handle these potential problems.

In the next step, I also add the anticipated change in the repo-rate path to the analysis. This should give a hint of the quality of our measure of the surprise part of the announced paths. Table 3.3 shows the results of the regressions in equation (3.2), still using equations (3.6) and (3.7) to define the anticipated change and the surprise. Note that this leads to a substantial increase in the R^2 , and for most horizons also in the estimates of the coefficient for the surprise, $\hat{\beta}_h^S$, compared to the case where the anticipated change is not included. This is a symptom that an omitted-variables bias was present but has now been overcome to some extent. The significant effect of the surprise now stretches up to a 9-to-10-quarter horizon. However, also note that the estimates of the coefficient for the anticipated change, β_h^A , are significantly larger than zero for most horizons. As discussed in section 3.2.1, this might be an indication that the measure of the variable $Anticipated_{h,t}$ is bad, and consequently also the measure of $Surprise_{h,t}$.

²²The interpretation of, for instance, $\hat{\beta}_1 = 0.55$ is that a surprise of 100 basis points in the repo-rate path one quarter ahead should move the market expectations 55 basis points in the same direction for the one-quarter horizon. Although the estimate for $\hat{\beta}_7 = 0.06$ is significantly larger than zero in a statistical sense, a movement of market expectations of 6 basis points in response to a 100 basis point surprise must be regarded as close to zero impact.

Table 3.3: Regression results of equation (3.2)

$h =$	1	2	3	4	5	6	7	8	9	10	11
β_h^S	0.47*** (0.09)	0.32*** (0.08)	0.22*** (0.07)	0.18*** (0.05)	0.14*** (0.03)	0.13*** (0.03)	0.13*** (0.03)	0.12*** (0.03)	0.09** (0.03)	0.06* (0.04)	0.05 (0.04)
β_h^A	0.22*** (0.06)	0.13*** (0.04)	0.11** (0.05)	0.12*** (0.04)	0.11*** (0.03)	0.12*** (0.03)	0.11*** (0.03)	0.11*** (0.03)	0.09*** (0.03)	0.07* (0.04)	0.07 (0.05)
$\gamma_{0,h}$	0.02* (0.01)	0.02 (0.01)	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.02** (0.01)	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.02 (0.01)
R^2	0.72	0.31	0.28	0.35	0.38	0.43	0.39	0.31	0.21	0.09	0.08
p_F	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.13	0.22
Obs.	41	46	47	48	48	47	46	46	46	47	40

Equation: $Impact_{h,t} = \beta_h^S Surprise_{h,t} + \beta_h^A Anticipated_{h,t} + \gamma_{0,h} + \varepsilon_{h,t}$.
Data sources: Bloomberg, Nasdaq OMX and the Riksbank.
Note: h refers to the horizon in quarters. ***, ** and * refer to significance at the 1%, 5% and 10% levels, respectively.

Given the potential problem identified above, the next natural step is to try to improve the measure of $Anticipated_{h,t}$ from the definition in equation (3.6). As described in section 3.2.1, one method, closely related to that suggested in Ferrero and Secchi (2009), is to define $Anticipated_{h,t}$ by equation (3.9), using the coefficients from regression equation (3.8). Note also that $Surprise_{h,t}$ is then defined as the unexplained part, or residual, of the same regression. Denoting the regression in (3.8) by *first stage* and the regression in (3.2) by *second stage*, the results are presented in table 3.4.

Let us consider the first stage. Recall that with $\alpha_h = 0$, $\mu_h^M = \mu_h^P = 1$, and $\mu_h^{M_p} = -1$, equations (3.6) and (3.9) are equivalent. It is apparent from table 3.4 that this is a bad assumption for all horizons beyond one quarter. Further, note that R^2 is high, indicating that the regressions of equation (3.8) capture the determination of $Anticipated_{h,t}$ quite well.

The second stage is presented in table 3.4 and in figure 3.2, where the estimates of β_h^S are illustrated with confidence intervals for each horizon h . The estimates of β_h^A are not as significantly different from zero as in table 3.3.²³ This also strengthens the hypothesis that equations (3.8) and (3.9) capture the variable $Anticipated_{h,t}$ well. Note also that for some quarters, the estimated impact of the surprise, $\hat{\beta}_h^S$, increases substantially as compared to table 3.3. The R^2 does also increase for some horizons, indicating that the regressions presented in table 3.4 fit the data better than those in table 3.3.

²³In fact, omitting $Anticipated_{h,t}$ from the second-stage equation does not change the results to any considerable extent.

Table 3.4: Regression results of equations (3.2) and (3.8)

$h =$	1	2	3	4	5	6	7	8	9	10	11
Second stage											
β_h^S	0.69*** (0.08)	0.49*** (0.12)	0.32*** (0.09)	0.21*** (0.06)	0.16*** (0.04)	0.13*** (0.04)	0.14*** (0.03)	0.14*** (0.03)	0.11*** (0.04)	0.08 (0.05)	0.06 (0.05)
β_h^A	0.19*** (0.04)	0.07 (0.04)	0.04 (0.05)	0.05 (0.04)	0.05 (0.03)	0.06* (0.04)	0.06* (0.03)	0.06* (0.03)	0.06 (0.04)	0.06 (0.05)	0.01 (0.05)
$\gamma_{0,h}$	-0.01 (0.01)	0.01 (0.01)	-0.00 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.00 (0.01)	-0.00 (0.01)	0.01 (0.01)	-0.00 (0.01)
R^2	0.82	0.41	0.37	0.36	0.37	0.34	0.36	0.33	0.21	0.09	0.04
p_F	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.14	0.44
Obs.	36	46	46	48	48	48	46	46	46	46	48
First stage											
α_h	-0.00 (0.03)	-0.01 (0.03)	-0.03 (0.04)	-0.07 (0.05)	-0.06 (0.07)	-0.04 (0.08)	0.05 (0.09)	0.09 (0.10)	0.11 (0.13)	0.09 (0.15)	0.08 (0.15)
μ_h^M	1.12 (0.10)	0.92 (0.06)	0.82** (0.07)	0.74*** (0.09)	0.68*** (0.10)	0.59*** (0.11)	0.59*** (0.10)	0.51*** (0.09)	0.45*** (0.08)	0.38*** (0.08)	0.35*** (0.07)
$\mu_h^{M_p}$	-0.90 (0.21)	-0.58*** (0.14)	-0.46*** (0.13)	-0.45*** (0.13)	-0.43*** (0.12)	-0.34*** (0.11)	-0.32*** (0.11)	-0.26*** (0.11)	-0.22*** (0.11)	-0.17*** (0.10)	-0.16*** (0.09)
μ_h^P	0.75 (0.19)	0.66** (0.15)	0.66*** (0.12)	0.74** (0.10)	0.78** (0.08)	0.77*** (0.07)	0.73*** (0.08)	0.74*** (0.08)	0.76** (0.09)	0.78** (0.09)	0.81** (0.08)
R^2	0.99	0.99	0.98	0.97	0.96	0.94	0.95	0.95	0.95	0.95	0.95
Obs.	36	46	46	48	48	48	46	46	46	46	48

First stage: $Path_{h,t}^{RB} = \alpha_h + \mu_h^M FRA_{h,t-\epsilon} + \mu_h^M FRA_{h,t_p} + \mu_h^P Path_{h,t_p}^{RB} + Surprise_{h,t}$.

Second stage: $Impact_{h,t} = \beta_h^S Surprise_{h,t} + \beta_h^A Anticipated_{h,t} + \gamma_{0,h} + \varepsilon_{h,t}$.

Data sources: Bloomberg, Nasdaq OMX and the Riksbank.

Note: h refers to the horizon in quarters. In the first-stage regression, significance levels refer to the difference from the reference levels (0, 1, -1, 1) for $(\alpha_h, \mu_h^M, \mu_h^{M_p}, \mu_h^P)$ respectively. ***, ** and * refer to significance at the 1%, 5% and 10% levels, respectively.




Figure 3.2: Estimates of β_h^S in equation (3.2), with 90%, 95% and 99% confidence intervals, for different horizons in quarters.

The results presented in table 3.4 and figure 3.2 may be viewed as the main results of this study. However, as mentioned above, there are still reasons to suspect bias in $\hat{\beta}_h^S$ due to omitted variables, and regression dilution. I also run the regressions including all control variables discussed in detail in section 3.2.2, i.e., the regression in equation (3.10). This is not only an attempt to decrease the omitted-variables bias, but also to study the interaction effects contained in equation (3.10). The full results are shown in table 3.5, and the estimates of β_h^S are illustrated, with uncertainty, in relation to the horizon h in figure 3.3. Note that the variables $Anticipated_{h,t}$ and $Surprise_{h,t}$ are still defined by the first-stage regressions of equation (3.8). However, nothing is changed in the first-stage regression from table 3.4, so table 3.5 only shows the second stage. It is not obvious whether table 3.5 and figure 3.3, or table 3.4 and figure 3.2, best illustrate the answer to the main question addressed in

this paper. Qualitatively there is a clear resemblance between the two, especially for shorter horizons, although the quantitative impact differs.




Figure 3.3: Estimates of β_h^S in equation (3.10), with 90%, 95% and 99% confidence intervals, for different horizons in quarters.

