Title
Essays on Macroeconomics and Monetary Policy

Author
Weber, Jacob Philip

Publication Date
2023

Peer reviewed|Thesis/dissertation

Essays on Macroeconomics and Monetary Policy

by

Jacob P Weber

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Christina Romer, Co-chair
Professor Emi Nakamura, Co-chair
Professor Jón Steinsson

Spring 2023

Essays on Macroeconomics and Monetary Policy

Abstract

Essays on Macroeconomics and Monetary Policy

by

Jacob P Weber

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Christina Romer, Co-chair

Professor Emi Nakamura, Co-chair

This dissertation improves our understanding of how the actions of the Federal Reserve and other central banks affect real activity in the economy. Chapters one and two show how secular change in the composition of investment spending in the United States has weakened the ability of conventional monetary policy tools (which alter the federal funds rate) to affect labor income and consumption. Specifically, Chapter one uses data on software developers collected from GitHub to show that firms are slow to adjust R&D activity and other so-called "intangible investment" in response to changes in interest rates because of congestion in onboarding the workers who produce it. Chapter two takes this result seriously as the explanation for the low observed responsiveness in the cross section for intangible investment as compared to tangible investment. Given the shift towards more and more intangible investment in the U.S. economy, this implies that investment spending overall is becoming less sensitive to changes in interest rates. Combined with two other secular changes documented in Chapter two—a rising import share in investment spending and a decline in the labor share of domestically produced investment—this implies that consumption and labor incomes for hand-to-mouth agents (and thus consumption and labor incomes overall) respond less to monetary policy shocks in general equilibrium. Finally, Chapter three uses novel historical data collected from the Bank of England to show that sterilized foreign exchange intervention (interventions in currency markets that hold policy rates fixed) can affect exchange rates.

Contents

Contents i

List of Figures ii

List of Tables iv

Introduction 1

1 Congestion in Onboarding Workers and Sticky R&D 4
- 1.1 Introduction 4
- 1.2 Simple Congestion Model 6
- 1.3 Evidence for Congestion in Onboarding from GitHub 14
- 1.4 Quantitative Model 26
- 1.5 Conclusion 34

2 Structural Changes in Investment and the Waning Power of Monetary Policy 35
- 2.1 Introduction 35
- 2.2 Investment and Consumption with Hand-to-Mouth Agents 38
- 2.3 Empirical Findings 40
- 2.4 General Equilibrium Response to Monetary Policy Shocks 51
- 2.5 Conclusion 61

3 How Powerful is Unannounced, Sterilized Foreign Exchange Intervention? 63
- 3.1 Introduction 63
- 3.2 Historical and Institutional Context 66
- 3.3 New Archival Data 67
- 3.4 IV Estimates 68
- 3.5 Conclusion 70

Conclusion 75

Bibliography 77

Appendices 87
- 3.6 Appendix to Chapter 1 87
- 3.7 Appendix to Chapter 2 115
- 3.8 Appendix to Chapter 3 119

List of Figures

0.1 Real GDP Growth and SEP Forecasts vs. The Effective Federal Funds Rate . . 2

1.1 With Congestion ($\rho(x)$ Decreasing): Firm Hires Slowly in Response to Positive Shocks . . . 12
1.2 Without Congestion ($\rho(x)$ Flat): Firm Immediately Adjusts to Positive Shocks 13
1.3 Comments On Proposed Contributions Made During Code Review on GitHub . 16
1.4 Distribution of Programming Languages . . . 18
1.5 Becoming Productive Requires Onboarding . . . 21
1.6 Most Work Is Done by Experienced Team Members . . . 22
1.7 Newcomers Either Contribute Once, or Stay a Long Time . . . 24
1.8 Non-parametric Estimate of the Onboarding Function ρ and Linear Approximation 27
1.9 Right-Skewed Endogenous Firm Distribution When ρ Slopes Down . . . 32
1.10 Model Responses to a Contractionary Monetary Policy Shock: Calibrated Congestion Model vs. Standard *ad hoc* Investment Adjustment Cost Model . . . 33

2.1 Compositional Shifts among Investment and Durables, 1947-2020 . . . 43
2.2 Effects of a 1% Increase in the Federal Funds Rate on Real Investment and Consumption (in Ppt.) . . . 45
2.3 Effect of a 1% Hike in the Federal Funds Rate on Investment and Consumer Durables in 1960 vs. 2020 . . . 47
2.4 The Domestic Labor Content of Non-Durables and Services and Components of Investment . . . 49
2.5 Breakdown of Domestic Labor Content into "Domestic Share of Expenditure" and "Labor Share of Domestic Production" by Sector, [1963-2014] . . . 50
2.6 The Production Structure of Section 2.4's Small Open Economy Model . . . 54
2.7 Response of Hand-to-Mouth Consumption and Determinants of Consumption . 58
2.8 Response of Net Exports and Net Foreign Assets . . . 60
2.9 Impulse Reponse of Exports and Imports to Romer & Romer Shocks . . . 61

3.1 The Pound/Dollar Exchange Rate Over Time . . . 71
3.2 Intervention Over Time . . . 72
3.3 The Components of US Fixed Investment . . . 88
3.4 Secular Rise in Software and R&D Investment . . . 89
3.5 Workflow on GitHub: Simplified Example . . . 98

3.6 Non-parametric Estimate of the Onboarding Function ρ 102
3.7 Model Responses to a Contractionary Monetary Policy Shock: Flatter ρ 103
3.8 Non-parametric Estimate of the Onboarding Function ρ with Individual Contributor Fixed Effects 105
3.9 Non-parametric Estimate of the Onboarding Function with Onboarding Success Determined by Junior's Activity Only 107
3.10 Non-parametric Estimate of the Onboarding Function with Onboarding Success Determined by Eventual Project Tenure Only 108
3.11 Becoming Productive Requires Onboarding (Results with 2016 Data) 110
3.12 Most Work is Done by Experienced Team Members (2016Q3 Sample) 111
3.13 Distribution of Programming Languages (2016Q3 Sample) 112
3.14 Newcomers Either Contribute Once, or Stay a Long Time (2016Q3 Sample) . . 113
3.15 Onboarding Requires Attention from Senior Workers (2016Q3 Sample) 114
3.7.1 Compositional Shifts among Investment Components, 1960-2020 115
3.7.2 The Diminishing Effect of a 1% Federal Funds Rate Hike Over Time? 116
3.8.1 Bias in IV under the Null Hypothesis that $\beta_1 = 0$ with Mean Reversion 136

List of Tables

1.1 Calibrated Parameters Common to Both Models 31

2.1 Selected Commodity Composition of Final Demand Components, 1997-2018 . . 42
2.2 Parameters for Section 2.4's Quarterly, Small Open Economy New Keynesian Model. 56
2.3 Decomposition of Dampening of Effects of Real Rate Shocks 59

3.1 Effect of Intervention on the Change in the Exchange Rate by Subsample [1952-1971] 73
3.2 First Stage Regressions: Effect on Intervention (Dollar Sales) by Subsample . . 73
3.3 Descriptive Statistics for Key Variables [1952-1971] 74
3.4 Composition of Non-Residential Investment in 2016 90
3.5 Panel Data on Individual Contributions (Example) 99
3.6 Panel Data on Juniors Joining Projects (Example) 101
3.7 Effect of Intervention on the Change in the Exchange Rate by Subsample [1952-1971] 120
3.8 Effect of Intervention on the Change in the Exchange Rate [1952-1971] 122
3.9 First Stage Regressions: Effect on Intervention (Dollar Sales) 122
3.10 Effect of Intervention on the Change in the Exchange Rate [1952-1971] 123
3.11 Effect of Intervention on the Change in the Exchange Rate by Subsample [1952-1971] 125
3.12 First Stage Regressions: Effect on Intervention (Dollar Sales) by Subsample . . 125
3.13 Effects of Sterilized Intervention on the Log Growth in the Value of the Dollar using Policy Rules and Forecasts Estimated via Adaptive Lasso: Results by Choice of Information Criterion 129
3.14 Effects of Sterilized Intervention on the Log Growth in the Value of the Dollar using Policy Rules and Forecasts Estimated via Adaptive Lasso: Results by Choice of Information Criterion and Subsample 130

Acknowledgments

I thank my committee members, Christy, Emi and Jón, for indispensable advice, discussion, and encouragement. If one measures committee membership by intensity of mentorship and advising, then David Romer ought certainly to be included in this list even if he has somehow escaped responsibility for signing off on this final dissertation.

Of my many incredible co-authors, I thank Justin Bloesch and Alain Naef in particular for allowing me to use our joint work in my dissertation.

I also thank Bob Barsky, Jeff Campbell, and Leo Melosi, whose advice and mentorship while I was a research assistant at the Federal Reserve Bank of Chicago was instrumental in helping me make the transition to graduate study. I thank my professors from my undergraduate studies at Brown University as well: Gauti Eggertsson, without whom I would never have known the Chicago Fed RA job existed, and Sylvia Kuo, without whom I would not have had the courage to apply.

I also gratefully acknowledge the hospitality and financial support I received from the Federal Reserve Bank of San Francisco while finishing this dissertation in the summer of 2022.

Most importantly, I thank my fiancée, Hannah Gribetz, for putting up with my inability to return library books and the other difficulties that inevitably arise when cohabitating with a PhD student. For her support and grace on our long journey together from Brown, to Chicago, and finally to Berkeley, I am eternally grateful.

Introduction

In the wake of the Great Recession (2007-2009), the Federal Reserve took the unprecedented step of cutting the federal funds rate to zero. Given past experience with dramatic changes in U.S. monetary policy in the 1980s and 1930s, policymakers expected such dramatic action to lead to a rapid expansion and recovery.[1] Instead, the recovery from the Great Recession dragged on for years, with policymakers consistently making overly-optimistic predictions for recovery while leaving interest rates at zero between 2009 and 2015. Conversely, upon raising interest rates in 2015, policymakers made overly pessimistic predictions about real GDP growth, as demonstrated in Figure 0.1.

This dissertation proposes a simple explanation for these forecast errors: leaving rates at zero was simply not as expansionary as policymakers expected, and raising rates was not as contractionary as expected, because hitherto unappreciated structural changes in the U.S. economy since the 1980s have made output less sensitive to changes in monetary policy.

One major structural change is that the composition of U.S. investment spending, the most sensitive component of output to monetary policy, has shifted away from traditional investment components (equipment and structures) and towards so-called "intangible" investment (R&D spending and software expenditures). Intangible investment appears to be less sensitive to monetary policy. Chapter 1 argues that this fact is due to high adjustment costs that firms face when adjusting intangible investment, and that these adjustment costs are a deep feature of labor-intensive R&D and other intangible investment. Specifically, Chapter 1 shows that in a simple partial equilibrium model rapidly adjusting R&D investment is costly if the probability of converting new hires into productive R&D workers ("onboarding") is decreasing in the number of new hires ("congestion"). Congestion thus causes R&D producing firms to slowly hire new workers in response to good shocks and hoard workers in response to bad shocks, providing a microfoundation for convex adjustment costs in R&D investment. Chapter 1 closes by using novel, high-frequency productivity data on individual software developers collected from GitHub, a popular online collaboration platform, to provide quantitative evidence for such congestion. Calibrated to this evidence, a sticky-wage new Keynesian model

[1] Nakamura and Steinsson (2018) report anecdotally that when prominent macroeconomists are asked what they find the most convincing evidence for monetary non-neutrality, the most common answers include the evidence presented in Friedman and Schwartz (1963) on the role of monetary policy in the Great Depression and the Volcker disinflation and recession of the early 1980s. Influential statistical analyses available at the time naturally drew most of their power from influential observations in the 1970s and 1980s; see the shock series presented in Romer and Romer (2004) (their figure 1a) and also discussion here in Chapter 2, Section 2.3.3 and accompanying analysis in Appendix 3.7.2.

Figure 0.1: Real GDP Growth and SEP Forecasts vs. The Effective Federal Funds Rate

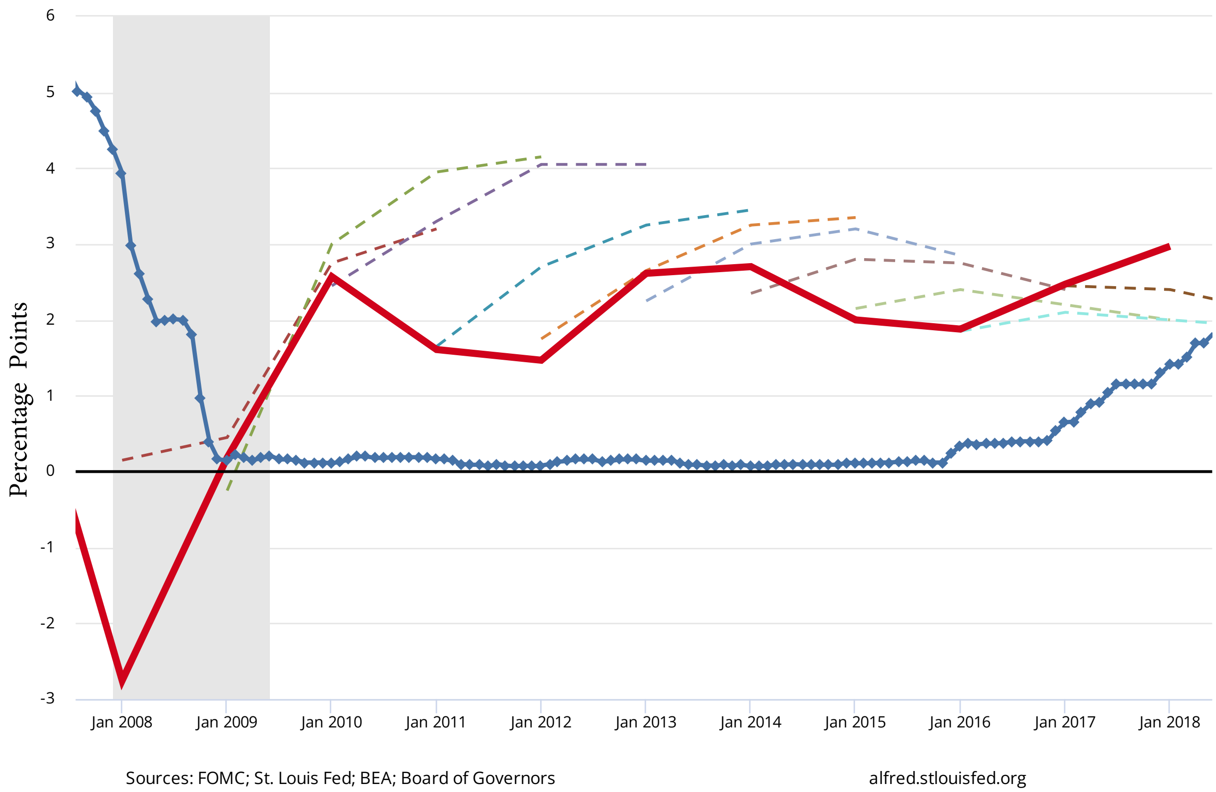

Notes: The thick red line plots Real Q4/Q4 GDP growth, alongside the Summary of Economic Projections (SEP) forecasts produced by the Federal Reserve Open Market Committee (as dashed lines, colored by vintage). The thick blue line with diamonds plots the effective federal funds rate. During the period when the federal funds rate was at zero, approximately 2009-2015, two- and three-year predictions were overly optimistic. After 2015, while raising interest rates, policymakers' projections were too pessimistic relative to the data. One potential explanation for this is that leaving rates at zero was not as expansionary as policymakers expected, and that raising rates was not as contractionary as expected. Figure based on analysis in Chang and Zimmermann (2019).

with heterogeneous investment-producing firms subject to congestion in onboarding and no other frictions yields hump-shaped responses of R&D investment to monetary policy shocks.