First note that the R^2 values increase for all horizons, and for some quite considerably, compared to table 3.4, where the control variables are not included. This suggests that the controls included are helpful in explaining the impact on market expectations. In other words, the model where the controls are included is probably closer to the true model explaining the impact on the FRA quotes than the one without controls. This indicates that I might have overcome some omitted-variables bias.

Table 3.5: Regression results of equation (3.10)

$h =$	1	2	3	4	5	6	7	8	9	10	11
	Second stage										
β_h^S	0.30*** (0.08)	0.39*** (0.12)	0.36*** (0.10)	0.34*** (0.07)	0.22*** (0.06)	0.15** (0.06)	0.03 (0.12)	-0.06 (0.09)	-0.12 (0.09)	-0.04 (0.08)	-0.01 (0.05)
β_h^A	0.08*** (0.03)	-0.00 (0.05)	0.02 (0.05)	0.01 (0.04)	0.00 (0.03)	0.00 (0.04)	0.03 (0.03)	0.06* (0.04)	0.05 (0.04)	0.04 (0.06)	0.06 (0.06)
$\gamma_{0,h}$	-0.00 (0.01)	0.02* (0.01)	0.02* (0.01)	0.02** (0.01)	0.02* (0.01)	0.01 (0.01)	0.02 (0.01)	0.01 (0.01)	0.01 (0.01)	0.02 (0.01)	0.01 (0.01)
$\gamma_{1,h}$	0.64*** (0.12)	0.50*** (0.18)	0.23 (0.19)	0.22* (0.13)	0.19* (0.10)	0.17 (0.11)	0.03 (0.14)	-0.01 (0.15)	-0.04 (0.12)	-0.01 (0.14)	0.02 (0.11)
$\gamma_{2,h}$	0.01 (0.05)	-0.03 (0.09)	-0.07 (0.09)	-0.20** (0.09)	-0.13 (0.10)	-0.07 (0.12)	0.12 (0.19)	0.24* (0.14)	0.27** (0.11)	0.07 (0.11)	0.11* (0.05)
$\gamma_{3,h}$	-0.01 (0.02)	-0.06*** (0.02)	-0.07*** (0.01)	-0.09*** (0.02)	-0.06*** (0.01)	-0.04*** (0.02)	-0.04** (0.02)	-0.02 (0.02)	-0.02 (0.02)	0.00 (0.03)	-0.02 (0.02)
$\gamma_{4,h}$	-0.01 (0.03)	0.02 (0.05)	0.06 (0.05)	0.05 (0.04)	0.04 (0.04)	0.02 (0.04)	0.01 (0.05)	-0.02 (0.05)	-0.02 (0.04)	-0.03 (0.05)	-0.04 (0.05)
$\gamma_{5,h}$	0.02 (0.09)	0.03 (0.14)	0.17 (0.13)	0.25*** (0.09)	0.18*** (0.06)	0.19*** (0.07)	0.22*** (0.08)	0.10 (0.08)	0.18** (0.09)	0.28 (0.18)	-0.07 (0.12)
$\gamma_{6,h}$	-0.70*** (0.18)	-1.13*** (0.44)	-0.74*** (0.25)	-0.46*** (0.13)	-0.33*** (0.10)	-0.28*** (0.11)	-0.18 (0.14)	-0.10 (0.14)	0.02 (0.14)	0.10 (0.20)	-0.03 (0.20)
R^2	0.92	0.68	0.60	0.68	0.68	0.63	0.50	0.41	0.38	0.20	0.17
p_F	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.01	0.33	0.47
Obs.	36	46	46	48	48	48	46	46	46	46	48

Equation: $Imp_{h,t} = \beta_h^S Surp_{h,t} + \beta_h^A Antic_{h,t} + \gamma_{0,h} + \gamma_{1,h} Surp_{0,t} + \gamma_{2,h} \frac{1}{10} \sum_{j \neq h} Surp_{j,t} + \gamma_{3,h} D_t^{ELB} + \gamma_{4,h} Disagr_t + \widetilde{D_{h,t}^{Closing}} Surp_{h,t} + \gamma_{6,h} \widetilde{Frac_t} Surp_{h,t} + \varepsilon_{h,t}$.

Data sources: Bloomberg, Nasdaq OMX and the Riksbank.

Note: h refers to the horizon in quarters. For the first-stage results, see table 3.4. ***, ** and * refer to significance at the 1%, 5% and 10% levels, respectively.

Now turn to the estimates of β_h^S . We see that these are smaller for short horizons, as well as long horizons, with than without controls included. However, they increase for some medium-term horizons, e.g., 4 quarters. The most likely explanation for why the estimates decrease for short horizons is that these are affected by the control for the surprise in the repo-rate decision. The estimates of these coefficients, $\gamma_{1,h}$, are larger than $\hat{\beta}_h^S$, and significantly different from zero for the shortest horizons, $h \in \{1, 2\}$. Judging from the size of the estimates, it seems that a surprise in the repo-rate decision is more effective than a surprise in the repo-rate path at managing the market expectations of very short horizons.

There are some more notable results in table 3.5. The standard errors of $\hat{\beta}_h^S$ increase compared to the case where the control variables are not included. This is probably the effect of more parameters being estimated from a small sample. As in the case without control variables, the estimates of β_h^A are well approximated by zero for most horizons, and the main results remain if *Anticipated* $_{h,t}$ is omitted from the equation.

The estimates of $\gamma_{2,h}$, capturing the impact of path surprises in all horizons *except* h , are significantly different from zero for some horizons. This suggests that there may be reactions to the entire curve rather than the specific quarterly timing of the repo-rate path. Some of the $\hat{\beta}_h^S$ might be overestimated in the sense that there is a counter impact in the other direction, while others might be underestimated. However, for most horizons, the estimates of $\gamma_{2,h}$ are well approximated by zero, so the main picture remains intact.

As mentioned above, the dummy variable for the effective lower bound, and the degree of disagreement, captured by coefficients $\gamma_{3,h}$ and $\gamma_{4,h}$, respectively, are included in an attempt to prevent an omitted-variables bias. Although the estimates of these might be interesting for other reasons, they are not important for our question of interest, and are hence not further discussed here.

The coefficient $\gamma_{5,h}$ captures the extra effect of a surprise if it is in the direction that closes the existing disagreement, or discrepancy, between the Riksbank's forecast and the market expectations. The estimates are significantly larger than zero for some horizons (and never significantly smaller than zero), suggesting that at least for some horizons, the impact of a surprise might be larger if the surprise is such that the Riksbank's new

forecast is more in line with the market expectations. Intuitively, it makes sense that a movement closer to the market expectations is viewed as more credible by the market, which is in line with the positive coefficients. Although not significant for all horizons, $\hat{\gamma}_{5,h}$ does, in general, have the correct sign. The lack of significance might arise from measurement errors and a small sample, as discussed above.

The estimates of $\gamma_{6,h}$, capturing if the timing of the announcement within the calendar quarter matters, are significantly smaller than zero for some, mainly short, horizons. This is intuitive, since the interpretation is that an announcement early within a quarter is viewed as less credible regarding the coming quarters. It is simply the case that there is more time left until the beginning of the calendar quarter that the announcement refers to, for all horizons. It also makes sense that the effect is larger for shorter horizons, since the relative difference caused by the announcement date within the quarter is larger the shorter is the horizon.

The main question of interest is what is the impact of a surprise in an announced repo-rate path on the market expectations of the future repo rate, i.e., the partial derivative $\frac{\partial \text{Impact}_{h,t}}{\partial \text{Surprise}_{h,t}}$. We recall equation (3.11), and by using the results from table 3.5, we can conclude the following; $\hat{\beta}_h^S$ is probably a good approximation of the effect in which we are interested, if we complement it with information on whether the surprise closes or opens the disagreement and where in the quarter the announcement is placed, slightly dependent on which horizons we are interested in.

Regardless of whether one finds the specification with a large set of control variables, presented in table 3.5, or the more scarce specification presented in table 3.4 more reliable, the overall impression of matching equation (3.2) to data is that it seems like the Riksbank has the ability to affect the market expectations with the repo-rate path. However, the effect is small or zero beyond one-and-a-half years, and even for shorter horizons the effect is not one-to-one. Less than half of a surprise is likely reflected by movements in the expectations, and the effect is decreasing with the horizon. These results may be viewed as lower bounds, since there is reason to suspect biased estimates of β_h^S towards zero due to regression dilution because of measurement errors.

Managing the expectations up to about half a year might be more effectively done by the repo-rate decision being a surprise, although the

path also has some bite here. One should also bear in mind that the surprise in the repo-rate decision is highly positively correlated with the surprise in the very short horizon of the repo-rate path, so, in practice, a combination of path and decision surprise is often the case.

Even with the more extensive set of controls, there is still reason to be concerned about omitted-variables bias. Especially, I would like to include controls for the surprise in other information released by the Riksbank simultaneously as the announcement of the repo-rate path and decision. As mentioned above, this includes forecasts for other macroeconomic variables, and an analysis of the current economic situation. However, this is unobserved and very difficult to proxy and hence, I have no other choice than to leave it omitted. This might bias the estimates of interest, and it is difficult to guess the sign and size of such a potential bias.

3.3.2 Robustness

In this section, I discuss the robustness of the results presented in section 3.3.1. This is particularly important since there is reason to believe that measurement errors are present, and important variables might be omitted. I begin the robustness exercises by introducing a control for the within-day movement caused by other macroeconomic news than the announcement. I also look at other measures of the surprise of an announcement than those defined in equations (3.7) and (3.8). I show that the results hold when a proxy for a time-and-horizon-specific premium is introduced. I also compare the measures of anticipated announcements presented in section 3.2.2 to a survey performed before each announcement. The control variables are relaxed one at a time to investigate the importance of each, and finally I try to analyze how robust the results are over time, which is difficult with such a small sample.

In order to overcome the problems arising from the potential impact of other macroeconomic news arriving within announcement days, I impose a proxy for the impact of news other than the announcement by the Riksbank. The proxy I use is the daily movement of the Norwegian FRA rates. Economic and financial conditions are very similar in the neighboring countries Norway and Sweden, and hence there is reason to believe that the Norwegian FRA market should react to news in a similar way as the Swedish FRA market, at least to news that is not Swedish-

specific or Norway-specific in its nature. Both Norway and Sweden are small open economies, and hence largely influenced by international news. The short-term rates, both in the interbank markets and the treasury bill markets, are highly correlated. During the period of interest, there have been no coinciding days of policy-rate announcements in the two countries. Hence, including the $Impact_{h,t}$, as defined in equation (3.5), for Norway as a control variable in $X_{h,t}$ in equations (3.1) and (3.2) might capture the non-announcement effect, if there is one.²⁴ This is possible since the Norwegian and Swedish FRAs are constructed in the exact same way, with the same settlement dates. Table 3.6 in appendix 3.A.4 shows the regression results of including the term $\gamma_{7,h} Impact_{h,t}^{NO}$ in equation (3.10). In Norway, FRAs are only available for a horizon of 8 quarters, hence the quarters 9–11 have been excluded. As before, the first-stage regression remains the same as in table 3.4, and is hence not shown. The effect on the results is very limited, indicating either that the daily FRA rates are good enough at isolating the effect of an announcement or that the impact on Norwegian FRAs is not good enough at capturing the effect of other news. It is also worth noting that the coefficients for the impact on Norwegian FRAs are not significant for most horizons.

I now turn to the measure of the variable $Anticipated_{h,t}$, and hence indirectly the variable $Surprise_{h,t}$. So far, these have been defined in two ways, either by equations (3.6) and (3.7) or by the regression equation (3.8) (together with (3.9)). I will investigate two more cases, suggested in Moessner and Nelson (2008): the path that the market expects the Riksbank to announce is given by the market pricing of the FRAs just prior to the announcement, and the path that the market expects the Riksbank to announce is the same as the one that was previously announced by the

²⁴ An endogeneity problem might also arise, if the announcement by the Riksbank also has an impact on the market expectations of future Norwegian policy rates. We would have what Angrist and Pischke (2008) refer to as a “bad control”. This is not unrealistic, since monetary policy is typically highly correlated in Norway and Sweden. However, including a dummy variable for the announcement days of Norges Bank, the central bank of Norway, in the regression equation (3.3) gives estimates that are not significantly larger than zero for any horizons except the 12-quarter horizon. This suggests that at least the Swedish market is not affected to any considerable extent by the communication of Norges Bank, so one might expect the reverse to be true as well.

Riksbank. More formally,

$$Anticipated_{h,t} = FRA_{h,t-\epsilon} - Path_{h,t_p}^{RB} \quad \text{and} \quad (3.12)$$

$$Anticipated_{h,t} = 0 \quad (3.13)$$

respectively. In the first case, the market expectation of communication and action by the Riksbank coincides. The market disregards the history of, an often systematic, discrepancy between the Riksbank forecast and the market expectations, and expects the Riksbank to fully change its forecast to be in line with the market's expectations. This is what Svensson (2015) refers to as *full predictability*. The second case assumes that the anticipated change in the Riksbank's communication is zero. Kuttner (2001) and Gürkaynak et al. (2005) argue that this is not likely the case. Both these assumptions might seem extreme and unrealistic, but have the advantage of being simple to relate to, and are hence worthwhile investigating. Note that when equation (3.12) or (3.13) is used, there is no need for a first-stage regression.

A summary of the results of the first case, expected communication coincides with the market expectations, i.e., the anticipated change in the communicated path is given by equation (3.12), is given in table 3.7 in appendix 3.A.4. Note that the R^2 is similar compared to the main results in table 3.5, and also the estimates of β_h^S are quite similar. However, also the estimates of β_h^A are in general significantly larger than zero, suggesting that this specification is probably not as good as the one used in the main analysis.

The results of the other alternative case, when the communicated path is expected not to change since the previous announcement, i.e., the anticipated change in the communicated path is zero, are summarized in table 3.8. Note that this is a regression of equation (3.1) rather than equation (3.2), since the anticipated change is defined to be zero in this case. The estimates of β_h are very different from the main results. This entire case is difficult to interpret and does not add much to the conclusion. However, it serves to emphasize the importance of finding a good measure of the anticipated and unanticipated parts of the Riksbank communication.

One obvious drawback of using FRAs as a measure of expectations of the future repo rate is that it might contain different kinds of premia, which was briefly discussed in section 3.2.2. In the main analysis, I compensate

for a *time-and-horizon-fixed* aggregate premium. An assumption that the premium does not change over time and is the same for all horizons might be too strong. Without commenting further on the type or nature of these possible premia, I follow Ferrero and Secchi (2009) in an attempt to allow the aggregate premium to vary over time and horizons. The idea is that, although market rates like the FRAs might contain premia, survey expectations should not. Hence, I construct the varying horizon-specific premium as the difference between FRA rates and the expected future repo rate according to a survey.²⁵ The survey is not conducted on the same dates as the announcements by the Riksbank, so I use linear interpolation in the time dimension to get a timely estimate. Moreover, the survey only concerns horizons of 1, 4 and 8 quarters. Linear interpolation is used also in the horizon dimension to get estimates for intermediate horizons. I extrapolate beyond 8 quarters by simply using the value of the 8-quarter horizon. Figure 3.6 in appendix 3.A.4 shows how these measures of the premium have evolved during the period of interest.

Table 3.9 presents the results when the varying premium is used. In line with Ferrero and Secchi (2009), the estimates of β_h^S are slightly higher for most short horizons. The R^2 is also increased for some short horizons. The overall impression remains when the time-varying premium is used. A notable feature here is that the estimates of $\gamma_{1,h}$, the impact of a surprise in the repo-rate decision, are significantly larger than zero only for the 1-quarter horizon, suggesting that the repo-rate decision might be an even worse instrument for affecting expectations beyond the immediate future.

It is important to be aware that using a time-varying premium, defined the way I have, does not resolve the problem that the announcement might have an effect on the premium itself; see equations (3.4) and (3.5). This weakness remains, but as has been mentioned above, it might not be important to distinguish between affecting the expectations component, or the premium component of the market rates. What matters for the effectiveness of monetary policy is the interest rates that the agents in the economy face.

²⁵A monthly survey performed by TNS Sifo Prospera (prospera.se) is used, where about 50 money market participants, mainly Swedish but also international, are asked to quantify their beliefs on different macroeconomic developments, including the repo rate, for different horizons. Up until the third quarter of 2009, the survey was conducted only quarterly.

Let us now return to the question of whether my measure of the anticipated (by the market) repo-rate path is appropriate. As mentioned above, the market expectation of which repo-rate path will be announced by the Riksbank, just prior to the announcement, is unobserved. However, I have defined and used two different proxies, one defined by equations (3.6) and (3.7) and the other defined by regression of equation (3.8) and then using equation (3.9). See also two robustness exercises discussed above in this section. One way of verifying the quality of these proxies is to compare them with surveys of the Riksbank's expected communication, although this is concerned with other potential problems. One such survey exists, where the Swedish commercial bank SEB asks the largest Swedish bond-market investors approximately one week prior to a new Riksbank announcement about their quantitative beliefs regarding not only the *actions* by the Riksbank but also regarding the *communication*, specifically the announcement of a repo-rate path. The survey does not cover the entire path, but typically asks about three specific horizons, and is hence not suitable as a substitute for the proxies used in this study. However, it is interesting to compare these specific observations to the proxies I use.

Figure 3.7 plots expectations of the announced path according to the survey, and according to my preferred proxy in equations (3.8) and (3.9), where colors represent different horizons. Running simple regressions where the proxy is explained by the survey and a constant suggests that a null hypothesis that the coefficient is one for the survey measure and the constant is zero cannot be rejected.²⁶ The results hold also when the proxy is defined by equations (3.6) and (3.7). Figure 3.7 illustrates the same thing for the surprise component at announcement, instead of the anticipated path. From this I conclude that the survey gives reason to believe that the proxies for path surprises used in this study are valid.