Having shown that investment adjustment costs are a deep feature of R&D spending, Chapter 2 formally explores how secular change in both the production and composition of investment goods has weakened private investment's role in the transmission of monetary policy to labor earnings and consumption. Specifically, Chapter 2 demonstrates analytically that fluctuations in the production of investment goods normally amplify the response of consumption to monetary policy shocks by varying labor income for hand-to-mouth agents. However, three secular changes have weakened this channel over time: (i) labor's share of value added in investment goods production has declined, (ii) the import share of investment goods has risen, and (iii) the composition of investment has shifted towards components that are less responsive to monetary policy (e.g., R&D and other intangible investment). A small open economy, two agent new Keynesian model calibrated to match these facts implies a 25% and 15% weaker response of labor income and aggregate consumption, respectively, to real interest rate shocks in a 2010's economy relative to a 1960's economy. A key maintained assumption in this exercise is that the cross-sectional differences observed for different kinds of investment—tangible and intangible—to monetary policy are due to deep, structural features (like high adjustment costs) that are invariant to changes in monetary policy. Chapter 1, by providing empirically-disciplined microfoundations for adjustment costs to changing intangible investment, gives us some reassurance that this critical assumption is correct.

The fact that many central banks found themselves constrained by the effective lower bound on interest rates and facing a slow recovery after the Global Financial Crisis has increased interest in "unconventional monetary policy" through which the central bank tries to stimulate the economy even when conventional short-term policy rates remain unchanged.[2] The final chapter of this dissertation explores unconventional monetary policy: specifically, sterilized foreign exchange intervention, in which a central bank intervenes in foreign exchange markets to attempt to alter the exchange rate while leaving policy rates unchanged. Though most central banks actively intervene on the foreign exchange market, the literature offers mixed evidence on their effectiveness: particularly for unannounced interventions. Chapter 3 uses declassified data from the archives of the Bank of England and the institutional features of the Bretton Woods era to estimate the effects of intervention on the exchange rate. The results suggest that a purchase of pounds equivalent to 1% of the money supply causes a statistically significant, 4-5 basis point appreciation in the pound.

While this dissertation presents new *positive* findings about how monetary policy affects the economy, it makes no *normative* claims about what policymakers facing an increasingly less interest rate sensitive economy should do. Specifically, these results raise important normative questions about the costs and future likelihood of "zero lower bound" episodes as the United States experienced from 2009 to 2015. This dissertation concludes by briefly discussing these normative questions, which point towards directions for future research.

[2]The "Effective Lower Bound" on interest rates was long thought to be zero, and is sometimes still referred to as the "Zero Lower Bound" despite recent evidence that market participants will accept small negative interest rates in exchange for the convenience of not having to hold cash directly.

Chapter 1

Congestion in Onboarding Workers and Sticky R&D[1]

1.1 Introduction

R&D investment, like other kinds of investment, is "sticky": the rate of investment spending is persistent both at the firm level and in the aggregate in response to shocks. To generate this result, a growing literature on intangible investment models R&D spending as subject to convex adjustment costs to the rate of investment spending.[2] More generally, mainstream macro models need these specific adjustment costs to capture the delayed and hump-shaped response of investment to monetary policy shocks. While helpful to fit the data in each case, this critical friction is *ad hoc*, meaning there are few explanations for its source.[3] Further, no proposed explanation focuses on R&D and other "Intellectual Property Products" (IPP) investment, which has grown steadily in importance and is now the single largest component of U.S. fixed investment.[4]

This paper provides an explanation for convex costs to adjusting the rate of R&D and other IPP investment. First, we show in a simple partial equilibrium model how such costs can arise from congestion in onboarding new workers. By onboarding, we mean that new, "junior" workers acquire firm- or project-specific skills on the job in order to transition to becoming productive "senior" workers. Since scarce attention and supervision from existing seniors is necessary for this transition, hiring many juniors at once decreases the probability that juniors successfully transition: a property we call congestion in onboarding. Firms subject to congestion in onboarding optimally hire new junior workers slowly in response to good shocks and hoard senior workers in response to bad shocks. Provided that the shocks affecting the firm are not too big, we show analytically that our congestion model is identical

[1] I thank Justin Bloesch for allowing me to use our joint work in this chapter (Bloesch and Weber, 2023).

[2] See e.g. Moran and Queralto (2018); Bianchi et al. (2019); Cloyne et al. (2022).

[3] See Christiano et al. (2005) and Smets and Wouters (2007) for this friction's importance; Christiano et al. (2018) review proposed foundations for these adjustment costs, which are distinct from intuitive features like convex capital installation costs (Hayashi, 1982), fixed adjustment costs, irreversible investment, etc.

[4] NIPA Table 1.1.5, years 2020 and 2021. Appendix 3.6.1 details the secular trend and components of IPP.

to a model of convex investment adjustment costs, thus providing a microfoundation for them.[5]

Next, we estimate the degree of congestion in onboarding for an important subset of these workers who produce IPP: software developers, who produce about 1/3 of all R&D investment and the majority of IPP investment.[6] We use data on individual software developers collected from GitHub, a popular online collaboration platform boasting over 80 million users across 4 million organizations as of 2022.[7] GitHub tracks the contributions of each user on software projects, documenting who authored each change to the code, allowing us to follow software developers and track their productivity over time on public, open source software projects. The data available from GitHub's Application Programming Interface (API) is on a terabyte scale. Rather than collect this data ourselves, we turn to the GHTorrent project (Gousios, 2013), long-used by software developers to study the productivity of other software developers.[8] Using this dataset, we find substantial congestion: when a project has many juniors joining at the same time, the probability that an individual junior successfully onboards and becomes a productive senior team member declines. The nature of the production process and narrative evidence suggest this stems from the fact that successful onboarding requires attention and supervision from senior workers while the junior worker acquires the project-specific knowledge necessary to contribute, as in our model.

Finally, we embed congestion in onboarding R&D workers in an otherwise standard new Keynesian model where R&D investment is produced by heterogenous firms facing large idiosyncratic productivity shocks. This allows us to consider the effects of monetary policy shocks in general equilibrium while relaxing the "small shocks" assumption made earlier for analytical tractability. We solve for the model's response to monetary policy shocks using sequence space methods (Auclert et al., 2021) and show that our calibrated onboarding frictions generate realistic, hump-shaped impulse responses.

This analysis supports a long-conjectured explanation for the observed stickiness in the empirical literature on R&D: that for firms engaged in knowledge production, substantial firm-specific human capital is bound up in the minds of workers and lost when workers leave. Firms thus behave "as if" they have high adjustment costs (Hall and Lerner, 2010; Kerr and Nanda, 2015). Consistent with this, recent empirical work establishes an important role for team- or firm-specific capital in knowledge creation (Jaravel et al., 2018; Kline et al., 2019). This explanation implicitly assumes that such firm-specific knowledge is difficult to transmit to newcomers – a property we establish as quantitatively relevant for an important subset of R&D workers.

Our empirical results provide a foundation specifically for convex costs to adjusting the

[5] Our use of labor adjustment costs to explain investment adjustment costs reflects the labor intensive nature of R&D, which requires specialized, project-specific knowledge to produce (Hall and Lerner, 2010).

[6] In the NIPAs Software is included in IPP both as R&D and in other subcategories. See Appendix 3.6.1.

[7] See `https://github.com/about` (accessed 10/24/22).

[8] As this public dataset has been largely overlooked by researchers in economics, and may be unfamiliar to many readers, Section 1.3.1 and Appendix 3.6.4 provide a thorough description with citations to more technical discussions published by software developers, which we hope will encourage researchers without a background in software development to work with this data.

rate of investment, which aggregate DSGE models incorporate *ad hoc* to capture the response of investment to monetary policy shocks (Christiano et al., 2005; Smets and Wouters, 2007). This friction is critical, and Smets and Wouters (2007) refer to it as the single most important real friction in improving model fit.[9] The secular rise of R&D and other IPP investment has not reduced the importance of these adjustment costs, as such intangible investment is if anything stickier than traditional tangible investment (equipment and structures): models fitting data for tangible and intangible investment separately find a much larger role for convex adjustment costs on intangible investment.[10] By providing an explanation for why R&D and other IPP investment is costly to adjust, we directly inform models of capital accumulation applied to such intangible investment, and for simpler aggregate models with only one type of investment spending, provide justification for the practice of retaining traditional frictions even as the nature of investment changes.[11] Specifically, we provide evidence that such adjustment costs for R&D and IPP production are "deep" features of the production process invariant to changes in government policy, which is an implicit assumption whenever using models with *ad hoc* adjustment costs to conduct any sort of counterfactual exercise or welfare analysis.

The paper proceeds as follows: Section 1.2 describes the problem of a firm producing a labor-intensive investment good (R&D or other IPP investment) subject to congestion in onboarding in partial equilibrium. We show that in a special case where shocks are sufficiently small, the model is identical to a model of investment adjustment costs. When shocks are large, we work through a numerical example in partial equilibrium to demonstrate that the firm still behaves "as if" it is subject to adjustment costs. Section 1.3 describes the GitHub data and estimates congestion in the onboarding of juniors on open source software projects. Section 1.4 calibrates the onboarding function in Section 1.2's problem to match Section 1.3's estimates and embeds it in an otherwise standard general equilibrium new Keynesian model with nominal wage rigidity and idiosyncratic risk in the production of investment goods. This model extends the partial equilibrium, numerical results of Section 1.2 to a general equilibrium setting, demonstrating that the response to monetary policy shocks is hump-shaped as in a model with convex investment adjustment costs. Section 1.5 concludes.

1.2 Simple Congestion Model

This section develops a simple partial equilibrium model of congestion in onboarding, with three main results. First, under a relatively strict set of assumptions, we show that subjecting

[9] Justiniano et al. (2010) argue that this stems from an overly smooth investment concept (excluding e.g. inventories) but continue to emphasize the critical role of investment in business cycle dynamics.

[10] Moran and Queralto (2018), Bianchi et al. (2019), and Cloyne et al. (2022) fit models to aggregate R&D, estimating much higher investment adjustment costs than for tangible investment (seven, four, and over twenty times as large, respectively). At the firm level, Peters and Taylor (2017) also estimate higher adjustment costs for intangible investment.

[11] Another important set of explanations includes Casares (2006), Edge (2007) and Lucca (2007) who illustrate how extensions of the "Time to Build" formulation of Kydland and Prescott (1982) can yield hump-shaped investment responses or are equivalent to convex adjustment costs.

investment-producing firms to congestion in onboarding yields an optimization problem which is equivalent to the problem of a firm facing convex investment adjustment costs, thus providing a microfoundation for such costs. Second, under a more general set of assumptions, we show numerically that firms subject to congestion in onboarding hire workers slowly in response to good shocks and hoard workers in response to bad shocks. This confirms that firms continue to behave as if they are subject to convex adjustment costs in partial equilibrium.[12] Finally, studying the firm's problem introduces the key, novel feature of the model that we can estimate in the data: the onboarding function ρ. It also formalizes a key testable assumption on the shape of the onboarding function, motivating the empirical analysis in Section 1.3.

We begin by outlining the firm's objective function and constraints. A representative investment-goods firm produces intangible investment (e.g., R&D or software) I_t and sells it to a representative household at price P_t^k.[13] There are decreasing returns to scale at the firm level and labor is the only factor of production. Letting S_{t-1} be the stock of onboarded (s)enior workers, firm output is $I_t = S_{t-1}^{\nu}$ with $\nu < 1$ as in e.g. Anzoategui et al. (2019) and Schmöller and Spitzer (2021).

So far, we have assumed nothing novel. The simplifying assumption that intangible output is produced with labor as the sole factor of production reflects the fact that a distinguishing feature of R&D spending is that the majority is spent on the wages and salaries of "highly educated scientists and engineers" (Hall and Lerner, 2010).[14] Diminishing marginal returns reflects results in Griliches (1990) on the relationship between patents and R&D spending. In practice there is much uncertainty about ν and we will calibrate it to be close to one, as the assumption of diminishing marginal returns is not critical to our results (see Section 1.4). What is critical to obtaining sticky behavior for investment spending I_t, and novel to this paper, is the assumption that the workers who produce it, S_{t-1}, are chosen by the firm subject to congestion in onboarding new workers.

Specifically, we assume senior workers come from junior workers J_t who will successfully onboard with endogenous probability ρ, which we will assume—and then test—is a declining function of J_t/S_{t-1}. The law of motion for S_t is thus

$$S_t \leq (1-d)S_{t-1} + \rho\left(\frac{J_t}{S_{t-1}}\right) J_t, \tag{1.1}$$

[12] By partial equilibrium, we mean that the analysis here considers the firm's response to an idiosyncratic shock holding critical prices, like the wage, fixed. Section 1.4 relaxes this assumption.

[13] I_t could either be accumulated into a capital stock and rented out directly as in a vertical model of innovation (Bianchi et al., 2019) or represent new "ideas" or varieties in a horizontal model of innovation, which produce monopoly rents that the household values at some P_t^k (Moran and Queralto, 2018). Section 1.4 will assume the latter.

[14] See Bloesch and Weber (2021) for estimates of the aggregate labor content of IPP overall, which is similar to construction after accounting for the input-output structure of investment spending. Altering the model to include capital in the production of intangible investment goods I_t would diminish the ability of congestion to explain sticky investment output I_t only to the extent that capital is both (a) substitutable with labor and (b) easy to adjust. We abstract from this possibility.

where $d \in (0,1)$ governs exogenous separations. Our preferred interpretation of endogenous probability ρ is that the onboarding process requires attention and supervision from workers while the juniors acquire firm- or project-specific capital necessary to become productive. Underlying this functional form, we can think of seniors as having a fixed time budget to allocate to onboarding juniors which is less effective when stretched across more and more juniors.[15]

Given these constraints, the firm maximizes the expected, present discounted value of current and future profits. Letting $\Lambda_{0,t}$ be the discount rate between time 0 and t, the firm maximizes

$$\mathbf{E}\left[\sum_{t=0}^{\infty} \Lambda_{0,t}\left[P_t^k I_t - W_t(S_{t-1} + J_t)\right]\right], \tag{1.2}$$

subject to the constraint that $J_t \geq 0$ and where the wage W_t paid to junior and senior workers is assumed to be identical. While unimportant for establishing the correspondence between our model and a model of convex investment adjustment costs, this simplifying assumption highlights the fact that when human capital acquired on the job is firm-specific, wages will not track productivity because workers cannot threaten to take their firm-specific capital to a different employer.[16] To avoid these difficulties with using on-the-job wage growth to infer the acquisition of firm-specific human capital, we turn to productivity data from GitHub. The model abstracts from human capital that is *not* firm-specific for simplicity.

We can gather these assumptions into the following optimization problem: firms choose paths for $\{I_{t+1}, J_t, S_t\}_{t=0}^{\infty}$ to solve

$$\max_{\{I_{t+1}, J_t, S_t\}_{t=0}^{\infty}} \mathbf{E}\left[\sum_{t=0}^{\infty} \Lambda_{0,t}\left[P_t^k I_t - W_t(S_{t-1} + J_t)\right]\right]$$

subject to

$$\begin{aligned}
I_t &= S_{t-1}^{\nu} \\
S_t &\leq (1-d)S_{t-1} + \rho\left(\frac{J_t}{S_{t-1}}\right) J_t \\
J_t &\geq 0.
\end{aligned}$$

We next elaborate a special case in which this problem simplifies to the problem of a firm choosing investment production subject to convex investment adjustment costs.

[15] An alternative foundation for this functional form could be that seniors' time and attention is necessary for on-the-job screening for highly idiosyncratic skills or idiosyncratic match quality, without which juniors will not be productive or cannot be trusted to work independently.

[16] We assume no R&D firm can pay below W_t due to the presence of an outside sector (producing consumption goods, in Section 1.4) which does not face congestion and treats all workers identically, so that any S or J worker can always immediately take a job at W_t in this sector. S workers can still threaten to leave in an attempt to convince the firm to pay $W_t + \epsilon$. The fact that we assume wage growth is zero as workers transition from J to S reflects a limiting case in which S workers have no bargaining power after they onboard ($\epsilon \to 0$) and are hence indifferent between staying, leaving for a job in the outside sector, or leaving to begin anew as a J worker at a different R&D firm.