A very interesting question, which has been left untouched so far, is whether the results are robust over time. Detmers and Nautz (2012) use similar methods and find that the ability of the Reserve Bank of New Zealand to affect market expectations with the policy-rate path has declined since the onset of the financial crisis in 2008. Unfortunately, this sample is too small to make any advanced exercises on the subject, like

²⁶This corresponds to all observations being on the 45-degree line in figure 3.7.

distinguishing between the periods before and after the financial crisis. However, the results are robust to omitting periods of one year at the time.²⁷

To summarize this section on robustness, the main results in this paper hinge on proxies for unobservable variables. Hence, there is reason to believe that the main explanatory variables suffer from measurement errors which could affect the conclusions in a meaningful way. Hence, it is important to investigate how sensitive the results are to different approaches and methods. This has been carried out in several ways, and the conclusion is that the main impression of the results survives.

3.4 Conclusions

Market expectations of future monetary policy, as measured by futures rates, tend to move more than normal on days when the Riksbank announces a new repo-rate decision, including a repo-rate path and a report analyzing the economic situation. This raises the question of what the ability of the Riksbank is to manage such expectations with the repo-rate path. In this paper, I have performed an event study with regression analysis partly relying on high-frequency identification to investigate whether unanticipated changes in the repo-rate path impact the market expectations of the future repo rate for different horizons. There are no perfect measures either of the anticipation of an announced repo-rate path or of the expected future repo rate, so I have constructed proxies for these. The proxy for expected future action, i.e., the future repo rate, is standard in the literature, and my results suggest that the proxy I use for the anticipated communicated path is valid.

The main finding is that the Riksbank has the ability to use the repo-rate path to affect market expectations up to between one and three years. The effect is less than one-to-one, and decreasing with the horizon. For very short horizons, one and two quarters, a surprise in the repo-rate decision might be a more effective tool for managing expectations. A combination of surprises in the decision and path, which is typically the case, seems to be an effective way of managing expectations.

The impact of a path surprise might be even larger if the surprise is in

²⁷These results are available upon request.

such a direction that it closes the gap between the Riksbank's announced forecast and the market expectations. If the surprise is such that the gap is further widened, the impact might be smaller. Also the timing within the calendar quarter of the announcement might affect the size of the impact, especially for short horizons.

There is a risk that the estimated effects are biased towards zero due to measurement errors in the proxy for surprise in the repo-rate path. Another potential source of biased estimates is that I have no means of controlling for the impact of surprises in other information released by the Riksbank in the same announcement as the repo-rate path and decision. Such information includes forecasts of other macroeconomic variables and an analysis of the current economic situation.

The main results are robust variations in the measurement of announcement surprises. This includes the time window size used for high-frequency identification, controlling for other news arriving within the window, and allowing for time-varying premia.

References

- Angrist, J. D. and Pischke, J.-S. (2008). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton university press.
- Archer, D. (2005). Central bank communication and the publication of interest rate projections. In *Sveriges Riksbank conference, "Inflation Targeting: Implementation, Communication, and Effectiveness"*, Stockholm, June.
- Björk, T. (2004). *Arbitrage Theory in Continuous Time*. Oxford university press, 2 edition.
- De Graeve, F. and Iversen, J. (2015). Central bank policy paths and market forward rates: A simple model. Working Paper 303, Sveriges Riksbank.
- Detmers, G.-A. and Nautz, D. (2012). The Information Content of Central Bank Interest Rate Projections: Evidence from New Zealand. *Economic Record*, 88(282):323–329.
- Dincer, N. N. and Eichengreen, B. (2007). Central Bank Transparency: Where, Why, and with What Effects? Working Paper 13003, National Bureau of Economic Research.
- Ehrmann, M. and Fratzscher, M. (2004). Taking Stock: Monetary Policy Transmission to Equity Markets. *Journal of Money, Credit and Banking*, 36(4):719–737.
- Ferrero, G. and Secchi, A. (2009). The announcement of monetary policy intentions. *Bank of Italy Temi di Discussione (Working Paper) No.* 720.
- Gürkaynak, R. S., Sack, B., and Swanson, E. T. (2005). Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements. *International Journal of Central Banking*, 1(1):55–93.
- Gürkaynak, R. S., Sack, B. P., and Swanson, E. T. (2007). Market-Based Measures of Monetary Policy Expectations. *Journal of Business & Economic Statistics*, 25(2):201–212.

- Gürkaynak, R. S. and Wright, J. H. (2013). Identification and inference using event studies. *The Manchester School*, 81:48–65.
- Ingves, S. (2007). Communication – what demands are made of an independent central bank? Speech at the Swedish Economics Association, Stockholm.
- Iversen, J. and Tysklind, O. (2017). The effect of repo rate path changes on asset prices. *Sveriges Riksbank Staff memo*, June.
- Krueger, J. T. and Kuttner, K. N. (1996). The Fed Funds Futures Rate as a Predictor of Federal Reserve Policy. *Journal of Futures Markets*, 16(8):865–879.
- Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics*, 47(3):523–544.
- Moessner, R. and Nelson, W. R. (2008). Central Bank Policy Rate Guidance and Financial Market Functioning. *International Journal of Central Banking*, 4(4):193–226.
- Nelson, C. R. and Siegel, A. F. (1987). Parsimonious modeling of yield curves. *Journal of business*, pages 473–489.
- Svensson, L. E. (1994). Estimating and Interpreting Forward Interest Rates: Sweden 1992-1994. Working Paper 4871, National Bureau of Economic Research.
- Svensson, L. E. (2015). Forward Guidance. *International Journal of Central Banking (forthcoming)*, 11.
- Sveriges Riksbank (2013). Perspectives on monetary policy expectations and forward rates. *Monetary Policy Report February 2013 (Box article)*, pages 48–52.
- Sveriges Riksbank (2015). The riksbank’s complementary monetary policy measures. *Monetary Policy Report February 2015 (Box article)*, pages 39–44.

- Winkelmann, L. (2010). The Norges Bank's key rate projections and the news element of monetary policy: A wavelet based jump detection approach. SFB 649 discussion paper 2010-062, Berlin.

3.A Appendix

3.A.1 Repo-rate paths and market pricing




Figure 3.4: Outcome of the **repo rate**, and forecasts by the Riksbank (black dotted) and the market (grey dotted) as quarterly averages at announcement dates.

3.A.2 Illustrating the control variable *Disagreement*




Figure 3.5: Illustration of how the control variable *Disagreement*, described in section 3.2.2, evolves over time.

3.A.3 Surprise in the repo-rate decision

One of the control variables that are included in the main analysis is the surprise component of a repo-rate decision. This is the difference between the change that is anticipated by the market just prior to a Riksbank announcement and the actual change of the repo rate at that announcement. Note that the surprise can be non-zero even if the change of the repo rate is zero, if the anticipated change is non-zero. The anticipated change is unobserved and hence, so is the surprise. However, market prices provide a good proxy for the anticipated change, and the method used to create this proxy is described in this section.

The method described here is the same as the one the Riksbank has been using for many years. However, the details have not been published, and hence I provide them here. The method is loosely built on Krueger and Kuttner (1996).

I use interest rate swaps to determine the surprise component of a repo-rate decision. The swaps I use are 30-day STINA swaps, with the STIBOR T/N as the underlying variable part.²⁸ The way the STINA swap works is that it breaks even if the (geometrical) average STIBOR T/N from $t+2$ to $t+32$ equals the swap rate. This makes the STINA swap rate the expected average of the STIBOR T/N rate over the duration. I disregard any swap premia, which I have reasons to believe are very close to zero in this case.

In turn, the STIBOR T/N rate is typically very close to the repo rate plus a fixed premium of 10 basis points.²⁹ In cases where the premium deviates from the normal 10 basis points, I make an implicit assumption that it will not be affected by the repo-rate announcement.

When a repo-rate decision is announced, it is not implemented until the next Wednesday after the day before the announcement. When using the STINA swaps to measure the expected repo-rate change, I consider two periods; the period from the announcement to the implementation of the (potentially) new repo-rate (I) and the period from the implementation to the settlement day of the STINA-swap (II). I exploit that the repo rate is known during period I , and once the decision is announced, it is

²⁸STIBOR T/N is a Swedish inter-bank interest rate used for overnight loans stretching from the next day to the day after that.

²⁹This is a consequence of the facilities offered to banks by the Riksbank.

also known in period II .³⁰ I consider the change in the STINA swap rate from just prior to an announcement at time t ($i_{t-\epsilon}^{\text{STINA}}$) to just after the announcement (i_t^{STINA}), and let τ_I and τ_{II} denote the number of days of period I and II , respectively.³¹ We have

$$1 + i_t^{\text{STINA}} \cdot \frac{30}{360} = \left(1 + E_t \left[i_I^{\text{T/N}} \right] \cdot \frac{1}{360} \right)^{\tau_I} \left(1 + E_t \left[i_{II}^{\text{T/N}} \right] \cdot \frac{1}{360} \right)^{\tau_{II}},$$

where $i_I^{\text{T/N}}$ and $i_{II}^{\text{T/N}}$ are the STIBOR T/N rates in periods I and II respectively, by the assumption that the STINA swap interest rate is the expected average of the STIBOR T/N over the duration. Taking natural logarithms and using the known approximation that $\ln(1+x) \approx x$ when x is small, we get

$$\begin{aligned} i_t^{\text{STINA}} \cdot \frac{30}{360} &\approx E_t \left[i_I^{\text{T/N}} \right] \cdot \frac{\tau_I}{360} + E_t \left[i_{II}^{\text{T/N}} \right] \cdot \frac{\tau_{II}}{360} && \Longleftrightarrow \\ i_t^{\text{STINA}} &\approx \frac{\tau_I E_t \left[i_I^{\text{T/N}} \right] + \tau_{II} E_t \left[i_{II}^{\text{T/N}} \right]}{30}. \end{aligned}$$

Equivalently, we have

$$i_{t-\epsilon}^{\text{STINA}} \approx \frac{\tau_I E_{t-\epsilon} \left[i_I^{\text{T/N}} \right] + \tau_{II} E_{t-\epsilon} \left[i_{II}^{\text{T/N}} \right]}{30}.$$

This gives us, for the change in the STINA swap rate,

$$\begin{aligned} i_t^{\text{STINA}} - i_{t-\epsilon}^{\text{STINA}} &\approx \frac{\tau_I E_t \left[i_I^{\text{T/N}} \right] + \tau_{II} E_t \left[i_{II}^{\text{T/N}} \right] - \left(\tau_I E_{t-\epsilon} \left[i_I^{\text{T/N}} \right] + \tau_{II} E_{t-\epsilon} \left[i_{II}^{\text{T/N}} \right] \right)}{30} \\ &= \frac{\tau_I \left(E_t \left[i_I^{\text{T/N}} \right] - E_{t-\epsilon} \left[i_I^{\text{T/N}} \right] \right) + \tau_{II} \left(E_t \left[i_{II}^{\text{T/N}} \right] - E_{t-\epsilon} \left[i_{II}^{\text{T/N}} \right] \right)}{30}. \end{aligned}$$

Since the repo rate is known in period I we have $E_t \left[i_I^{\text{T/N}} \right] = E_{t-\epsilon} \left[i_I^{\text{T/N}} \right]$,

³⁰I am using 30-day swaps, and regular repo-rate announcements of the Riksbank are more than one month apart. However, two times during the period of interest, the repo rate has been changed between regular announcement days. If the market expects this to be the case, the results may be invalid. Since repo-rate changes between meetings seem to be very rare and difficult to predict, I assume that the market expectations rely on this not being the case.

³¹Since the announcements by the Riksbank take place on different weekdays, τ_I and τ_{II} vary, but we always have $\tau_I + \tau_{II} = 30$.

and since it is not expected to change within period II we have

$$\begin{aligned} E_t \left[i_{II}^{T/N} \right] - E_{t-\epsilon} \left[i_{II}^{T/N} \right] &= E_t \left[i_{II}^{\text{repo}} + \eta \right] - E_{t-\epsilon} \left[i_{II}^{\text{repo}} + \eta \right] \\ &= E_t \left[i_{II}^{\text{repo}} \right] - E_{t-\epsilon} \left[i_{II}^{\text{repo}} \right] \\ &= \text{Surprise}_{0,t}, \end{aligned}$$

i.e., the unanticipated change in the repo rate at time t , where η is the premium of the STINA T/N rate. This is where the assumption that the premium is unaffected by the announcement is crucial. Solving finally gives the surprise as a measure of the observed change in the STINA swap rate,

$$\begin{aligned} i_t^{\text{STINA}} - i_{t-\epsilon}^{\text{STINA}} &= \frac{0 + \tau_{II} \text{Surprise}_{0,t}}{30} && \Longleftrightarrow \\ \text{Surprise}_{0,t} &= \frac{\tau_{II}}{30} \left(i_t^{\text{STINA}} - i_{t-\epsilon}^{\text{STINA}} \right). \end{aligned}$$

In practice, the Riksbank uses the change in the STINA swap rate from the day prior to a repo-rate announcement to the end of the announcement day, for technical reasons. The formula then has to be adjusted with another term, since the two STINA swaps compared have different settlement days. The derivation is straightforward and this is also the measure that is used in this study.

3.A.4 Results of robustness checks; tables and figures

Table 3.6: Regression results of equation (3.10) including the movement of corresponding Norwegian futures rate as a control variable

$h =$	1	2	3	4	5	6	7	8
	Second stage							
β_h^S	0.17** (0.07)	0.43*** (0.13)	0.36*** (0.10)	0.34*** (0.08)	0.21*** (0.06)	0.12 (0.07)	0.03 (0.12)	-0.02 (0.09)
β_h^A	0.06** (0.03)	-0.00 (0.05)	0.02 (0.05)	0.01 (0.04)	0.00 (0.03)	-0.00 (0.04)	0.03 (0.03)	0.05 (0.03)
$\gamma_{7,h}$	0.49** (0.20)	-0.21 (0.32)	-0.02 (0.38)	0.13 (0.23)	0.29** (0.13)	0.17 (0.13)	0.17 (0.14)	0.37*** (0.11)
R^2	0.93	0.69	0.60	0.69	0.70	0.64	0.52	0.50
p_F	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Obs.	36	45	46	48	48	48	46	46

Equation: $Imp_{h,t} = \beta_h^S Surp_{h,t} + \beta_h^A Antic_{h,t} + \gamma_{0,h} + \gamma_{1,h} Surp_{0,t} + \gamma_{2,h} \frac{1}{10} \sum_{j \neq h} Surp_{j,t} + \gamma_{3,h} D_t^{ELB} + \gamma_{4,h} \widetilde{Disagr}_t + \gamma_{5,h} \widetilde{D}_{h,t}^{Closing} Surp_{h,t} + \gamma_{6,h} \widetilde{Frac}_t Surp_{h,t} + \gamma_{7,h} Imp_{h,t}^{NO} + \varepsilon_{h,t}$.

Data sources: Bloomberg, Nasdaq OMX and the Riksbank.

Note: h refers to the horizon in quarters. The regressions also include the coefficients $\gamma_{n,h}$, $n \in \{0, 1, \dots, 6\}$. The estimates of these have been omitted. ***, ** and * refer to significance at the 1%, 5% and 10% levels, respectively.

Table 3.7: Regression results of equation (3.10) when the anticipated change of the path is given by (3.12)

$h =$	1	2	3	4	5	6	7	8	9	10	11
β_h^S	0.15** (0.06)	0.24*** (0.08)	0.44*** (0.09)	0.29*** (0.05)	0.20*** (0.05)	0.10** (0.05)	0.05 (0.06)	-0.06 (0.07)	-0.03 (0.08)	0.07 (0.07)	0.07 (0.07)
β_h^A	0.05* (0.03)	0.01 (0.06)	0.10** (0.04)	0.10*** (0.03)	0.08** (0.03)	0.10*** (0.03)	0.12*** (0.03)	0.11*** (0.02)	0.08*** (0.03)	0.06 (0.04)	0.06 (0.05)
R^2	0.93	0.71	0.69	0.68	0.63	0.57	0.49	0.46	0.36	0.20	0.20
p_F	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.34	0.48
Obs.	42	47	47	48	48	48	47	47	47	47	40

Equation: $Imp_{h,t} = \beta_h^S Surph_{h,t} + \beta_h^A Antich_{h,t} + \gamma_{0,h} + \gamma_{1,h} Surp_{0,t} + \gamma_{2,h} \frac{1}{10} \sum_{j \neq h} Surp_{j,t} + \gamma_{3,h} D_{\ell}^{ELB} + \gamma_{4,h} Disagr_t + \gamma_{5,h} \widetilde{D}_{h,t}^{Closing} Surph_{h,t} + \gamma_{6,h} \widetilde{Frac}_t Surph_{h,t} + \varepsilon_{h,t}$.

Data sources: Bloomberg, Nasdaq OMX and the Riksbank.

Note: h refers to the horizon in quarters. The regressions also include the coefficients $\gamma_{n,h}$, $n \in \{0, 1, \dots, 6\}$. The estimates of these have been omitted. ***, ** and * refer to significance at the 1%, 5% and 10% levels, respectively.

Table 3.8: Regression results of equation (3.10) when the anticipated change of the path is zero

$h =$	1	2	3	4	5	6	7	8	9	10	11
β_h	0.11*** (0.03)	-0.03 (0.08)	-0.03 (0.08)	0.17*** (0.07)	0.26*** (0.06)	0.26*** (0.08)	0.18*** (0.08)	0.13* (0.07)	-0.00 (0.07)	0.00 (0.07)	0.01 (0.06)
R^2	0.94	0.52	0.43	0.49	0.58	0.56	0.43	0.34	0.26	0.13	0.13
p_F	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.09	0.59	0.56
Obs.	42	47	47	48	48	48	47	47	47	47	48

Equation: $Imp_{h,t} = \beta_h \widetilde{Surp}_{h,t} + \gamma_{0,h} + \gamma_{1,h} Surp_{0,t} + \gamma_{2,h} \frac{1}{10} \sum_{j \neq h} Surp_{j,t} + \gamma_{3,h} D_t^{ELB} + \gamma_{4,h} \widetilde{Disagr}_t + \gamma_{5,h} \widetilde{D}_{h,t}^{Closing} Surp_{h,t} + \gamma_{6,h} \widetilde{Frac}_t Surp_{h,t} + \varepsilon_{h,t}.$

Data sources: Bloomberg, Nasdaq OMX and the Riksbank.