1.2.1 Congestion in Onboarding and Exact Equivalence

Under some mild assumptions regarding $\rho(x)$, when shocks are small this problem is identical to the problem of a firm choosing the optimal level of investment subject to convex adjustment costs. To see this, assume the law of motion for S binds so that equation (1.1) becomes:

$$S_t = (1-d)S_{t-1} + \rho\left(\frac{J_t}{S_{t-1}}\right)J_t. \tag{1.3}$$

and assume that optimal $J_t > 0$, so that we can ignore the constraint that $J_t \geq 0$. In other words, assume that bad shocks are always small enough that the firm only ever reduces its size by slowing the pace of hiring to below the quantity necessary to replace exogenous separations, and not by implementing a hiring freeze (i.e. $J_t = 0$) or firing senior workers (i.e. choosing $S_t < (1-d)S_{t-1}$). In this case, the following proposition holds:

Proposition 1. *Consider the problem of a firm choosing paths $\{I_{t+1}, J_t, S_t\}_{t=0}^{\infty}$ subject to the law of motion (1.1) and the production function $I_t = S_{t-1}^{\nu}$ to maximize the expected, present discounted value of current and future profits (1.2). In a solution where (1.1) binds always and $J_t > 0$ always, then the firm's problem can be written as:*

$$\max_{\{I_{t+1},J_t,S_t\}_{t=0}^{\infty}} \mathbf{E}\left[\sum_{t=0}^{\infty}\Lambda_{0,t}\left[P_t^k I_t - W_t(S_{t-1}+J_t)\right]\right]$$

subject to

$$I_t = S_{t-1}^{\nu}$$
$$S_t = (1-d)S_{t-1} + \rho\left(\frac{J_t}{S_{t-1}}\right)J_t$$

where $\rho(x) \in [0,1]$ on $x \in [0,\infty)$ and $\rho'(x) < 0$. Let $f(x) \equiv \rho(x)x$ be strictly increasing on some domain D that does not restrict the firm's optimal choice. Then there exists an equivalent maximization problem yielding the same solution for I_t:

$$\max_{\{I_{t+1}\}_{t=0}^{\infty}} \mathbf{E}\left[\sum_{t=0}^{\infty}\Lambda_{0,t}\left[P_t^k I_t - W_t\left(1 + \underbrace{\Phi\left(\frac{I_{t+1}}{I_t}\right)}_{\substack{\text{Convex Adjustment Costs}\\ \text{from Onboarding}}}\right)I_t^{\frac{1}{\nu}}\right]\right]$$

and a domain G which does not restrict firm's optimal choice and where $\Phi' > 0$ on G. Further, if $f''(x) < 0$ on D then $\Phi'' > 0$ on G.

See Appendix 3.6.2 for proof and a discussion which demonstrates that the assumption that $f(x) \equiv \rho(x)x$ is strictly increasing and strictly concave on some interval D does not restrict $\rho(x)$ to some exotic function, and would be satisfied by $\rho(x) = b - ax$ or $\rho(x) = \frac{1}{ax-b} + c$, for example. We will show that $\rho(x)$ is likely better approximated by the latter function (i.e.

ρ is not globally linear) but use the former in our quantitative exercises. The key testable assumption is that $\rho(x)$ is decreasing.

To understand why our investment adjustment costs are denominated in terms of the wage, W_t, note that in an intermediate step we show that $\rho(x)$ decreasing implies the existence of convex *labor adjustment costs* to changing the stock of S workers. To see this, note we can plug in the binding law of motion (1.3) to eliminate J_t and recast the firm's problem in terms of choosing S_t and J_t to maximize

$$\max_{\{I_{t+1},S_t\}_{t=0}^{\infty}} \mathbf{E}\left[\sum_{t=0}^{\infty}\Lambda_{0,t}\left[P_t^k I_t - W_t\left(S_{t-1} + \underbrace{\boldsymbol{\mathcal{F}}\left(\frac{S_t}{S_{t-1}}\right)S_{t-1}}_{\substack{\text{Labor}\\\text{Adjustment}\\\text{Costs}}}\right)\right]\right],$$

subject to the constraint that $I_t = S_{t-1}^{\nu}$ with $\nu < 1$. It can be shown that the labor adjustment cost function $\boldsymbol{\mathcal{F}}(\cdot)$ is an increasing, convex function whose existence and properties rely on a key testable assumption: that $\rho\left(\frac{J_t}{S_{t-1}}\right)$ is decreasing (see Appendix 3.6.2).

Given these convex labor adjustment costs, using the constraint $I_t = S_{t-1}^{\nu}$ to substitute out S_t yields the maximization problem in Proposition 1:

$$\max_{\{I_{t+1}\}_{t=0}^{\infty}} \mathbf{E}\left[\sum_{t=0}^{\infty}\Lambda_{0,t}\left[P_t^k I_t - W_t\left(1 + \underbrace{\Phi\left(\frac{I_{t+1}}{I_t}\right)}_{\substack{\text{Investment}\\\text{Adjustment}\\\text{Costs}}}\right)I_t^{\frac{1}{\nu}}\right]\right],$$

where the investment adjustment cost function $\Phi(\cdot)$ is again convex if $\rho\left(\frac{J_t}{S_{t-1}}\right)$ is decreasing.

The next section explores the implications of decreasing $\rho(x)$ (congestion) in a numerical setting with occasionally-binding constraints which relaxes the "small shocks" assumption made here.

1.2.2 Congestion in Onboarding With Large Idiosyncratic Shocks

Relaxing the assumption that the R&D firm never lays off workers or implements a hiring freeze does not qualitatively change the results. To see this, consider a numerical solution to the firm's problem where prices (P_t^k and W_t) are taken as given and constant, but there is exogenous risk in the production process for R&D. Output is now given by:

$$I_t \equiv e_t S_{t-1}^{\nu},$$

where e_t is a productivity shock which follows a persistent Markov process. For simplicity in this section, we assume e_t only takes on two states: high or low.[17] Finally, we assume the

[17] It may seem superflous to introduce the new variable e_t given that the firm's problem treats changes in P_t^k and e_t as identical shocks to the marginal revenue product of S workers. However, we will need this formulation when introducing idiosyncratic risk in Section 1.4's general equilibrium model with heterogenous firms, where P_t^k is endogenous. Thus, we introduce productivity shocks e_t now.

firm discounts the future at an interest rate $1+r$ also taken as given and constant. We can then solve for the firm's optimal choices given an appropriate calibration. Critically, this calibration assumes $\rho(x)$ is a decreasing function.[18]

A firm at time t with idiosyncratic productivity e_t and incumbent, senior workers S_t has the following value function: plugging in the constraint $I_t = e_t S_{t-1}^{\nu}$,

$$V_t(e_t, S_{t-1}) = \max_{J_t, S_t} \left\{ P^k e_t S_{t-1}^{\nu} - W(S_{t-1} + J_t) + \frac{E_t[V_{t+1}(e_{t+1}, S_t)]}{1+r} \right\}$$

Senior workers separate at rate d and juniors J_t become productive at endogenous rate ρ:

$$S_t \leq (1-d)S_{t-1} + \rho\left(\frac{J_t}{S_{t-1}}\right) J_t$$
$$J_t \geq 0$$

An increase in e_t is a positive shock to the marginal revenue product of the firm's workers. Accordingly, transitioning from the low to the high state will cause the firm to increase in size. The grey arrows in Figure 1.1 illustrate the adjustment of a firm that has been in the low productivity state for a long time transitioning to the high productivity state. Conditional on remaining in the high state, the firm slowly hires new workers, since congestion in onboarding means that hiring many J's at once is costly, eventually converging to the long run optimum given by S_∞.

To show the delayed response to negative shocks, we can work through the opposite case of a firm that has been in the high productivity state for a long time (choosing S_∞ in Figure 1.1) and transitions to the bad state with low marginal revenue products. The firm "hoards" S workers and responds by implementing a hiring freeze ($J = 0$), letting exogenous separations bring the firm to the long run optimum for the low productivity state (S_0 in Figure 1.1). This behavior is optimal because the S workers have option value: if the firm returns to the high state, it will have to pay heavy costs to rebuild the team, and so it avoids letting the size of the team get too small too quickly. Indeed, the subtle kink in the firm's policy function in the low state (the blue dotted line) reflects the point at which setting $J = 0$ sees the firm shrink too quickly, so the firm chooses $J > 0$.

Note that this labor hoarding behavior does not depend on congestion, and would be present even in a standard fixed hiring cost model. However, without congestion, there is a strong asymmetry as the firm adjusts immediately to positive shocks. Figure 1.2 shows this by repeating the exercise in Figure 1.1 but for the case where $\rho(x)$ is nearly constant, i.e. $\rho'(x) \approx 0$, which is identical to a model with a fixed cost of hiring new workers.

This model of congestion in onboarding was motivated by key features of the process of software development observed in GitHub data, which we have shown can map into a model of convex investment adjustment costs given appropriate concavity of $\rho(x)x$. The next section uses data on software developers collaborating on GitHub to investigate whether ρ is a function of J_t/S_{t-1} with $\rho'(x) < 0$ by estimating ρ as a function of J_t/S_{t-1} non-parametrically.

[18] Other than the number of states in the Markov process, the calibration for $\rho(x)$ and parameters ν and d follows the general equilibrium model in Section 1.4. Prices W, P^k and r are calibrated here to the endogenous steady-state solutions that arise from this calibration when solving the model in Section 1.4.

Figure 1.1: With Congestion ($\rho(x)$ Decreasing): Firm Hires Slowly in Response to Positive Shocks

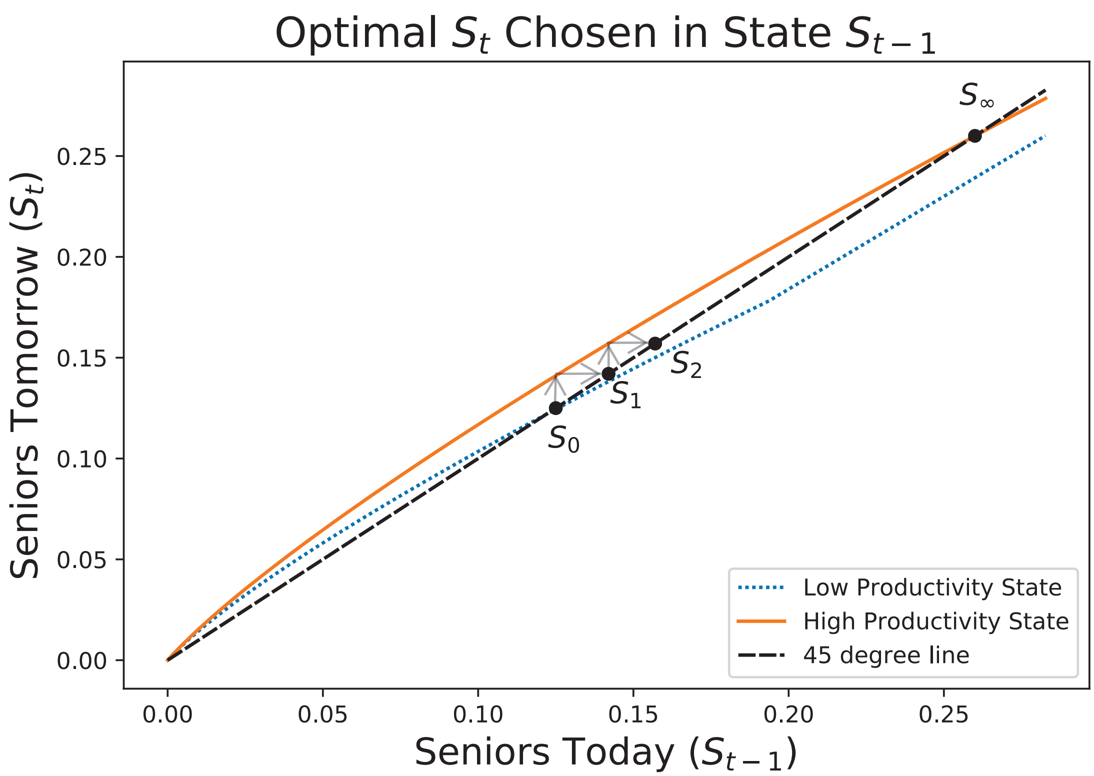

Notes: The firm's optimal choice of senior workers tomorrow, S_t, given seniors today, S_{t-1}. There are two lines because this choice depends on workers' productivity, which can be low or high. The figure illustrates slow adjustment for a firm with S_0 senior workers transitioning from the low productivity to the high productivity state in period $t = 1$. Grey arrows trace out the firm's choices at $t = 1$ and then $t = 2$ assuming it remains in the high state. Note that these choices S_1 and S_2 remain far below the long-run value S_∞. The firm also slowly adjusts to negative shocks by "hoarding labor" in case it transitions back to the good state. Adjustment is slow because the firm implements a hiring freeze ($J_t = 0$) and lets exogenous separations slowly reduce the stock of senior workers S_t.

Figure 1.2: Without Congestion ($\rho(x)$ Flat): Firm Immediately Adjusts to Positive Shocks

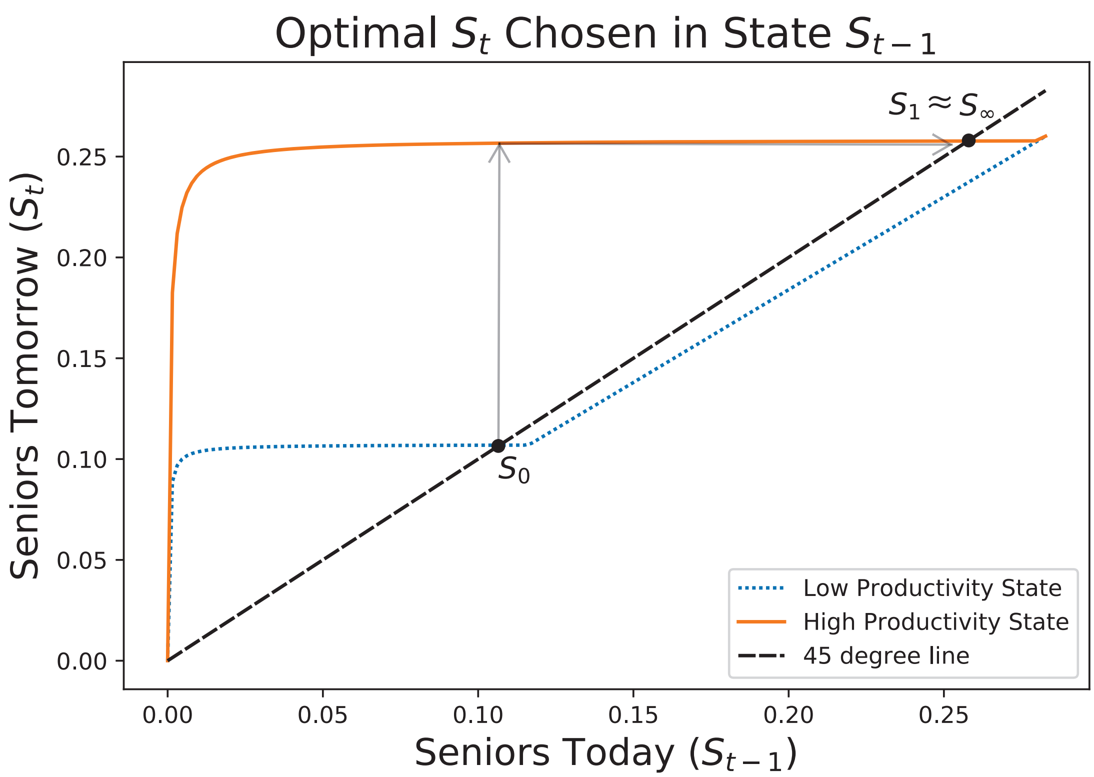

Notes: This figure repeats the exercise in Figure 1.1 for the case where $\rho(x)$ is nearly constant, i.e. $\rho'(x) \approx 0$. This is identical to a fixed cost hiring model, yielding slow adjustment to negative shocks from labor hoarding (not shown), but rapid adjustment to positive shocks (grey arrows).

1.3 Evidence for Congestion in Onboarding from GitHub

Investigating $\rho\left(\frac{J_t}{S_{t-1}}\right)$ requires two steps. First, we identify J workers and S workers. Second, we non-parametrically estimate the probability that a J worker successfully transitions to an S worker as a function of current J_t/S_{t-1} to evaluate the shape of the ρ function.

Leaving aside issues of identification in the second step for the moment, note that measurement of J and S is difficult, since the distinction between J and S that we wish to explore is the acquisition of team- or project-specific capital (which the model collapses to a binary for tractability). Wages imperfectly track marginal productivity increases resulting from the acquisition of this kind of human capital (Caplin et al., 2022; Kline et al., 2019) and need not do so at all as in the limiting case described above where the firm has all the bargaining power. While the limiting case may seem extreme, the empirical prediction that wages do not rise with the initial on-the-job acquisition of project-specific capital seems appropriate for highly educated knowledge workers who are often salaried and/or take compensation as stock options, exercised long after the date of hiring.[19] As we will show, there are substantial productivity gains within the first six months of joining a project in the sample of R&D workers that we study, so that using salaried workers' annual wages to investigate on-the-job productivity growth would be restrictive. To establish this fact and establish a definition for J and S workers that we can use to estimate ρ, we turn to productivity data from GitHub.

1.3.1 GHTorrent Data

GitHub is an online collaboration platform and version control service founded in 2008. It was acquired by Microsoft in 2018 for $7.5 billion USD, reflecting the platform's popularity both for the development of proprietary projects and Open Source Software (OSS). We use data on OSS projects collected systematically from GitHub by Gousios (2013) and made available through Google BigQuery.[20] Collection in GHTorrent began in February of 2012, with information extended back to 2008, and data is available up through 2019Q2.[21] The GHTorrent dataset grows exponentially in size over time and is large (on a terabyte scale). We provide a brief, high-level description of the dataset here; Appendix 3.6.4 provides a detailed description of how the GHTorrent data is structured, accessed, and cleaned by us for the purposes of estimating the regressions described below.

GitHub is the dominant version control service in use today: in a 2021 survey, 91% of software developers globally reported using GitHub for either personal projects or at work.[22]

[19]See e.g. Mehran and Tracy (2001). Sun and Xiaolan (2019) show formally how such long-term wage contracts are optimal when human capital acquired on the job is imperfectly portable (i.e. is firm specific).

[20]GHTorrent is the most popular source for researchers using GitHub data as Cosentino et al. (2016), document. For a comparison of the costs and benefits of other methods, see Mombach (2019).

[21]We access GHTorrent through BigQuery. Since collection began in 2012, we do not have information on projects e.g. created in 2008 and deleted in 2010.

[22]JetBrains conducts an annual "State of Developer Ecosystem" industry survey, which in 2021 included responses from "31,743 developers in 183 countries or regions" (JetBrains, 2021).

While not every company uses GitHub, the production and code review process that GitHub enables – the "Pull/Merge" model of development – is ubiquitous; 84% of developers reported using this model while at work, which makes it nearly as common as email at 90% (JetBrains, 2021). This development process works as follows:

1. A user creates a project ("repository") and allocates power to other trusted users to approve changes (seniors).
2. Potential contributors, junior and senior, propose changes (through "pull requests").
3. Seniors examine the submitted code, leave comments and request alterations before approval ("merging the pull request") in a process called "code review."

Code review is thus an opportunity for juniors to learn how to contribute and signal competence. Over time, a good track record leads juniors to be promoted to seniors. However, juniors do not "graduate" from code review: it is common practice for all code to be at least nominally reviewed, no matter how experienced the contributor, on both OSS projects and in private sector, commercial code development.[23] We thus observe, for each user, their history of attempted contributions to various projects, if and when those changes were approved, and the comments made during code review. Figure 1.3 presents a selection of these comments.

Is Pull/Merge development in OSS projects representative of private sector, commercial development? We consider the following dimensions: the way GitHub and the Pull/Merge model is used; the nature of the users; and the nature of the projects.

Regarding the Pull/Merge model, survey evidence suggests that the way pull requests are used in private GitHub projects—to self-assign tasks and facilitate code review—is identical for both OSS and commercial development. This reflects the fact that most commercial software development is collaborative and that most commercial software developers on GitHub report contributing to OSS projects as well (Kalliamvakou et al., 2015). The now widespread commercial adoption of the Pull/Merge model and OSS development methods for the use of *proprietary* software development (so-called "Inner Source" development) reflects the historical success of the open source model in developing a number of high quality, successful products including Linux, Apache, MySQL and PHP/Perl/Python (Stol et al., 2014). The adoption of these methods was often driven from the "bottom up" by developers who realized they would be helpful for proprietary software development (Stol et al., 2014; Kalliamvakou et al., 2015); see Appendix 3.6.3 for additional detail on the history of industry adoption of OSS development methods.

Regarding users, note that some OSS projects are in fact maintained and developed by paid employees. To understand this, note that while the cost to the firm of making code open source is an inability to charge for it later, open source development creates the opportunity for users to alert the firm to problems (free debugging and testing) or to add features (free development) which benefits the firm when the OSS project is a tool used internally; see e.g. Lerner and Tirole (2005) for a deeper explanation of why firms may want their paid employees

[23] See Kalliamvakou et al. (2015) Figure 1 for developer workflow in commercial projects using GitHub.

Figure 1.3: Comments On Proposed Contributions Made During Code Review on GitHub

body
(In short, `()` + arguments to constructor are missing there in fields definition.) Again, I am not writing code for you. I just show you ideas. You should understand what code is doing and finish things.
Ahh, the coobooks in infra/cookbooks are public cookbooks and I just cloned the locally and copied them over. We could introduce another tool called [Librarian](https://github.com/applicationsonline/librarian) which is like Bundler for chef. I'm cool with
@irium Just push new commits to the *same branch* where you originally sent the [pull request](https://help.github.com/articles/using-pull-requests) from. They'll automatically show up on this page.
In case features.json is not present or broken (for example the first time you run it) the .bak file is used... Should I rename it?
1: D200 One-line docstring should not occupy 3 lines 4: I102 copyright year is outdated, expected 2014 but got 2013 43:16: E203 whitespace before ':' 44:20: E203 whitespace before ':' 45:18: E203 whitespace before ':' 46:18: E203 whitespace before ':' 47:
这个是draggable运行的demo页面，可以不review，发布前会制作更好的demo。
Let's remove this file.

Notes: Code review is not simple yes/no approval. It requires time and attention from seniors as they interact with juniors, giving juniors the opportunity both to learn how to contribute and signal competence. A good track record leads juniors to be promoted to seniors. Source: Pull Request Comments on GHTorrent, accessed via Google BigQuery.

to work on OSS projects, or to make proprietary projects OSS. Consistent with this, on both OSS and "Inner Source" projects within large, private firms, it is widely acknowledged that it is the users of a project who become contributors through discovering bugs or out of a desire to improve functionality for their own purposes (Stol et al., 2014).

Moreover, many government agencies develop code openly and provide it as a public good: Mergel (2015) finds over 7,000 government owned OSS repositories on GitHub (87% of which were for the development of software, as opposed to e.g. sharing data or joint editing of text documents) including the Department of the Interior, NASA, and the Department of Defense.[24] In practice most observed activity on OSS GitHub projects occurs during business hours, dropping on holidays and weekends (Gousios and Spinellis, 2012; McDermott and Hansen, 2021), which suggests that much OSS development happens at work.

However, most contributors are not directly hired to work by an OSS project's owners, who are often private individuals rather than companies or governments. Beyond the fact that volunteers contribute to the OSS projects that they use in order to adapt them for their own ends, as just discussed, motives for volunteer OSS contribution can include learning or reputation building, and turnover on projects is likely high relative to the private sector; see

[24] Our sample of "active" projects as of 2019Q2 includes 36,537 repositories, prohibiting individual inspection. However, we can easily identify some government-maintained projects by filtering for repository names which contain "gov" or ".gov". Fewer than 1% of all repository names contain these strings, which includes repositories belonging to the cities of Boston and Philadelphia, and also a significant UK presence: the Government Digital Service is responsible for over 100 repositories; see `https://github.com/alphagov`.

Vasilescu et al. (2015) for a discussion. This does not mean that most users are students, and indeed most OSS contributors are professionals: survey evidence suggests that the median GitHub user is 29 years old (mean 30) with 8 years of IT experience (mean 10.5) in the United States or Europe (Vasilescu et al., 2015). Moreover, survey evidence generally reveals a contributor's own need for software as the primary reported motivation for OSS contributions.[25] Though their activities have significant positive spillovers to other users, volunteer contributors are not pure altruists.

Finally, OSS projects in our sample are not dominated by personal projects or spam websites. This is because our sample of repositories restricts to large repositories with at least 100 contributions (i.e. merged pull requests) and which we call *projects*.[26] Focusing on such "active" projects with a minimum number of contributions is considered best practice in the literature on OSS software development to ensure that we are isolating projects which are true attempts to collaboratively develop software, although there is no specific guideline for what counts as "active" (Kalliamvakou et al., 2014). Given this, we have checked that our results are robust to changes in this threshold value (we tried 100, 120, and 200; see Appendix 3.6.4 for details). This approach naturally restricts the sample or repositories to large collaborative projects because it is technically possible to work on GitHub without using the pull request model, which effectively adds extra steps to aid in code review. While almost no commercial projects use GitHub this way (Vasilescu et al., 2015) many small or personal projects proceed by making changes ("commits") without pull requests and are thus effectively excluded from our analysis. As Figure 1.4 shows, the most popular primary languages in our sample are Javascript, C, and Python; projects written in CSS or HTML make up only 5% of our sample (e.g., large, jointly-developed websites).

Even after keeping only "large" repositories that contain many merged pull requests, there are a few "test repositories" that do not represent collaborative software development. These are characterized by many pull requests with very short merge times. Similarly, some users are actually bots. These are not difficult to detect and we remove them manually by filtering for repositories with the phrase "test" in the name or with implausibly low average approval times, and by removing users with variations of the name "bot" following Wyrich et al. (2021).[27] However, this highlights the fact that the exact way GitHub is used may vary across both users and projects, which informs our analysis below.

[25]See Hertel et al. (2003) for a study of Linux contributors; Lakhani and Wolf (2003) for various projects on Sourceforge; and Hann et al. (2004) for the Apache project.

[26]Note that Kalliamvakou et al. (2014) warn against conceptually thinking of "repositories" as projects like we do here. This is because certain activity measures, like the number of commits, are mis-measured unless one combines each repository with all related "forked" repositories. A downside is that some forks are indeed new projects. We refrain from combining repositories with their forks for analysis since our measure of time-to-merge for merged pull requests (discussed in Section 1.3.2) does not suffer this measurement problem.

[27]Some GitHub accounts are "Organizational" and stand in for groups of users. We drop such "users" from our analysis.

Figure 1.4: Distribution of Programming Languages

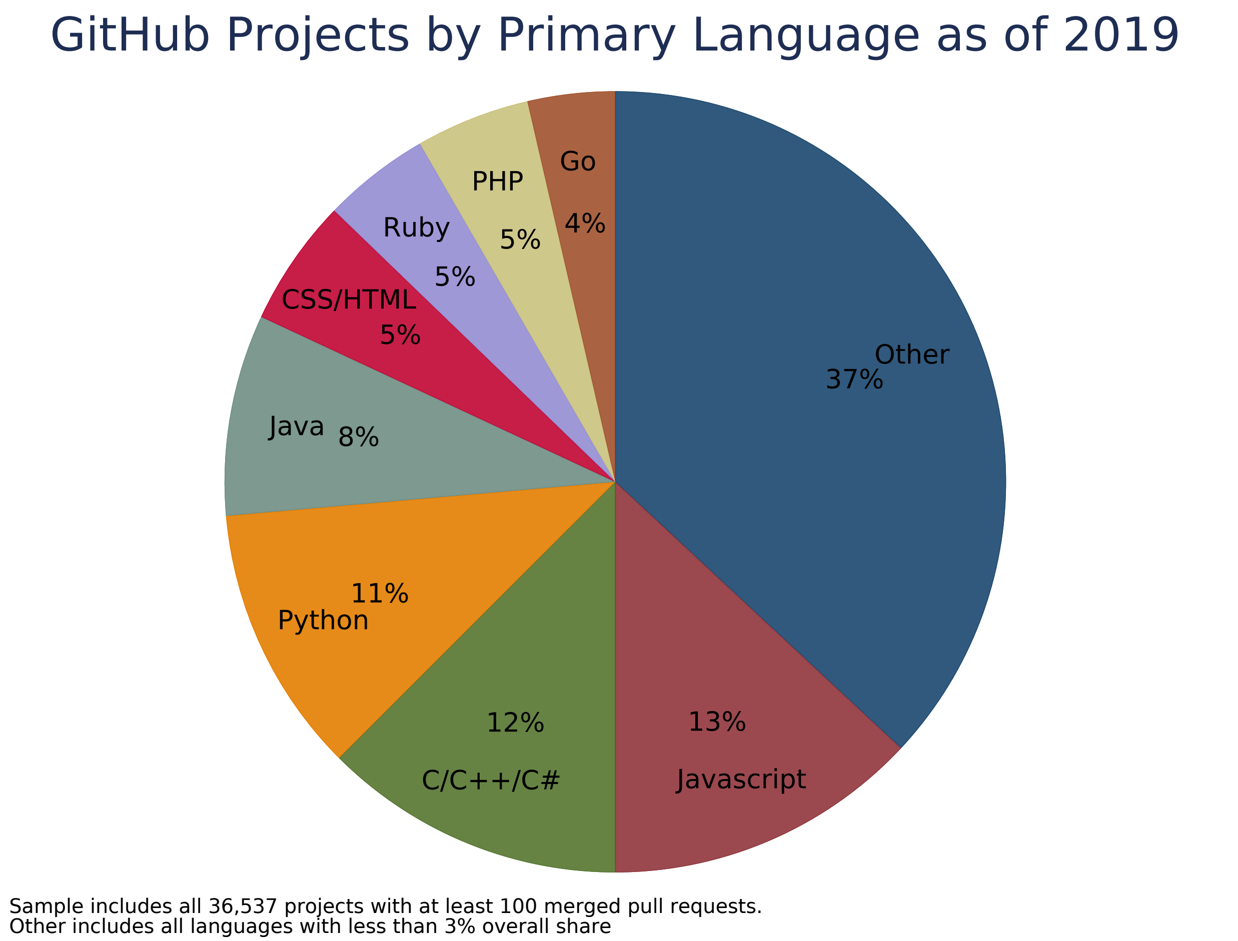

Notes: Weights each project by total number of contributions (merged pull requests). Unweighted results are similar. "Other" includes languages like R and Matlab which are a very small share of the projects in our sample. Source: GHTorrent and authors' calculations.

1.3.2 Onboarding: Identifying J vs. S in GHTorrent

This section estimates how productivity evolves over time on OSS projects in GHTorrent, establishing that there are non-trivial productivity gains over time in the first six months of experience. With this fact in hand, we will define workers J who successfully onboard and become S as those newcomers that remain over six months and/or begin to engage in reviewing the code of other contributors. This definition will then enable us to observe how this onboarding probability varies with the ratio of newcomers to incumbents, J_t/S_{t-1}.

Of the various productivity metrics considered in the literature, we use approval time (i.e. the length of the code review process) for a user's contributions as our measure of that user's productivity.[28] Since this is both a direct measure of how long it takes a user to close an issue and a direct measure of how much "hand-holding" the team thinks a user needs, it is a natural metric to study the onboarding process. As we will show, approval time shrinks dramatically with initial increases in project-specific tenure. Consistent with this interpretation, we examine the number of comments each contribution receives during code review, finding that there is less discussion as juniors gain experience on a project.