Note: h refers to the horizon in quarters. The regressions also include the coefficients $\gamma_{n,h}$, $n \in \{0, 1, \dots, 6\}$. The estimates of these have been omitted. ***, ** and * refer to significance at the 1%, 5% and 10% levels, respectively.




Figure 3.6: Time-varying premia constructed as the difference between market pricing and a survey.

Table 3.9: Regression results of equation (3.10) with a time-and-horizon-varying premium illustrated in figure 3.6

$h =$	1	2	3	4	5	6	7	8	9	10	11
Second stage											
β_h^S	0.32*** (0.06)	0.36*** (0.08)	0.41*** (0.08)	0.41*** (0.08)	0.25*** (0.05)	0.18*** (0.05)	-0.05 (0.12)	-0.10 (0.11)	-0.15 (0.11)	-0.03 (0.07)	-0.03 (0.05)
β_h^A	0.07* (0.03)	0.04 (0.03)	0.04 (0.04)	0.06 (0.04)	0.05 (0.04)	0.02 (0.04)	0.05 (0.03)	0.06** (0.03)	0.07*** (0.03)	0.08 (0.05)	0.07 (0.05)
$\gamma_{1,h}$	0.51*** (0.10)	0.26 (0.16)	0.07 (0.17)	0.18 (0.13)	0.16 (0.11)	0.14 (0.13)	-0.07 (0.16)	-0.09 (0.16)	-0.07 (0.13)	0.04 (0.18)	0.01 (0.12)
R^2	0.94	0.80	0.70	0.71	0.68	0.65	0.50	0.44	0.39	0.14	0.17
p_F	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.67	0.46
Obs.	36	46	46	48	48	48	46	46	46	46	48
First stage											
α_h	-0.03 (0.04)	-0.05* (0.03)	-0.08** (0.03)	-0.13*** (0.05)	-0.14*** (0.06)	-0.14** (0.07)	-0.08 (0.08)	-0.03 (0.09)	0.00 (0.12)	0.01 (0.15)	-0.02 (0.15)
μ_h^M	1.26* (0.14)	0.97 (0.09)	0.91 (0.08)	0.90 (0.09)	0.86 (0.10)	0.80* (0.11)	0.76** (0.10)	0.67*** (0.08)	0.60*** (0.08)	0.51*** (0.07)	0.47*** (0.08)
$\mu_h^{M_p}$	-0.64 (0.26)	-0.44*** (0.13)	-0.49*** (0.12)	-0.53*** (0.15)	-0.51*** (0.13)	-0.43*** (0.12)	-0.36*** (0.12)	-0.25*** (0.11)	-0.22*** (0.12)	-0.17*** (0.12)	-0.20*** (0.10)
μ_h^P	0.39*** (0.21)	0.53*** (0.14)	0.64*** (0.10)	0.70** (0.11)	0.72*** (0.10)	0.71*** (0.08)	0.66*** (0.09)	0.63*** (0.10)	0.65*** (0.12)	0.70** (0.12)	0.78*** (0.09)
R^2	0.99	0.99	0.99	0.98	0.97	0.96	0.96	0.96	0.96	0.96	0.96
Obs.	36	46	46	48	48	48	46	46	46	46	48

First stage: $Path_{h,t}^{RB} = \alpha_h + \mu_h^M FRA_{h,t-\epsilon} + \mu_h^{M_p} FRA_{h,t_p} + \mu_h^P Path_{h,t_p}^{RB} + Surprise_{h,t}$.

Second stage: $Impact_{h,t} = \beta_h^S Surprise_{h,t} + \beta_h^A Anticipated_{h,t} + \gamma_{0,h} + \epsilon_{h,t}$.

Data sources: Bloomberg, Nasdaq OMX, the Riksbank and TNS Sifo Prospera.

Note: h refers to the horizon in quarters. In the first-stage regression, significance levels refer to significant difference from the reference levels $(0, 1, -1, 1)$ for $(\alpha_h, \mu_h^M, \mu_h^{M_p}, \mu_h^P)$, respectively. The regressions also include the coefficients $\gamma_{n,h}$, $n \in \{0, 2, 3, \dots, 6\}$. The estimates of these have been omitted. **, * and * refer to significance at the 1%, 5% and 10% levels, respectively.




Figure 3.7: Comparison of a survey measure of expected communicated repo-rate path and the corresponding proxy used in the quantitative exercises, defined by equations (3.8) and (3.9). Colors represent horizons, from **shortest** (1 quarter) to longest **longest** (11 quarters). Note that the dashed line is a 45° line, not a regression line.




Figure 3.8: Comparison of the announcement surprises associated with a survey measure of expected communicated repo-rate path and the corresponding proxy used in the quantitative exercises, defined by equation (3.8). Colors represent horizons, from **shortest** (1 quarter) to longest **longest** (11 quarters). Note that the dashed line is a 45° line, not a regression line.

Sammanfattning

Den här avhandlingen består av tre fristående kapitel, av vilka de första två är tätt förknippade med varandra, och bygger på samma makroekonomiska modell. Alla tre kapitel har det gemensamt att de ursprungligen motiveras av mitt intresse för ekonomisk politik. Min bakgrund som anställd vid Sveriges Riksbanks avdelning för penningpolitik har lärt mig vikten av att det bedrivs forskning som är relevant för det beslutsfattande som påverkar väldigt många människors liv.

I kapitel 1, **Lönestelheter och heterogena hushåll**, utgår jag från en standardmodell med heterogena hushåll, där det inte finns någon marknad för försäkringar som kan neutralisera den idiosynkratiska osäkerheten. Detta skapar motiv för försiktighetssparande. Jag antar att produktivitet är idiosynkratisk och fluktuerar till följd av slumpmässiga chocker. Jag antar också att löner är stela på individnivå.³² Den typ av lönestelhet jag antar är vanlig i den ny-Keynesianska litteraturen, men på aggregerad nivå. De frågor jag försöker besvara i kapitlet är om lönestelheter på individnivå ger upphov till mikroekonomiska mekanismer som uppfattas som realistiska och intuitiva. Vilka antaganden om lönestelheterna behöver jag göra för att uppnå detta? Vad är det som förklarar hushållens beteende och beslut i en sådan ekonomi, och vad blir följderna för enskilda hushålls utfall på arbetsmarknaden? Dessa frågor analyserar jag i partiell jämvikt, där lönen på aggregerad nivå, och räntan i ekonomin antas vara exogent givna.

Den typ av lönestelhet som jag analyserar medför två konceptuella utmaningar. Den första är att det finns två olika löner att hantera. Idiosynkratiska produktivitetschocker innebär att olika hushåll producerar olika mycket på samma tid. Det innebär också att hur mycket ett enskilt hushåll kan producera under en given tidsrymd, säg en timme, varierar över tid. Det betyder att det finns två olika löner: *kompensation per arbetad tidsenhet*, vilket jag refererar till som *timlön*, och *kompensation per producerad enhet*, vilket jag refererar till som *effektivitetsslön*. Lönestelheten kan bara gälla för en av dessa, och vilken av dem det är får väldigt olika konsekvenser för ekonomin. Föreställ dig ett hushåll

³²I modellen är ett hushåll att likställa med en individ, och därmed är individnivå och hushållsnivå ekvivalenta, och kommer att användas om vartannat.

vars lön sattes i en tidigare period. Nu faller hens produktivitet, men hen har inte möjlighet att uppdatera sin lön. Hen producerar nu mindre än tidigare för varje arbetad timme, och om det är hens timlön som är orörlig så innebär detta att hens arbetsgivare har samma lönekostnad för den tid som hen arbetar. Däremot är arbetsgivarens lönekostnad per enhet hen producerar högre, vilket får till följd att hens arbetsinsats är lägre värderad av arbetsgivaren. Efterfrågan på hens arbete sjunker, vilket gör att hen drabbas av vad jag refererar till som *underarbete*, ett begrepp nära besläktat med arbetslöshet.³³ Om det, å andra sidan, är hens effektivitetslön som är orörlig så producerar hen förvisso mindre per arbetad timme, men arbetsgivarens lönekostnad för en arbetad timme är också lägre. Det är inte självklart i vilken riktning denna situation påverkar dels arbetsgivarens efterfrågan på hens arbete, dels hens egen arbetsvilja. Situationerna beskrivna ovan är väldigt olika på mikronivå, och vilken av dem som antas gälla spelar stor roll för modellens resultat.