Many factors determine approval time beyond individual competence, which motivates the inclusion of controls in our regressions. Forsgren et al. (2021) criticize single-factor measures of performance for the purposes of employee evaluation on the grounds that they are influenced by project-specific factors beyond the control of individual programmers, aligning with prior work on the determinates of approval times in OSS projects from GHTorrent. While changes of good quality and changes that match a project's "roadmap" have a better chance of being accepted, and while a developer's track record can positively influence approval time, project size and complexity also affect approval times. While these can be handled with project fixed effects, there are also project-individual specific features which may cause both longer tenure and faster approval times: for example, a strong pre-existing social connection between the contributor and the project manager. For a survey of papers establishing these facts, see Wyrich et al. (2021). Moreover, a good match in terms of skills between a junior and a particular project (Lazear, 2009) could lead to both longer tenure and faster approval times. An advantage of our setting is that we have rich enough data to estimate individual-by-project fixed effects, controlling for all such confounders.

Finally, we may observe that juniors improve over time on a project because they are acquiring general software development experience. To disentangle the effects of *overall* experience from *project-specific* experience, we control for the overall age of a user's GitHub account, or total tenure on GitHub, in addition to project-specific experience. This is made possible by the fact that we observe the same user working on multiple projects, potentially at the same time, over the course of their career.[29]

[28] This is measured as time-to-merge, or the time between opening and merging a pull request. For an overview of this and other commonly used software productivity metrics, see Forsgren et al. (2021). Time-to-merge also has practical advantages in our longitudinal context, as footnote 26 notes. Commits are also a common metric in the literature, partly reflecting a focus on cross-sectional, project-level analyses; longitudinal studies like this paper's following individual developers are rarer (Cosentino et al., 2016).

[29] It is common for developers on OSS projects to work on several projects at once, and even in firms

Let $y_{i,p,t}$ be either the approval time or total comments received for a contribution opened by user i on project p at time t. We can then estimate the following model via linear regression:

$$\begin{aligned} y_{i,p,t} = &\sum_{j=1}^{13} D(\text{Months Project Experience} = j)_{i,p,t} \\ &+ \sum_{k} D(\text{Months Programming Experience} = k)_{i,t} \\ &+ D_{i,p} + \beta_{PA,p} ProjectAge_{p,t} + \epsilon_{i,p,t}. \end{aligned} \tag{1.4}$$

The first sum consists of dummy variables for having between one and thirteen or more months of experience on project p at time t, and the second sum consists of dummy variables for overall programming experience measured by GitHub account age at time t.[30] We also allow for individual-by-project fixed effects ($D_{i,p}$) and project-specific linear time trends ($\beta_{PA,p} ProjectAge_{p,t}$). See Appendix 3.6.4 for additional detail.

Figure 1.5 uses the marginal effects estimated from equation (1.4) to compare the unconditional mean values for a user with zero months of project-specific experience to predictions for an otherwise identical user with varying degrees of project-specific experience. This reveals that approval time falls precipitously in the first six months of project-specific experience, roughly leveling off thereafter (though standard errors increase). Newcomers also need less "hand-holding" over the same period of time, as the average number of comments per contribution declines for the first six months before leveling off.

Consistent with this, Figure 1.6 demonstrates that most work is done by users with at least six months of experience, though precisely quantifying "work done" is difficult as we do not observe the content of each contribution. Given that large or complex tasks take longer to be approved (Gousios et al., 2014; Wyrich et al., 2021) and that more experienced developers take on more difficult tasks (Torkar et al., 2011; Subramanian, 2020), this should bias our results against finding positive effects from tenure. In light of this, we view our results as a plausible lower bound on the returns to project-specific tenure.

We interpret these documented returns to project-specific tenure as reflecting a combination of skill-acquisition and earned trust or reputation within a team, which our model in Section 1.2 is general enough to encompass. We emphasize the acquisition of project-specific skills, as this frequently arises in interviews with practitioners. Appendix 3.6.3 elaborates on this narrative evidence. All this suggests that attention from incumbents should matter for successful onboarding. We test this in the next section.

where developers are unable to do so, most wish that they could; see Torkar et al. (2011) Appendix B5. More recent survey evidence suggests that most commercial software developers on GitHub report contributing to OSS projects as well (Kalliamvakou et al., 2015), consistent with the view that such multitasking is normal.

[30] This framing reflects the fact that any mis-measurement due to individuals creating accounts at different stages in their career is absorbed by individual-by-project fixed effects.

Figure 1.5: Becoming Productive Requires Onboarding

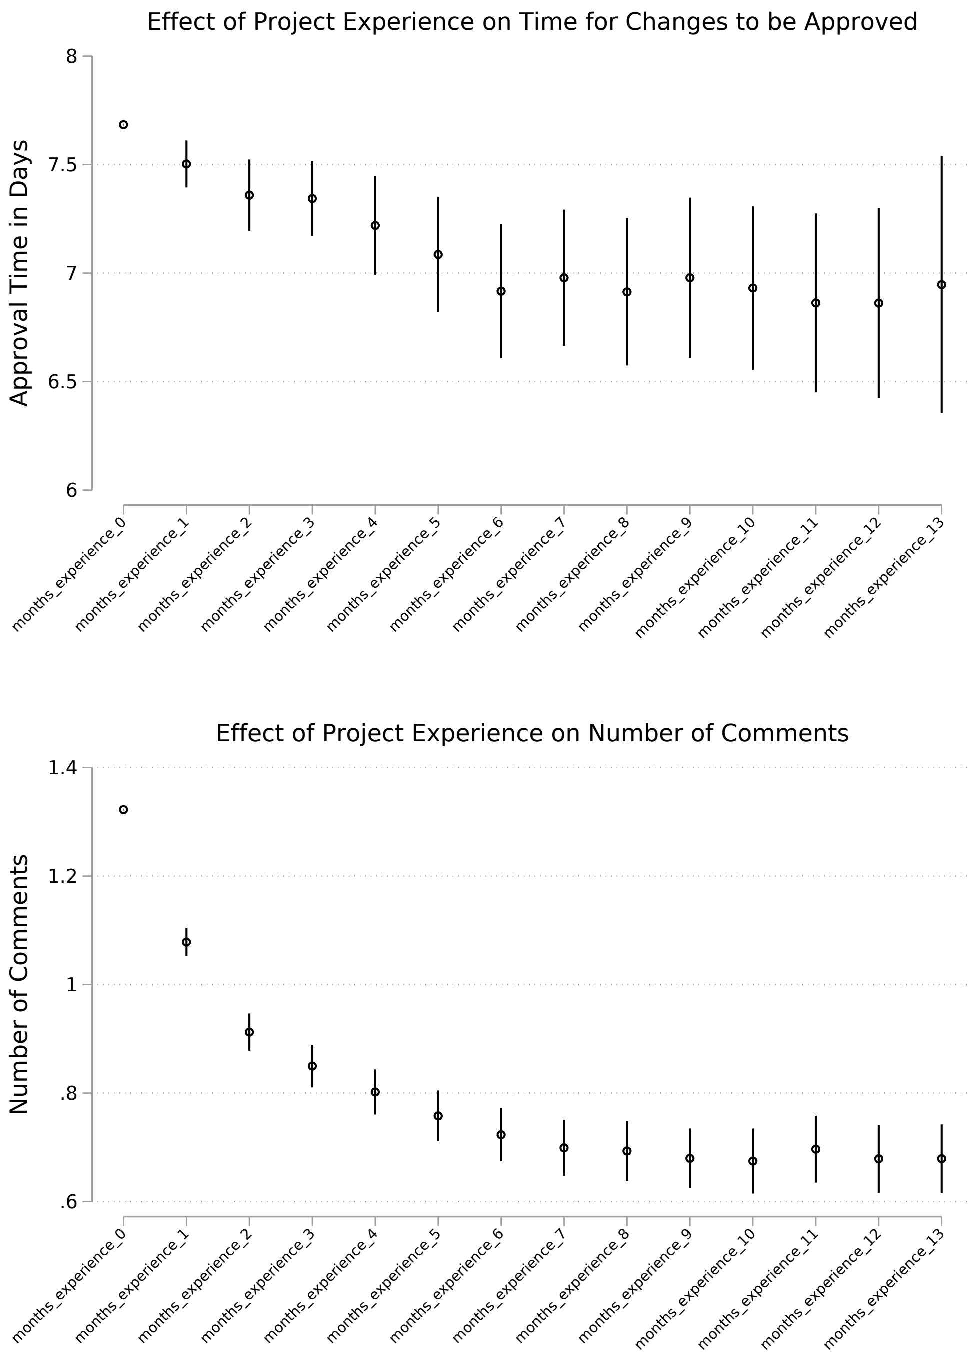

Notes: Over time, new contributors' proposed changes are approved faster, with less discussion. Unconditional mean values for a user with zero months of project-specific experience (first hollow dot) compared to predictions for an otherwise identical user with one or more months of project-specific experience (capped at 13). See text. Standard errors are clustered at the project level. Source: GHTorrent and authors' calculations.

Figure 1.6: Most Work Is Done by Experienced Team Members

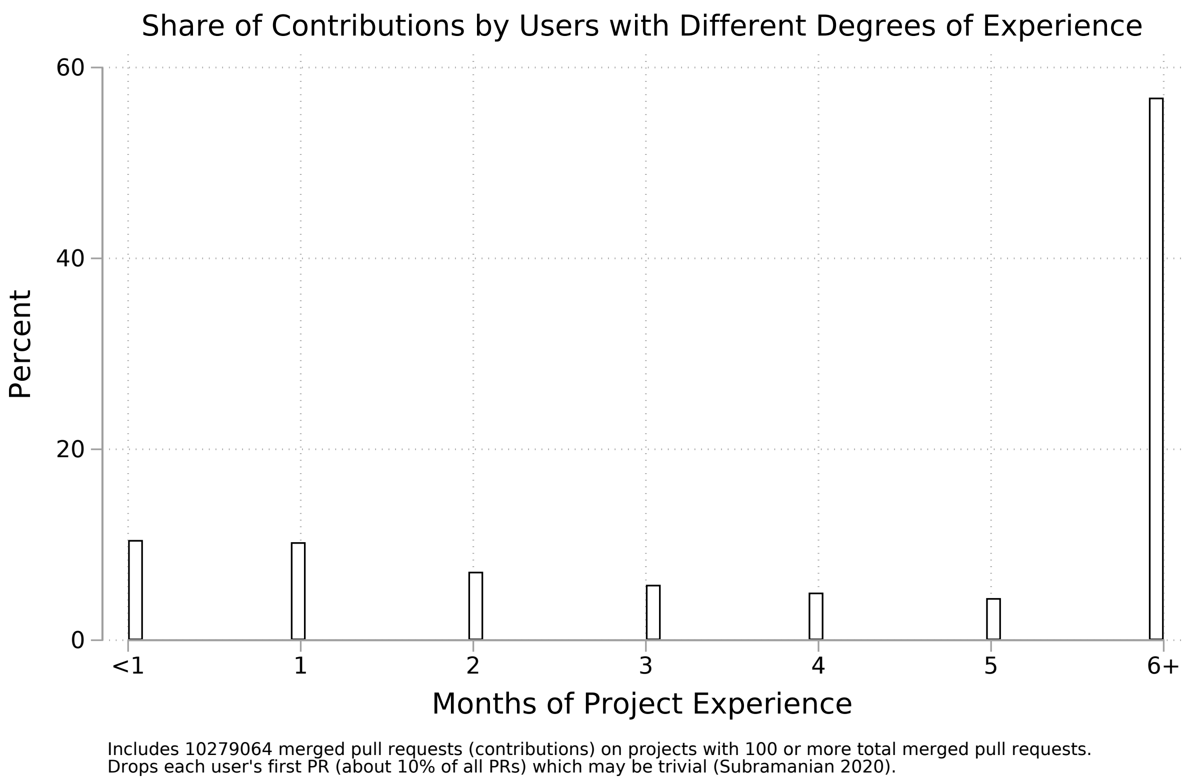

Notes: This figure plots the share of contributions by users with different degrees of project experience at the time of that contribution, showing that most work is done by those who have at least six months of project-specific experience. Since we do not otherwise control for complexity or importance of these contributions, and given that longer-tenure workers take on more complex and important tasks, this figure likely understates the importance of work done by senior contributors. Source: GHTorrent and authors' calculations.

1.3.3 Congestion in Onboarding: Estimating $\rho\left(\frac{J_t}{S_{t-1}}\right)$

In this section, we demonstrate that there is congestion in onboarding by estimating ρ non-parametrically as a function of J_t/S_{t-1}. Specifically, we provide evidence that ρ is a decreasing function, which Section 1.2 shows implies that a firm will behave "as if" it had high adjustment costs to changing the level of investment.

We begin by identifying junior J type and senior S type workers. In each calendar month t and each project p, we assign each user with activity on at least one pull request in p at t into either category J or category S. We drop users who never contribute, and restrict attention to those who will eventually contribute at least once (i.e. open a pull request that is merged). A J type transitions to an S type on a particular project either when they have reached tenure of at least six months, or when we observe them reviewing code written by others. Formally, we identify code review in the data when we observe a user merging/closing/commenting on pull requests authored by other users, and project tenure is measured as the length of time between a user's first observed activity and their last observed activity on a project.

Note that this definition implies that some workers are S types from the beginning – presumably e.g. project founders – and never transition.[31] Our binary definition reflects the fact that a majority of juniors have negligible tenure and contribute precisely once, presumably to fix a bug or add a feature they need, while a nontrivial subset continue to contribute for at least six months. These two groups comprise over 80% of all junior-project observations; see Figure 1.7.

We define the quantity of J types on project p at time t as $J_{p,t}$, tabulated as the number of users who have contributed to that project (i.e. authored at least one pull request that was eventually merged) at time t with less than six months of tenure and who do not engage in code review (i.e. who have not been observed merging/closing/commenting on a pull request opened by someone else). The other active users are summed into $S_{p,t}$. We then estimate a linear probability model: let $\mathbf{1}(i \text{ joining } p \text{ at } t \text{ onboards})$ denote an indicator function for whether a newcomer i on project p (counted in the sum $J_{p,t}$) will eventually transition to being an S type on project p. We then estimate the following via linear regression:

$$\begin{aligned}\mathbf{1}(i \text{ joining } p \text{ at } t \text{ onboards}) &= \sum_b D\left(\frac{J_{p,t}}{S_{p,t}} \text{ in bin } b\right) \\ &+ D_p + \beta_{PA,p} ProjectAge_{p,t} + X_t + \gamma_{i,t} + \epsilon_{i,p,t}. \end{aligned} \tag{1.5}$$

We estimate the effect of $J_{p,t}/S_{p,t}$ non-parametrically by measuring it as a set of dummy variables representing equidistant bins for junior-senior ratios. Project specific dummies D_p control for unobservable project-specific features that may make some projects easier to join, while $ProjectAge$ is a project-specific time trend meant to capture the project life cycle, since some projects may become harder to join as they age; $\gamma_{i,t}$ captures newcomer-specific factors,

[31] Also note that once a J type worker transitions on a project, they are counted as an S type in any calendar month when they "re-appear" on that project in the pull request data.

Figure 1.7: Newcomers Either Contribute Once, or Stay a Long Time

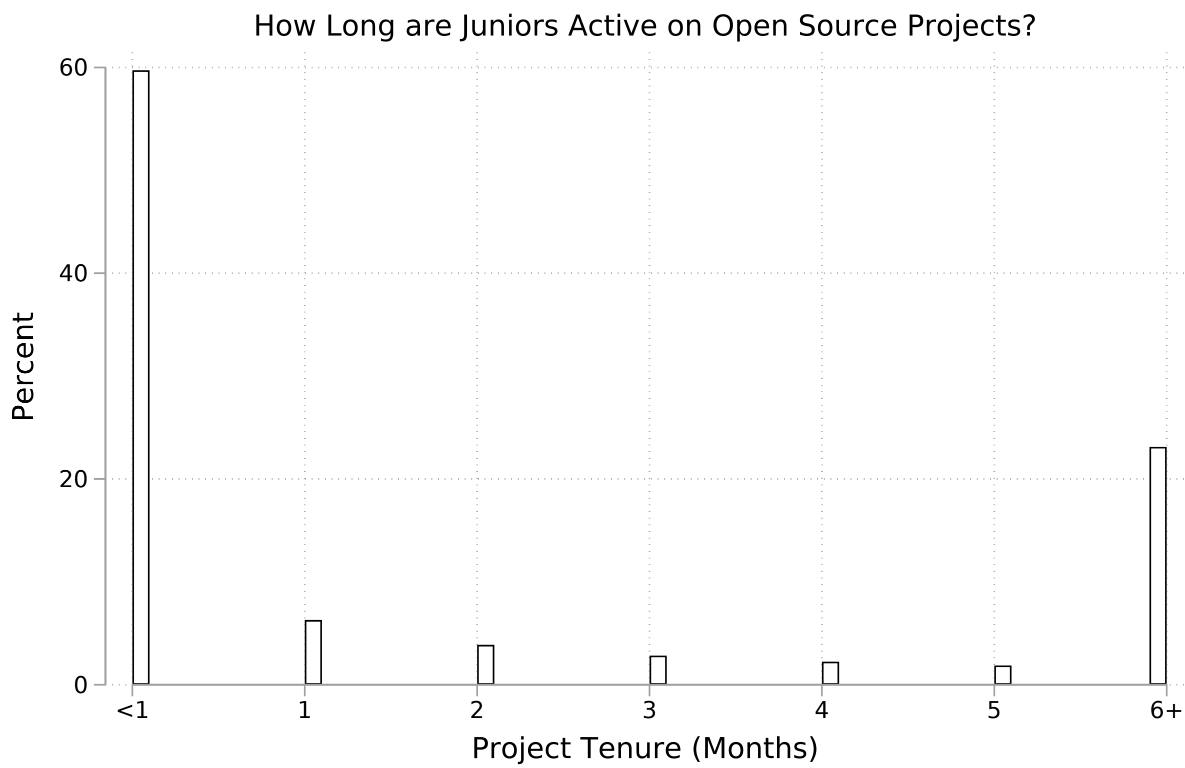

Notes: This figure plots the share of all new, junior contributors on various projects by their subsequent observed tenure on that project. Tenure is measured as the length of time between a user's first observed activity and their last observed activity on a project. Most juniors will have very short tenure (rounded to the nearest month) and contribute once, followed by a nontrivial second group who remain much longer. Source: GHTorrent and authors' calculations.

such as account age, which change over time, and X_t are year fixed effects. See Appendix 3.6.4 for additional detail.

We cannot include user-project specific fixed effects here because they are collinear with the outcome variable (we only observe one outcome per project for each individual: either they onboard, or they do not). Relatedly, we cannot well-estimate individual fixed effects because in practice most individuals join very few OSS projects in sample over time. Note if someone joins only one project in our sample of large OSS projects, we cannot estimate a fixed effect for them. Appendix 3.6.5.2 discusses this and shows that our results are qualitatively unchanged by adding individual fixed effects, though the sample size shrinks.

Figure 1.8 plots the results for equidistant bins of ratios from just above zero to just over 1:1. In practice, over 75% of all project-month observations have $J/S \in [0, 1]$ but there is a significant fat-tail.[32] The figure compares the unconditional mean onboarding probability for a junior on a project in the smallest bin for $\frac{J}{S}$ ratios (first hollow dot) compared to predicted probabilities for an otherwise identical user as the ratio $\frac{J}{S}$ increases (capped at 1.105), and standard errors are clustered at the project level. The results demonstrate that as the ratio of juniors to seniors increases, the onboarding probability falls. Note the jump in the probability for the bin which contains the exact ratio 1:1, which the regression intuitively associates with a relatively higher onboarding probability. Interpreted causally, these estimates literally show us the shape of ρ.

This causal interpretation requires that the ratio J/S be uncorrelated with the error term $\epsilon_{i,p,t}$. In considering potential violations, it seems most natural to worry that juniors not only choose projects which may be easy to join (captured by project fixed effects) but also choose to join projects *at points in time when projects are easy to join.* For example, certain points in a project's development might make for natural "entry points" and our project-specific time trends may imperfectly capture this. Thus, high $\frac{J_{p,t}}{S_{p,t}}$ may occur when newcomers flock to a project at t to take advantage of a high draw for $\epsilon_{i,p,t}$ which is common to many people, and thus correlated with $J_{p,t}$. Thus, it is possible that we are biased towards finding an opposite result, or upward-sloping curve instead of the downward sloping nonlinear relationship in Figure 1.8.

In practice the decision to contribute to an OSS project seems highly idiosyncratic and is often driven by a desire to add needed features or improve functionality for one's own use, as described in Section 1.3. Consistent with this, the inclusion of controls does not do much to change the shape of the relationship in Figure 1.8, as changes in the J/S ratio are not correlated with project characteristics. The fact that a large share of project-month observations occupy the space where $J/S > 1$—where onboarding workers is particularly difficult—further suggests that project maintainers (S) do not have much control over how and when newcomers arrive; indeed, in a model with ρ calibrated to match this data, profit-maximizing firms will generally avoid this region. This highlights an advantage of using OSS projects as opposed to *proprietary* projects: to the extent that firms hire at points in time

[32] See Appendix 3.6.5.1 for results with more bins, capturing this tail. This results in a slightly flatter estimate for ρ which leaves the quantitative results in section 1.4 qualitatively unchanged, though a flatter ρ means less congestion and a less-hump shaped response for investment to monetary policy shocks.

when it is particularly easy to onboard juniors, we would expect this bias towards a flatter or upward-sloping ρ to be severe. Since project maintainers do not have control over when newcomers join, this bias is mitigated in our setting.

To bring Figure 1.8's empirical results into the model, we specify a simple linear functional form for ρ which approximates the nonlinear relationship:

$$\mathbf{1}(i \text{ joining } p \text{ at } t \text{ onboards}) = .47 - 0.7\frac{J_{p,t}}{S_{p,t}} \equiv \rho\left(\frac{J_{p,t}}{S_{p,t}}\right).$$

This is plotted as a blue line in Figure 1.8. We proceed to use this in the next section to illustrate the ability of congestion to generate hump-shaped responses to monetary policy shocks in line with the data.[33]

1.4 Quantitative Model

This section builds a quantitative model where a continuum of firms produce intangible investment subject to idiosyncratic risk and congestion in onboarding. This extends the results in Section 1.2.2 to a general equilibrium setting and shows that congestion continues to produce dynamics which are similar to those in a model with standard investment adjustment costs. Specifically, this section shows that congestion yields a hump-shaped and delayed response of intangible investment to monetary policy shocks. We abstract from standard frictions often used to fit the data (e.g. endogenous capital utilization, habit formation, etc.) to isolate the effect of congestion in creating persistent dynamics in the model. We will solve and compare two different models: one with "I-dot" investment adjustment costs following Christiano et al. (2005) and standard investment production, and one where the only friction in investment production comes from congestion as described above.

In both models, a representative household solves a standard optimization problem. The household accumulates intangible capital K_t through intangible investment $INTAN_t$. We assume this is the only capital used by firms and abstract from tangible capital for simplicity. It also trades a riskless bond in zero net supply B_t which pays real interest rate r. The household solves

$$\max_{\{C_t\}_{t=0}^{\infty},\{B_t\}_{t=0}^{\infty},\{INTAN_t\}_{t=0}^{\infty},\{K_t\}_{t=0}^{\infty}} \mathbf{E}\left[\sum_{t=0}^{\infty}\beta^t\left(\frac{C_t^{1-\sigma}}{1-\sigma} - \omega\frac{L_t^{1+\eta}}{1+\eta}\right)\right] \tag{1.6}$$

subject to standard budget and capital accumulation constraints,

$$C_t + B_t + P_t^k INTAN_t = (1+r_{t-1})B_{t-1} + R_t^k K_{t-1} + W_t L_t + DIV_t \tag{1.7}$$

$$K_t = (1-\delta)K_{t-1} + \left(1 - \frac{\phi}{2}\left(\frac{INTAN_t}{INTAN_{t-1}} - 1\right)^2\right) INTAN_t \tag{1.8}$$

[33] Appendix 3.6.2 discusses other potential functional forms and the implications of our linear choice for the firm's problem.

Figure 1.8: Non-parametric Estimate of the Onboarding Function ρ and Linear Approximation

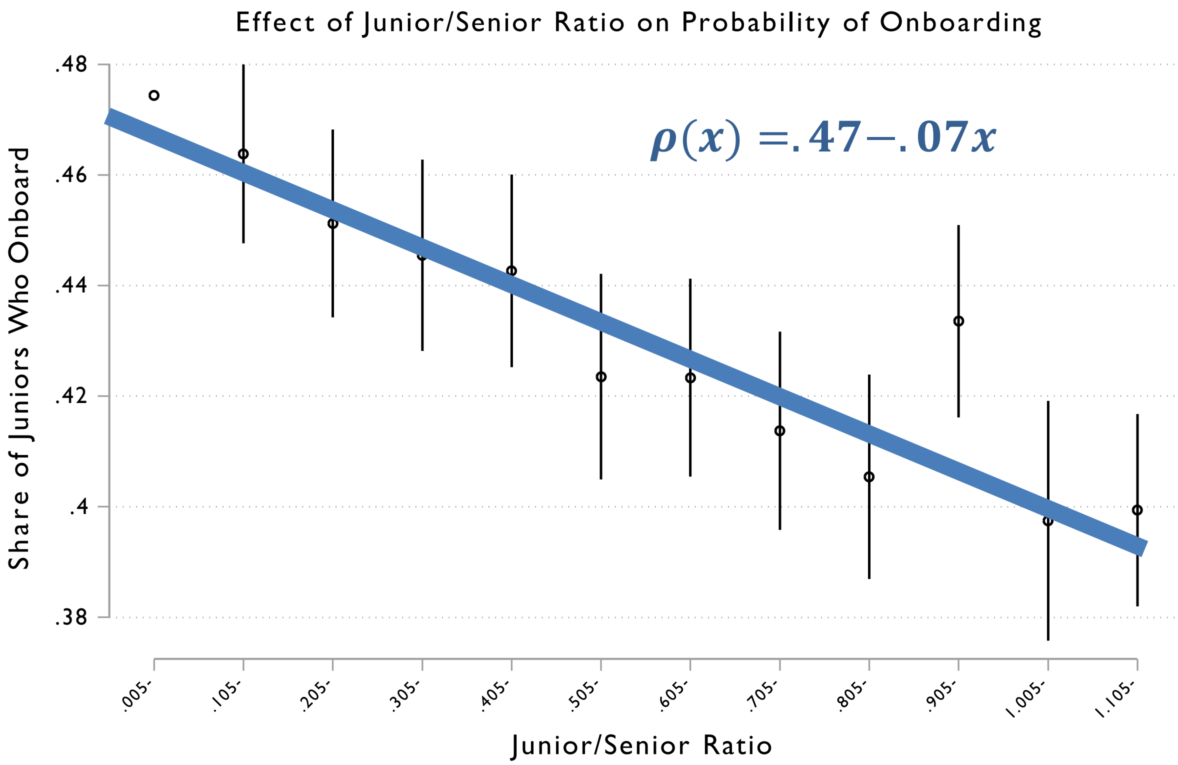

Notes: Unconditional mean onboarding probability for a junior joining a project at date t in the lowest bin for junior to senior ratio $\frac{J}{S}$ at t (first hollow dot) compared to predicted probabilities for an otherwise identical user as the ratio $\frac{J}{S}$ increases (capped at 1.105). We count juniors as successfuly onboarding if they remain with the project for at least six months or begin reviewing code written by others (merging/closing/commenting on pull requests authored by other users). Note the jump in the probability for the bin which contains the exact ratio 1:1 (i.e. bin .905-1.005) which the regression intuitively associates with a relatively higher onboarding probability. The blue line linearly approximates this relationship for use in Section 1.4's calibration. See text. Standard errors are clustered at the project level. Source: GHTorrent and authors' calculations.

where $\phi = 0$ in the congestion model with heterogenous firms. The household earns income from supplying capital K_t and labor L_t to firms, and also potentially from dividends paid by investment-goods producing firms, DIV_t.

A perfectly competitive, representative final consumption good firm produces Cobb-Douglas,

$$C_t = Z_t K_{t-1}^{\alpha} J_{c,t}^{1-\alpha}$$

with the following standard factor demands for capital and labor:

$$R_t^k = \alpha \left(\frac{K_{t-1}}{J_{c,t}} \right)^{\alpha - 1}$$
$$W_t = (1-\alpha) \left(\frac{K_{t-1}}{J_{c,t}} \right)^{\alpha}.$$

A continuum of investment goods firms indexed by $i \in [0,1]$ produce intangible investment. Their problem is formally stated below, and involves hiring juniors $j_t(i)$ and seniors $s_t(i)$ to produce intangible investment. We thus define aggregate labor used in the intangible sector as

$$J_t + S_{t-1} \equiv \int_0^1 (j_t(i) + s_{t-1}(i))\, di,$$

so that aggregate labor demand from all firms is given by

$$L_t = J_{c,t} + J_t + S_{t-1}.$$

Regarding wages, we continue to make the simplifying assumption that all workers receive the same wage W_t. This can be thought of as a limiting case of the bargaining problem between each onboarded S worker and their firm, given other assumptions. To see this, note that any worker can take a job in the perfectly competitive consumption goods sector, which does not face congestion and pays all workers their (identical) marginal product. Since any S or J worker can always immediately take a job at W_t in this sector, no R&D firm can pay below W_t. However, S workers in the congestion model can still threaten to leave the firm and attempt to convince the firm to pay $W_t + \epsilon$. The fact that we assume wage growth is zero as workers transition from J to S reflects a limiting case in which S workers have no bargaining power after they onboard ($\epsilon \to 0$) and are hence indifferent between staying, leaving for a job in the outside sector, or leaving to begin anew as a J worker at a different R&D firm.

To introduce wage stickiness, we assume a wage Phillips curve following Erceg et al. (2000). Denoting gross nominal wage inflation as π_t^w,

$$\pi_t^w(\pi_t^w - 1) = \frac{\epsilon}{\psi} \left(\omega L_t^{1+\eta} - \frac{\epsilon - 1}{\epsilon} W_t L_t C_t^{-\sigma} \right) + \beta \pi_{t+1}^w (\pi_{t+1}^w - 1).$$

Our reliance on this standard formulation (and standard values for ϵ and ψ) for wage stickiness serves the goal of highlighting the role congestion plays in determining aggregate dynamics.

Finally, we assume the central bank sets the nominal interest rate $1 + i_t$ according to a standard Taylor rule. Denoting gross price inflation as $\pi_t \equiv \frac{P_t}{P_{t-1}}$,

$$i_t - i_{ss} = \phi_\pi(\pi_t - 1) + \epsilon_t$$

where ϵ_t is shock following an AR(1) process and ϕ_π determines the responsiveness of the central bank to inflation. The two models we compare differ only in their production of investment goods and choice for adjustment costs, ϕ.

Model 1: Representative Firm with I-dot Adjustment Costs ($\phi > 0$)

The first model assumes simply that all investment firms i are identical and that $INTAN_t = S_{t-1}^\nu$ where S_t is chosen freely each period. J_t is zero here always, so aggregate labor demand is simply $L_t = S_{t-1} + J_{c,t}$. To get hump-shaped impulse response functions, this model needs I-dot adjustment costs with $\psi > 0$.

This model serves as a benchmark for the congestion model, described next.

Model 2: Heterogenous Firms with Congestion in Onboarding

Intangible investment goods firms are owned by households (or, equivalently, a representative venture capital firm that maximizes household utility) and maximize the expected present value of current and future dividends. These firms solve the same optimization problem described in Section 1.2.2, but with new notation since we now consider a continuum of firms $i \in [0, 1]$ optimizing given time-varying prices. These firms choose individual stocks of senior workers $s_t(i)$ and junior workers $j_t(i)$ which aggregate up to total J_t and S_t in the intangible investment sector. They face a common price for their output, P_t^k, but the productivity shock $e_t(i)$ is now firm-specific. This means that the onboarding constraint and non-negativity constraint on J_t will bind for some firms but not others in the stochastic steady state that we linearize around.