Den andra konceptuella utmaningen gäller drastiskt *överarbete*, vilket är motsatsen till *underarbete*, som introducerades ovan. Föreställ dig ett hushåll vars produktivitet stiger, medan lönen inte kan uppdateras. För varje arbetad timme producerar hen nu mer än tidigare, och om det är timlönen som är orörlig så är arbetsgivarens lönekostnad för den arbetade timmen densamma som tidigare. Uppenbarligen värderar arbetsgivaren hen nu högre, och efterfrågar mer arbete av hen. Om hen är tvingad att möta den stigande efterfrågan, vilket är det vanligaste antagandet i den befintliga litteraturen om makroekonomiska modeller med lönestelheter, så kommer hen att tvingas arbeta mer än vad hen själv skulle vilja. Denna situation, där ett hushåll arbetar mer än vad hen skulle vilja på grund av att hen inte kan uppdatera sin lön, är vad jag refererar till som *överarbete*. Produktivitet tenderar att variera kraftigare på individnivå än på aggregerad nivå, vilket gör att överarbete blir ett mer betydelsefullt fenomen då idiosynkratiska produktivitetschocker modelleras. Det kan betraktas som problematiskt att vissa hushåll tvingas arbeta mångdubbelt mer än vad de själva skulle vilja. Ett sätt att undvika detta är att tillåta hushåll att arbeta färre timmar än vad som efterfrågas från arbetsgivarens

³³En viktig skillnad är att arbetslöshet ofta betraktas som ett binärt tillstånd, det vill säga antingen är man helt arbetslös, eller inte alls arbetslös. Begreppet *underarbete*, som jag använder mig av, är snarare en kontinuerlig skala där ett hushåll kan vara mer eller mindre drabbat av *underarbete*.

håll. Jag refererar till ett sådant antagande som *icke-tvingande* arbetsefterfrågan. Det skapar en asymmetri där drastiskt underarbete förekommer, som följd av att vissa löner är för höga, medan drastiskt överarbete, till följd av för låga löner, inte är möjligt.

De två avvägningarna som beskrivs ovan — stela timlöner eller stela effektivitetslöner, respektive tvingande eller icke-tvingande arbetsefterfrågan — ger upphov till totalt fyra möjliga versioner av modellen. Mina resultat visar att endast en av dessa versioner ger upphov till trovärdiga mikroekonomiska mekanismer och plausibla korrelationer på mikronivå: stela timlöner, och icke-tvingande arbetsefterfrågan. Bara den versionen resulterar i negativ korrelation mellan idiosynkratisk produktivitet och underarbete, i kombination med ett obrutet positivt samband mellan välfärd och produktivitet på individnivå. Sjunkande produktivitet när lönen inte kan uppdateras skapar ett löne-produktivitetsgap. Det är optimalt för arbetsgivaren att substituera till andra arbetare, vilket leder till lägre efterfrågan på arbete från hushållet i fråga, som då drabbas av en period av ofrivilligt underarbete. Perioden av underarbete upphör antingen om ett tillfälle att uppdatera lönen dyker upp, eller om hushållets produktivitet stiger så att löne-produktivitetsgapet försvinner. Bägge dessa situationer gör att efterfrågan på hushållets arbete stiger.

I ett omvänt scenario, där ett hushåll blir mer produktivt medan timlönen ligger fast, uppstår ett löne-produktionsgap åt andra hållet. Detta gör att hens arbete, till den givna lönen, värderas högre av arbetsgivaren, och efterfrågan på hens arbete stiger. Hen väljer att utnyttja möjligheten att avvika nedåt från efterfrågan på hen s arbete, och i själva verket påverkas inte hens arbetstid mer än marginellt av den positiva produktivitetsschocken. Om hen däremot skulle få chansen att uppdatera sin lön så skulle hen välja en högre lön som skulle öka hens arbetsinkomst.

I kapitel 2, **Lönestelheter och heterogena hushåll i allmän jämvikt**, så tar jag samma modell som i kapitel 1, men betraktar den i allmän jämvikt. Det innebär att aggregerad lön och ränta sätts så att arbetsmarknaden och marknaden för statsobligationer klarerar, det vill säga, dessa priser bestäms endogen i ekonomin. Detta kräver en mer explicit modellering av företagssektorn, där priserna är stela på samma sätt som lönerna är. Att betrakta allmän jämvikt möjliggör en analys av hur aggregerade storheter i ekonomin påverkas av de lönestelheter på

mikronivå som jag införde i kapitel 1. Hur påverkas försiktighetssparande av risk att drabbas av underarbete, och hur påverkar detta i sin tur jämviktsräntan? Hur påverkas aggregerad produktion av att utbud och efterfrågan tillåts avvika från varandra på individnivå, vilket orsakas av lönestelheten?

Jag tar en ny-Keynesiansk modell med heterogena agenter, som är standard i litteraturen, och lägger till två saker: stela löner på hushållsnivå och ett antagande att arbetsefterfrågan är icke-bindande. Utöver att besvara frågor om aggregerade ekonomiska utfall så lämpar sig denna modell väl för att adressera fördelningsfrågor i ekonomin, som exempelvis hur arbetsinsatser och förmögenheter är fördelade.

Huvudresultaten som jag finner är att modellen i allmän jämvikt ger upphov till mikroekonomiska fördelningar som kvalitativt matchar mikrodata över USA:s ekonomi. En betydande risk för underarbete förstärker försiktighetssparandet. Jämfört med motsvarande modell utan lönestelheter leder detta till en lägre jämviktsränta. Lönestelheten leder också till en mer ojämlik fördelning av arbetsinsatserna i ekonomin, vilket resulterar i lägre aggregerad arbetsinsats i produktionen, mätt i effektivitetstermer, och i förlängningen lägre produktionsnivå. Ytterligare en följd av de mer ojämlikt fördelade arbetsinsatserna som underarbete ger upphov till är att inkomstskillnaderna är större med än utan lönestelheter. Däremot så bidrar det högre försiktighetssparandet till en mer koncentrerad fördelning av förmögenheter. Jag finner att fattiga och lågproduktiva hushåll tenderar att vara överrepresenterade bland de som lider av underarbete, vilket är i linje med mikrodata på arbetslöshet. Jag visar också att mina huvudresultat är robusta mot variationer i de viktigaste parametrarna i modellen.

I kapitel 3, **Hur stor verktygslåda har en centralbank? Att styra förväntningar med prognoser för styrräntan: evidens från Sverige**, genomför jag en eventstudie av hur svenska marknadsräntor påverkas av Sveriges Riksbanks räntepubliceringar. En räntepublicering består av minst tre delar: ett styrräntebeslut, en prognos för styrräntan, och en penningpolitisk rapport som analyserar det aktuella ekonomiska läget. Jag isolerar effekten som kommer från överraskningar i prognosen för styrräntan. Identifikationen av de eftersökta effekterna bygger delvis på högfrekventa förändringar av finansiella priser runt en räntepublicering.

Jag använder förändringar i terminsräntor på olika horisonter som mått på hur förväntningarna påverkas.

Huvudresultatet är att en överraskning i Riksbankens publicerade prognos för styrräntan har en påverkan på marknadens förväntningar om den framtida styrräntan. Effekten är dock inte ett till ett, och är bara signifikant upp till ungefär ett och ett halvt års horisont, vilket är kortare än Riksbankens prognoshorisont. Resultatet är robust mot ett antal variationer i hur variablerna som ingår i analysen mäts.

IIES monograph series

1. Michael, Michael, *The Theory of Commercial Policy: Trade and Protection*, 1973
2. Söderström, Hans Tson, *Studies in the Microdynamics of Production and Productivity Change*, 1974
3. Hamilton, Carl B., *Project Analysis in the Rural Sector with Special Reference to the Evaluation of Labour Cost*, 1974
4. Nyberg, Lars and Staffan Viotti, *A Control Systems Approach to Macroeconomic Theory and Policy in the Open Economy*, 1975
5. Myhrman, Johan, *Monetary Policy in Open Economies*, 1975
6. Krauss, Melvyn, *International Trade and Economic Welfare*, 1975
7. Wihlborg, Clas, *Capital Market Integration and Monetary Policy under Different Exchange Rate Regimes*, 1976
8. Svensson, Lars E.O., *On Competitive Markets and Intertemporal Resources Allocation*, 1976
9. Yeats, Alexander J., *Trade Barriers Facing Developing Countries*, 1978
10. Calmfors, Lars, *Prices, Wages and Employment in the Open Economy*, 1978
11. Kornai, János, *Economics of Shortage, Vols I and II*, 1979
12. Flam, Harry, *Growth, Allocation and Trade in Sweden. An Empirical Application of the Heckscher-Ohlin Theory*, 1981
13. Persson, Torsten, *Studies of Alternative Exchange Rate Systems. An Intertemporal General Equilibrium Approach*, 1982
14. Erzan, Refik, *Turkey's Comparative Advantage, Production and Trade Patterns in Manufactures. An Application of the Factor Proportions Hypothesis with Some Qualifications*, 1983
15. Horn af Rantzen, Henrik, *Imperfect Competition in Models of Wage Formation and International Trade*, 1983