Each firm $i \in [0, 1]$ takes the price of intangible capital P_t^k, wages W_t, and interest rates r_t as given. Production of aggregate investment is $INTAN_t \equiv \int_0^1 e(i)s_{t-1}(i)^\nu di$, where idiosyncratic productivity $e_t(i)$ takes on discrete values and follows a Markov process calibrated to match a persistent AR(1) process. A firm with idiosyncratic productivity $e_t(i)$ and incumbent, senior workers $s_{t-1}(i)$ has the following value function:

$$V_t\left(e_t(i), s_{t-1}(i)\right) = \max_{j_t(i), s_t(i)} \left\{ P_t^k e_t(i) s_{t-1}(i)^\nu - W_t\left(s_{t-1}(i) + j_t(i)\right) + \frac{E_t[V_{t+1}\left(e_{t+1}(i), s_t(i)\right)]}{1 + r_t} \right\}$$

where workers separate at rate d and new hires j_t become specialized at endogenous rate ρ:

$$\begin{aligned} s_t(i) &\leq (1-d)s_{t-1}(i) + \rho\left(\frac{j_t(i)}{s_{t-1}(i)}\right) j_t(i) \\ j_t(i) &\geq 0. \end{aligned}$$

Finally, in this model we "turn off" adjustment costs in the household budget constraint and set $\phi = 0$.

1.4.1 Calibration

For our quarterly calibration we choose standard values whenever possible. The household's discount factor is set to $\beta = .99$ and the inverse intertemporal elasticity of substitution is set to $\sigma = 2$ simply to be away from the log case $\sigma = 1$. The elasticity of labor supply is set to $\eta = 1$. The depreciation rate of intangible capital is set at the standard value used in the literature for capital of $\delta = .025$.[34] The capital share of income in the consumption sector is set to $\alpha = .3$. For nominal rigidities, we choose $\epsilon = 10$ and $\psi = 100$ to target a wage Phillips curve slope of 0.1. Finally, we set the Taylor rule parameter to be $\phi_\pi = 1.5$. Table 1.1 summarizes these choices.

For the production of intangible investment goods, we choose $\nu = .95$ implying that production is close to linear in labor s_{t-1}. There is much uncertainty surrounding this parameter, which governs the returns to scale in R&D: structural models fit to aggregate data often require lower estimates ranging from 0.3-0.5 (Moran and Queralto, 2018; Anzoategui et al., 2019; Schmöller and Spitzer, 2021), while Griliches (1990) presents cross-sectional evidence that suggests a wide range inclusive of one may be appropriate. We choose a high number to make it clear that the muted response to shocks in our model is not coming from excessive diminishing marginal returns, as low choices for ν can reduce the volatility of R&D as noted by Comin and Gertler (2006).

In Model 2 with idiosyncratic risk and congestion in onboarding, we assume e_t follows a nine-state Markov process calibrated to match a persistent AR(1) process.[35] We choose a separation rate $d = .08$ to match data on the quarterly separation rate of "Professional, Scientific & Technical Services" workers.[36]

Recall for ρ we use a linear form as described above in Section 1.3 calibrated to $\rho = .47 - .07 \left(\frac{j_t}{s_{t-1}}\right)$. Note that this linear specification does exogenously cap the firms ability to grow at any cost, since at some point it counterfactually predicts that $\rho = 0$, and we exploit this feature during grid search in solving the firm's problem. This limitation not terribly restrictive: in our calibration, this implies optimal choices for J_t/S_{t-1} lie in $[0, \frac{b}{2a}] = [0, 5.625]$ (see Appendix 3.6.2).

1.4.2 Results

We solve the model in sequence space to first order around this steady state with idiosyncratic risk given an exogenous path for a shock to the monetary policy rule, ϵ_t (Auclert et al., 2021). That steady state features an endogenous distribution of R&D firms, which Figure 1.9 plots. The distribution of firm sizes is right-skewed, despite the fact that idiosyncratic shocks are

[34] Intangible capital like R&D depreciates much faster than this (Li and Hall, 2020). Using higher values for δ does not qualitatively change the results.

[35] e_t follows a nine-state Markov process calibrated to discretize $X_t = .95X_{t-1} + \gamma_t$ with $\gamma \sim \mathcal{N}(0, .025)$. For the *two-state* calibration presented in Section 1.2, this implies a high-state productivity of 5% more than in the low state. In practice, the level of idiosyncratic risk barely matters for the aggregate model's dynamics.

[36] This value reflects the average of post-2008, aggregate data from the BLS and LEHD; separation rates were slightly higher prior to this. Using a higher value ($d = .10$) does not materially change the results.

Table 1.1: Calibrated Parameters Common to Both Models

Parameter	Description	Value
β	Household's discount factor	.99
σ	Inverse intertemporal elasticity of substitution	2
δ	Depreciation rate of capital	.025
α	Capital share of consumption goods sector	.3
η	Inverse Frisch labor supply elasticity	1
ϵ/ψ	Slope of the wage Phillips curve	.1

Notes: Standard parameters in the quarterly new Keynesian Model. See text for details.

symmetric, because it is harder for firms to grow than to shrink: firms scale up in the face of positive shocks more slowly than they downsize in response to negative shocks.

Figure 1.10 presents the congestion model's quarterly impulse responses to an contractionary monetary policy shock ϵ_t calibrated to decay at rate of 10% per quarter. These responses are the red, dotted lines in Figure 1.10. The shock causes consumption and inflation (not shown) to fall while the real wage slightly rises due to nominal wage rigidity. In the aggregate, intangible investment firms adjust output by firing juniors J_t, which results in a slow response of seniors S_t and their output, intangible investment $INTAN_t$. Since most workers in the intangible investment sector are senior in steady state, given our calibrated values, the aggregate labor supply response looks more like the hump-shaped response of S workers. Finally, since we are interested in the model's ability to capture the sticky and hump-shaped response of R&D in the data, which is measured at cost, we show that the wage bill of workers in the intangible investment sector $(W_t(J_t + S_{t-1}))$ is also hump-shaped.

The congestion model's responses are comparable to the standard *ad hoc* model of I-dot adjustment costs, which are shown by the blue, crossed lines in Figure 1.10. This shows that the results in Section 1.2 are not reversed in a general equilibrium setting where large idiosyncratic shocks violate the assumptions made in Proposition 1. We conclude that our congestion model provides a highly plausible explanation, or microfoundations, for the investment adjustment costs used in quantitative DSGE models to capture the dynamics of R&D and other intangible investment.

Figure 1.9: Right-Skewed Endogenous Firm Distribution When ρ Slopes Down

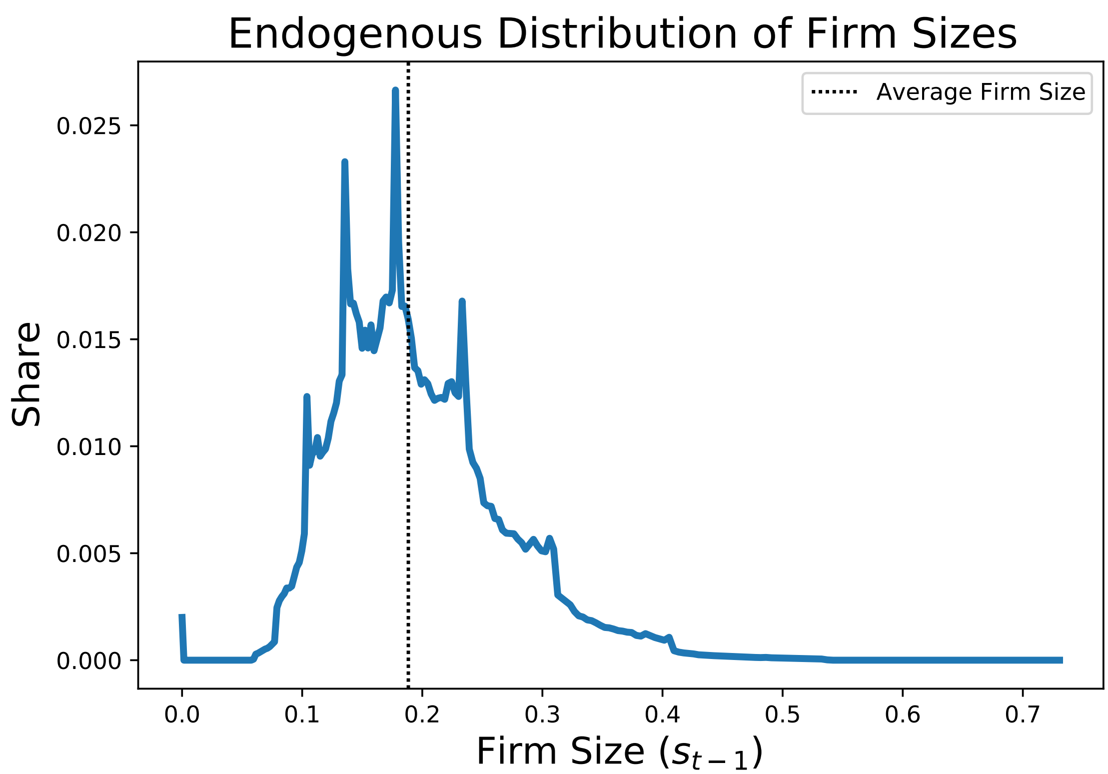

Notes: Endogenous distribution of firms in the congestion model. Firms are ex-ante identical but ex-post different in size due to idiosyncratic productivity shocks which follow a nine-state Markov process. This is calibrated to match a persistent AR(1) process, and the “spikes” in the distribution are the long-run values for firms that have been in a particular productivity state for a long time (and could be “smoothed out” by adding more states). The vertical line plots the average value of S across all firms (the steady state value of S in the model). The distribution is right-skewed because scaling up in the face of positive shocks takes a long time (due to congestion) but layoffs can happen more quickly.

Figure 1.10: Model Responses to a Contractionary Monetary Policy Shock: Calibrated Congestion Model vs. Standard *ad hoc* Investment Adjustment Cost Model

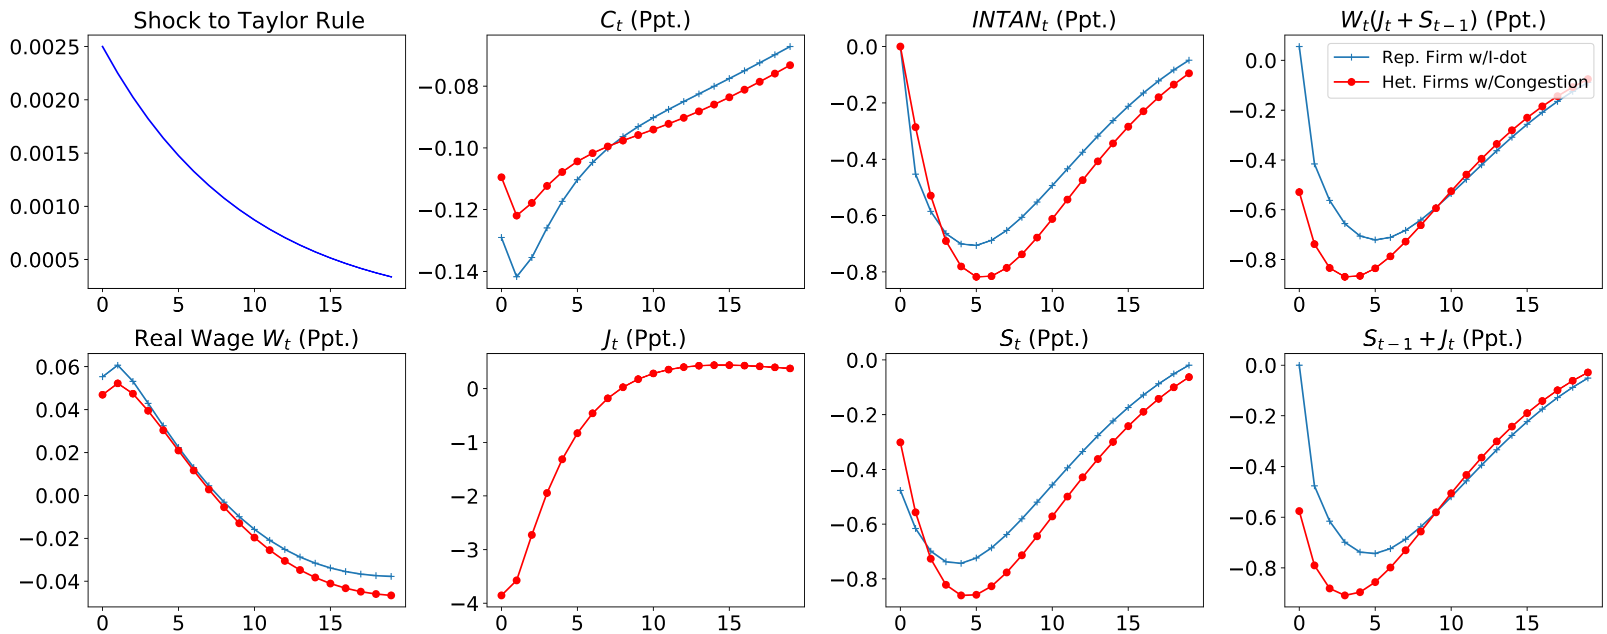

Notes: When ρ slopes down, the congestion model's impulse responses to a contractionary monetary policy shock (red dotted lines) are delayed and hump shaped as in a standard *ad hoc* investment adjustment cost model (blue crossed lines). Quarterly impulse response functions in the model with calibrated $\rho = .47 - .07\left(\frac{j_t}{s_{t-1}}\right)$. See Figure 3.7 in Appendix 3.6.5.1 for results with a flatter ρ function, which results in less sticky investment responses.

1.5 Conclusion

This paper provides a microfoundation for convex adjustment costs to changing the level of R&D investment and other IPP investment, now the single largest component of U.S. investment spending. We showed formally how such costs arise naturally from congestion in onboarding new workers for firms that produce such investment goods. We then provided empirical evidence that such congestion is a significant feature of R&D and IPP production by studying the evolution of individual software developers' productivity on GitHub. Calibrating a specific functional form for our onboarding function to match this GitHub data, we embed it in an otherwise-standard dynamic stochastic general equilibrium model bereft of other real frictions. This model delivers hump-shaped responses of key macroeconomic aggregates in line with the *ad hoc* adjustment costs widely used in aggregate models. By opening up the "black box" of *ad hoc* investment adjustment costs and providing a microfoundation for them, we can confirm that the sluggish adjustment of IPP is invariant to changes in monetary policy. Thus, the common assumption that such investment is sticky *ad hoc* for structural reasons appears appropriate.

This empirical analysis supports a long-conjectured explanation for the observed stickiness in the empirical literature on R&D: that for firms which engage in knowledge production, substantial firm-specific human capital is bound up in the minds of workers and lost if the worker leaves. Firms thus behave "as if" they have high adjustment costs (Hall and Lerner, 2010; Kerr and Nanda, 2015). This paper formalizes and provides empirical evidence on this idea, illustrating how a model of congestion in acquiring firm-specific human capital can map into a model of adjustment costs in the production of investment goods.

Relatedly, note that this paper presents a theory of labor adjustment costs, which we then disciplined on rich data for workers who produce R&D and other IPP investment. Given the nature of the data, this paper focused on explaining the dynamics of such intangible investment. However, the congestion dynamics and narrative evidence presented here seem plausibly applicable to other occupations. Recent work suggests that such congestion, if a broad feature of labor markets, could well-explain the dynamics of unemployment in the aggregate (Mercan et al., 2021). Empirically investigating the extent to which the relative prevalence of congestion and firm-specific capital could explain the relatively muted business cycle dynamics of high-skill employment represents an intriguing path for future work, which our analysis of software developers suggests is promising.