16. Nandakumar, Parameswar, *Macroeconomic Effects of Supply Side Policies and Disturbances in Open Economies*, 1985
17. Sellin, Peter, *Asset Pricing and Portfolio Choice with International Investment Barriers*, 1990
18. Werner, Ingrid, *International Capital Markets: Controls, Taxes and Resources Allocation*, 1990
19. Svedberg, Peter, *Poverty and Undernutrition in Sub-Saharan Africa: Theory, Evidence, Policy*, 1991
20. Nordström, Håkan, *Studies in Trade Policy and Economic Growth*, 1992
21. Hassler, John, Petter Lundvik, Torsten Persson, and Paul Söderlind, *The Swedish Business Cycle: Stylized facts over 130 years*, 1992
22. Lundvik, Petter, *Business Cycles and Growth*, 1992
23. Söderlind, Paul, *Essays in Exchange Rates, Business Cycles and Growth*, 1993
24. Hassler, John A.A., *Effects of Variations in Risk on Demand and Measures of Business Cycle Comovements*, 1994
25. Daltung, Sonja, *Risk, Efficiency, and Regulation of Banks*, 1994
26. Lindberg, Hans, *Exchange Rates: Target Zones, Interventions and Regime Collapses*, 1994
27. Stennek, Johan, *Essays on Information-Processing and Competition*, 1994
28. Jonsson, Gunnar, *Institutions and Incentives in Monetary and Fiscal Policy*, 1995
29. Dahlquist, Magnus, *Essays on the Term Structure of Interest Rates and Monetary Policy*, 1995
30. Svensson, Jakob, *Political Economy and Macroeconomics: On Foreign Aid and Development*, 1996
31. Blix, Mårten, *Rational Expectations and Regime Shifts in Macroeconometrics*, 1997
32. Lagerlöf, Nils-Petter, *Intergenerational Transfers and Altruism*, 1997
33. Klein, Paul, *Papers on the Macroeconomics of Fiscal Policy*, 1997
34. Jonsson, Magnus, *Studies in Business Cycles*, 1997
35. Persson, Lars, *Asset Ownership in Imperfectly Competitive Markets*, 1998
36. Persson, Joakim, *Essays on Economic Growth*, 1998

37. Domeij, David, *Essays on Optimal Taxation and Indeterminacy*, 1998
38. Flodén, Martin, *Essays on Dynamic Macroeconomics*, 1999
39. Tangerås, Thomas, *Essays in Economics and Politics: Regulation, Elections and International Conflict*, 2000
40. Lidbom, Per Pettersson, *Elections, Party Politics and Economic Policy*, 2000
41. Vestin, David, *Essays on Monetary Policy*, 2001
42. Olofsgård, Anders, *Essays on Interregional and International Political Economics*, 2001
43. Johansson, Åsa, *Essays on Macroeconomic Fluctuations and Nominal Wage Rigidity*, 2002
44. Groth, Charlotta, *Topics on Monetary Policy*, 2002
45. Gancia, Gino A., *Essays on Growth, Trade and Inequality*, 2003
46. Harstad, Bård, *Organizing Cooperation: Bargaining, Voting and Control*, 2003
47. Kohlscheen, Emanuel, *Essays on Debts and Constitutions*, 2004
48. Olovsson, Conny, *Essays on Dynamic Macroeconomics*, 2004
49. Stavilöt, Ulrika, *Essays on Culture and Trade*, 2005
50. Herzing, Mathias, *Essays on Uncertainty and Escape in Trade Agreements*, 2005
51. Bonfiglioli, Alessandra, *Essays on Financial Markets and Macroeconomics*, 2005
52. Pienaar, Natalie, *Economic Applications of Product Quality Regulation in WTO Trade Agreements*, 2005
53. Song, Zheng, *Essays on Dynamic Political Economy*, 2005
54. Eiseensee, Thomas, *Essays on Public Finance: Retirement Behavior and Disaster Relief*, 2005
55. Favara, Giovanni, *Credit and Finance in the Macroeconomy*, 2006
56. Björkman, Martina, *Essays on Empirical Development Economics: Education, Health and Gender*, 2006
57. Larsson, Anna, *Real Effects of Monetary Regimes*, 2007
58. Prado, Jr., Jose Mauricio, *Essays on Public Macroeconomic Policy*, 2007
59. Tonin, Mirco, *Essays on Labor Market Structures and Policies*, 2007
60. Queijo von Heideken, Virginia, *Essays on Monetary Policy and Asset*

Markets, 2007

61. Finocchiaro, Daria, *Essays on Macroeconomics*, 2007
62. Waisman, Gisela, *Essays on Discrimination and Corruption*, 2008
63. Holte, Martin Bech, *Essays on Incentives and Leadership*, 2008
64. Damsgaard, Erika Färnstrand, *Essays on Technological Choice and Spillovers*, 2008
65. Fredriksson, Anders, *Bureaucracy, Informality and Taxation: Essays in Development Economics and Public Finance*, 2009
66. Folke, Olle, *Parties, Power and Patronage*, 2010
67. Drott, David Yanagizawa, *Information, Markets and Conflict*, 2010
68. Meyersson, Erik, *Religion, Politics, and Development*, 2010
69. Klingelhofer, Jan, *Models of Electoral Competition*, 2010
70. Perrotta, Maria Carmela, *Aid, Education and Development*, 2010
71. Caldara, Dario, *Essays on Empirical Macroeconomics*, 2011
72. Mueller, Andreas, *Business Cycles, Unemployment and Job Search*, 2011
73. von Below, David, *Essays in Climate and Labour Economics*, 2011
74. Gars, Johan, *Essays on the Microeconomics of Climate Change*, 2012
75. Spiro, Daniel, *Some Aspects of Resource and Behavioral Economics*, 2012
76. Ge, Jinfeng, *Essays on Macroeconomics and Political Economy*, 2012
77. Li, Yinan, *Institutions, Political Cycles and Corruption*, 2012
78. Håkanson, Christina, *Changes in Workplaces and Careers*, 2013
79. Qin, Bei, *Essays on Empirical Development and Political Economics*, 2013
80. Jia, Ruixue, *Essays on the Political Economics of China's Development*, 2013
81. Campa, Pamela, *Media Influence on Pollution and Gender Equality*, 2013
82. Seim, David, *Essays on Public, Political and Labor Economics*, 2013
83. Shifa, Abdulaziz B., *Essays on Growth, Political Economy and Development*, 2013
84. Panetti, Ettore, *Essays on the Economics of Banks and Markets*, 2013
85. Schmitt, Alex, *Beyond Pigou: Climate Change Mitigation, Policy Making and Decisions*, 2014

86. Rogall, Thorsten, *The Economics of Genocide and War*, 2015
87. Baltrunaite, Audinga, *Political Economics of Special Interests and Gender*, 2016
88. Harbo Hansen, Niels-Jakob, *Jobs, Unemployment, and Macroeconomic Transmission*, 2016
89. Stryjan, Miri, *Essays on Development Policy and the Political Economy of Conflict*, 2016
90. Karadja, Mounir, *On the Economics and Politics of Mobility*, 2016
91. Kitamura, Shuhei, *Land, Power and Technology, Essays on Political Economy and Historical Development*, 2016
92. Malmberg, Hannes, *Human Capital in Development Accounting and Other Essays in Economics*, 2017
93. Öberg, Erik, *On Money and Consumption*, 2017
94. Lane, Nathaniel, *States of Development: Essays on the Political Economy of Development in Asia*, 2017
95. Prawitz, Erik, *On the Move: Essays on the Economic and Political Development of Sweden*, 2017
96. Iwanowsky, Mathias, *Essays in Development and Political Economics*, 2018
97. Dehdari, Sirus Håfström, *Radical Right, Identity, and Retaliation*, 2018
98. Harmenberg, Karl, *Essays on Income Risk and Inequality*, 2018
99. Kilström, Matilda, *Households' Responses to Policy in Labor and Credit Markets*, 2018
100. Meriläinen, Jaakko, *Essays in Political Economics*, 2019
101. Sigurdsson, Jósef, *Essays on Labor Supply and Adjustment Frictions*, 2019
102. Cocciolo, Serena, *Participatory Governance and Public Service Provision*, 2019
103. Mitrunen, Matti, *Essays on the Political Economy of Development*, 2019
104. Olsson, Jonna, *Work, wealth, and well-being: Essays in macroeconomics*, 2019
105. Daroughch, Saman, *Search and Mismatch*, 2019
106. van Vlokhoven, Has, *On the Cost of Capital, Profits and the Diffusion of Ideas*, 2020

107. Foltyn, Richard, *Essays in Macroeconomics and Household Finance*, 2020
108. Lerva, Benedetta, *Worms, Farms and Schools: Three Essays on Investment Choices in Uganda*, 2020
109. Kragh-Sørensen, Kasper, *On Housing, Mortgages, and Taxation*, 2020
110. Karlman, Markus, *Essays on Housing: tax treatment, prices, and macroeconomic implications*, 2020
111. Kinnerud, Karin, *Financial Choice and Public Policy*, 2020
112. Ghisolfi, Selene, *Fairness, technology adoption, water sanitation and pandemic control: Six essays on four topics in Development Economics*, 2020

This thesis consists of three self-contained essays, of which the first two are closely connected and handle individual-level wage stickiness, while the last chapter handles forward guidance by central banks.

Wage stickiness and household heterogeneity introduces wage stickiness on the household level into a standard macroeconomic model with household heterogeneity, and concludes that when realistically modeled, this environment gives rise to severe underemployment spells to the individual households.

Wage stickiness and household heterogeneity in general equilibrium analyzes the general-equilibrium implications of the heterogeneous-agents model with sticky wages developed in chapter 1, and finds that the household-level wage stickiness has a large impact on a worker's precautionary motive to save.

How big is the toolbox of a central banker? uses an event study to investigate if the central bank of Sweden can use its policy-rate forecast to affect market expectations of the future policy rate, and concludes in the affirmative.



Magnus Åhl

holds a B.Sc. with a Major in Economics from Stockholm University, and a M.Sc. in Mathematics from Stockholm University

ISBN 978-91-7911-340-7
ISSN 0346-6892