Chapter 2

Structural Changes in Investment and the Waning Power of Monetary Policy[1]

2.1 Introduction

Growing evidence suggests that monetary policy shocks have smaller effects on economic activity now than in the past, even putting aside issues of an effective lower bound on interest rates. Multiple authors, using various empirical techniques, report declining responsiveness of real output and inflation (Boivin, Kiley, and Mishkin, 2010), consumer durables (Van Zandweghe and Braxton, 2013), employment (Willis and Cao, 2015), and investment (Baldi and Lange, 2019) to U.S. monetary policy shocks.[2]

This paper proposes an explanation: secular change in both the production and composition of investment goods has weakened private investment's role in the transmission of monetary policy to labor earnings and consumption. The importance of investment in driving consumption fluctuations in heterogenous agent models where some households have high marginal propensities to consume (MPC's) out of labor income has recently been demonstrated quantitatively by Auclert, Rognlie, and Straub (2020) and analytically by Bilbiie, Känzig, and Surico (2020). In such models, investment amplifies fluctuations in consumption by generating labor income for the high-MPC households. The high volatility of investment in U.S. data means that investment fluctuations play an outsized role in driving consumption fluctuations in such calibrated models.

We revisit this mechanism in a parsimonious, two-agent framework that links the consumption by hand-to-mouth agents to investment. We depart from the analyses of the previous authors in studying an open economy environment, revealing an important role for imports. We show that the consumption response of hand-to-mouth agents to changes in real interest rates depends on (i) the responsiveness of investment, (ii) the size of nominal investment spending relative to nominal consumption spending, and (iii) the extent to which investment

[1] I thank Justin Bloesch for allowing me to use our joint work in this chapter (Bloesch and Weber, 2021).

[2] See also Boivin et al. (2010) for a summary of an older literature on the declining interest rate sensitivity of the U.S. economy.

generates labor income domestically. We label this third term the *domestic labor content*, which we measure from publicly available data. We then show how secular changes have lead to declines in both the responsiveness of investment to real interest rate shocks and the domestic labor content of investment.

We begin by reviewing the composition changes of investment and consumer durables between 1947 and 2020. The most notable change is the rise in "intellectual property products" (IPP) which has grown from less than 1% of GDP around 1950 to nearly 5% of GDP by the beginning of 2020, now accounting for a full fifth of nominal spending on investment and durables. Estimating empirical impulse response functions for each component of investment and durables, we find that IPP is an order of magnitude less responsive to monetary policy shocks than the other components of investment, consistent with firm-level evidence that suggests "intangible" investment spending is relatively insensitive to monetary policy (Caggese and Pérez-Orive, 2020; Döttling and Ratnoski, 2020) and behaves more like a fixed cost (De Ridder, 2019). A simple shift-share analysis implies that the responsiveness of total investment to monetary policy shocks would fall by 20%, under the assumption that the responses of each component are constant over time and as shares change.

Next, we measure the domestic labor content of investment spending and its subcomponents using publicly available Input-Output (I-O) tables. We decompose the domestic labor content into the domestic share of expenditure and the labor share in domestic production, revealing that the domestic labor content of investment and durable goods has fallen from 59 cents on the dollar to 46 cents on the dollar from 1963-2014, driven almost entirely by a decline in the domestic labor content of equipment and durable goods.

To study the effects of these observed trends in a general equilibrium setting, we develop a two-agent, three-sector, small open economy new Keynesian model. Households are split between intertemporally optimizing "Ricardian" agents and hand-to-mouth agents, and production is partitioned between a domestic investment good sector, a domestic consumption good sector, and a traded export good sector. Calibrating the investment good sector's labor and import shares to reflect observed declines in the domestic labor content from the 1960's to the 2010's leads to significant dampening of the response of domestic labor income and hand-to-mouth consumption. Finally, to derive our headline estimates, we compare "1960s" and "2010s" economies calibrated to match observed trends in the domestic labor content of all final demand components, as well as a modest change in the depreciation rate reflecting the shift in composition to shorter-lived IPP investment. This experiment suggests a decline in the responsiveness of labor income and consumption of 25% and 15%, respectively.

Lastly, we find that an increase trade share of the US economy and the possibility of a stronger exports channel of monetary policy does not offset our findings. The model predicts that in response to an expansionary monetary policy shock, net exports *declines* in the short run and later turns positive. This is because rising demand for imports immediately following the shock more than offsets an increase in exports from a weaker exchange rate. We show that this response is supported empirically in our estimated impulse responses and as well as by other research such as Kim (2001).

The remainder of the paper is organized as follows. Section 2.2 uses a minimal number of model assumptions to derive a simple decomposition relating the effect of monetary policy shocks on aggregate consumption to the objects we measure in the data, framing our empirical work. Section 2.3 reviews the changing composition of investment, estimates empirical impulse responses of the components of investment, and documents secular changes in the domestic labor content of the components of investment. Section 2.4 presents a complete two agent, three sector, small open economy new Keynesian model and conducts monetary policy experiments to illustrate how observed changes in the domestic labor content of investment mute the effects of monetary policy on consumption in general equilibrium. Section 3.5 concludes.

2.1.1 Related Literature

This paper is related to recent work using firm-level data to study the response of intangible investment to monetary policy. Döttling and Ratnoski (2020) show that the stock prices of firms with a greater share of intangible capital respond less to monetary policy shocks, citing less of an ability to fund intangible assets with collateral, thereby weakening the credit channel, as well as intangible capital's higher depreciation rate than tangible capital. Caggese and Pérez-Orive (2020) show that firms with more intangible assets depend substantially more on internal savings than collateralized financing, dampening the response of investment to the collateral channel of monetary policy, which Cloyne, Ferreira, Froemel, and Surico (2018) identify as a quantitatively significant channel through which monetary policy affects investment by publically traded firms in the U.S. Our parsimonious general equilibrium model abstracts from these changes in the nature of investment, focusing instead on the decline in the domestic labor content and capturing the reduced interest-rate sensitivity of investment simply through higher adjustment costs; see Bloesch and Weber (2023) for a thorough discussion of investment adjustment costs and intangible investment.

A large body of related work uses I-O tables to study the changing production structure of investment goods. Herrendorf, Rogerson, and Valentinyi (2020) show that investment is increasingly produced in the service sector, and Hubmer (2020) studies changes in the labor content of production for various categories of U.S. final demand going back to 1982 as part of his exploration of the decline in the aggregate U.S. labor share, including investment. The growth implications of the increasing importance of imports in satisfying domestic investment demand has been studied by Cavallo and Landry (2018, 2010) while House, Mocanu, and Shapiro (2017) provide evidence that positive shocks to the demand for investment goods result in substantially higher imports. We depart from these studies by bringing in additional data to extend our analysis with I-O tables back to the 1960s and by focusing on the implications for the transmission of monetary policy. For a related study which uses I-O tables to to study the implications for optimal fiscal policy in the presence of households with heterogenous marginal propensities to consume, see Flynn, Patterson, and Sturm (2020).

Finally, this paper is related to a growing body of work on the importance of indirect effects in the transmission of monetary policy in heterogeneous-agent new Keynesian models (Alves et al., 2020; Kaplan et al., 2018). We diverge from this literature in our focus on documenting secular change in the transmission of monetary policy, which leads us to abstract from many of the model features that prove important for determining the level of the consumption response to monetary policy, such as the distribution of profits and incidence of changes in aggregate labor income across different types of agents.

2.2 Investment and Consumption with Hand-to-Mouth Agents

This section clarifies both the role of investment in amplifying the effects of monetary policy shocks on consumption and how this role is dampened by the decline in the labor share of investment goods, thus motivating our empirical analyses in section 2.3.4 and section 2.3.3. This section also highlights critical assumptions behind the numerical results we present in Section 2.4, which contains a full description of the general equilibrium model.

Consider a model with a unit mass of infinitely-lived households indexed by i wherein a share of households χ are hand-to-mouth, meaning they cannot save and earn only labor income. We call these "Keynesian" households. The remaining share $1-\chi$ we call "Ricardian" households who can save in bonds or capital. Within each type, all households solve the same optimization problem and thus choose the same consumption. We can thus write aggregate consumption (C_t) as a weighted average of consumption across the two types, using $C_{k,t}(i)$ to denote individual Keynesian or hand-to-mouth consumption and $C_{r,t}(i)$ to denote individual Ricardian consumption:

$$C_t = \int_0^{\chi} C_{k,t}(i)di + \int_{\chi}^{1} C_{r,t}(i)di = \chi C_{k,t} + (1-\chi)C_{r,t},$$

where we use the fact that $C_{k,t}(i) = C_{k,t}$ for all Keynesian households and $C_{r,t}(i) = C_{r,t}$ for all Ricardian households. We also assume that the quality and quantity of labor supplied is identical across all households i, so that all earn the same labor income.[3] Letting W_t be the real wage and N_t be aggregate labor supply, aggregate consumption demand is:

$$C_t = \chi W_t N_t + (1-\chi)C_{r,t}, \tag{2.1}$$

which reflects the fact that a share χ of real labor income is always immediately consumed by the hand-to-mouth agents. We further assume that Ricardian agents satisfy the following intertemporal optimality condition in all time periods: letting R_t be the gross real interest rate at time t,

$$u'(C_{r,t}) = \beta R_t E_t[u'(C_{r,t+1})].$$

Solving this standard Euler equation forward, observe that Ricardian consumption today depends only on the expected path of R_t and the long-run value of consumption (which we

[3]Identical labor supply is a common assumption in heterogenous agent models with sticky wages, which allow for idiosyncratic labor income risk by multiplying identical labor supply by a scalar idiosyncratic productivity term; see e.g. Auclert et al. (2020). Allowing for fixed productivity differentials across types would introduce complexity without qualitatively changing the analysis here. See Section 2.4 for details.

assume is unaffected by monetary policy).[4] Thus, (2.1) states that aggregate consumption C_t, given a path of real interest rates pinning down $C_{r,t}$, depends on how *labor income* responds to real interest rates – assuming the share of hand-to-mouth agents is strictly positive ($\chi > 0$).

The relationship of real labor income to real interest rates depends on the economy's production structure. We assume different sectors produce each final demand component: domestic investment and durables consumption, I_t, domestic nondurables and services consumption, C_t, and exports, X_t. Using P_t^k, P_t^c and P_t denote the nominal price in home currency of investment, consumption and exports, respectively, define real investment as $INV_t \equiv P_t^k I_t / P_t^c$ and real exports as $EXP_t \equiv P_t X_t / P_t^c$, obtaining

$$W_t N_t \quad = \quad dlc_t^I INV_t + dlc_t^x EXP_t + dlc_t^c C_t \tag{2.2}$$

where each dlc_t^j denotes the (potentially time varying) *domestic labor content* for sector j. The domestic labor content is defined as the nominal quantity of labor income generated domestically from a dollar of final demand of sector j. As we will see in later sections, the domestic labor content can be expressed as $(1 - m_t^j)(1 - \alpha_t^j)$, where m_t^j is the import share for sector j, and $(1 - \alpha_t^j)$ is the labor share of domestic production for sector j.

Substituting the above equation for labor income into (2.1) and rearranging yields

$$C_t \quad = \quad \frac{\chi dlc_t^I}{1 - \chi dlc_t^c} INV_t + \frac{\chi dlc_t^x}{1 - \chi dlc_t^c} EXP_t + \frac{1 - \chi}{1 - \chi dlc_t^c} C_{r,t}. \tag{2.3}$$

To decompose the effects of monetary policy in the model, consider a sequence of expected, gross real interest rates $\{\mathbb{E}_t[R_{t+j}]\}_{j=0}^{\infty}$ for which our model will yield some equilibrium outcome for C_t and the other endogenous variables. How does C_t change in response to an incremental change in the current gross real interest rate R_t while holding R_{t+j} fixed for $j \geq 1$? Allowing for investment (INV_t) and exports (EXP_t) in equation (2.3) to be functions of R_t, and assuming that the domestic labor content of each sector (dlc_t^j) does not respond much to monetary policy, we take logs and take the derivative of aggregate consumption with respect to the log of the gross real interest rate, such that for any sequence of expected rates, $\mathrm{d} \ln R_t$ is the incremental difference between realized and expected log interest rates in period t:[5]

$$\begin{aligned} \frac{\mathrm{d} \ln C_t}{\mathrm{d} \ln R_t} = & \frac{\chi dlc_t^I}{1 - \chi dlc_t^c} \times \frac{INV_t}{C_t} \times \frac{\partial \ln INV_t}{\partial \ln R_t} \\ & + \frac{\chi dlc_t^x}{1 - \chi dlc_t^c} \times \frac{EXP_t}{C_t} \times \frac{\partial \ln EXP_t}{\partial \ln R_t} \\ & + \frac{1 - \chi}{1 - \chi dlc_t^c} \times \frac{C_{r,t}}{C_t} \times \frac{\partial \ln C_{r,t}}{\partial \ln R_t}, \end{aligned} \tag{2.4}$$

[4]Log-linearizing this equation, as we do in Section 2.4, this statement is equivalent to the insight from the literature on the forward guidance puzzle that current and expected deviations of the real interest rate from the natural rate are the sole determinant of consumption in standard representative agent models where the Euler equation takes this form; see McKay et al. (2016) and Del Negro et al. (2015).

[5]Section 2.4's model assumes both Cobb-Douglas production and that each factor earns its marginal product, implying constant labor and import shares and hence a constant dlc^j for each sector j.

where $\mathrm{d}\ln X_t$ is the deviation of logs of variable X to the baseline in which there were no shocks to the path of expected interest rates $\{\mathbb{E}_t[R_{t+j}]\}_{j=0}^{\infty}$. Thus, the response of consumption can be decomposed as a weighted average of the response of real investment, exports, and Ricardian nondurables consumption to shocks to the real interest rate.

To understand our concern with investment, note that recent evidence provided by Cloyne, Ferreira, and Surico (2020) suggests that in practice the response of non-durables and services consumption with respect to interest rates for households who are not borrowing constrained ($\frac{\partial \ln C_{r,t}}{\partial \ln R_t}$) is trivial in magnitude relative to e.g. investment or durables (see their figure 3). The exports component plays a quantitatively important role but much less so than investment in the US for two reasons. First, the response of exports to interest rates is weaker than the response of investment and durables by nearly an order of magnitude, as we will show empirically. Second, the US is relatively closed, so that the nominal ratio of exports to non-durables and services consumption ($\frac{EXP_t}{C_t}$) is less than half that of investment and durables ($\frac{INV_t}{C_t}$). Thus in any calibration of a two-agent model of the class described here which fits these empirical facts, including the model presented and solved numerically in Section 2.4, the contribution of investment to the general equilibrium response of consumption to real interest rates will be significant.

The next section presents evidence that both the response of investment to real interest rates $\frac{\partial \ln INV_t}{\partial \ln R_t}$ and its multiplier $\frac{dlc_t^I}{1-\chi dlc_t^c}$ have fallen between 1963 and 2014.[6]

2.3 Empirical Findings

We document both that the composition of investment to is shifting towards components that are less responsive to monetary policy and that the domestic labor content of investment goods is declining. Section 2.3.1 describes the data sources. Sections 2.3.2 and 2.3.3 break investment demand into its components and show that an increasing share of investment spending is classified as "Intellectual Property Products", which are quite unresponsive to monetary policy. Section 2.3.4 documents that domestic labor content of investment goods has declined, due to both a rising import share and falling labor share of value added, and that these changes are concentrated in the production of equipment and durable goods.

2.3.1 Data

Composition of Investment We obtain quarterly data on the components of investment and durable goods from 1947-2020 from the National Income and Product Accounts (NIPA).

[6]Another fact that implies weakened monetary policy is the shift towards a services economy, which shows up here as a decrease in $\frac{INV_t}{C_t}$. We do not discuss this here purely because it is not a novel observation: that a declining share of interest-sensitive sectors in output may imply weaker monetary policy is well understood, as pointed out recently in e.g. Summers and Stansbury (2019). Though it is worth noting that this channel is not present in a representative agent economy (to see this, set $\chi = 0$).

Monetary Policy Shocks We use the narrative shock series constructed by Romer and Romer (2004) and updated by Wieland and Yang (2020) from 1967-1998, which Section 2.3.3 describes in more detail.

Domestic Labor Content To compute the domestic labor content of various components of final demand over time, we use a new series of annual Input-Output (I-O) tables released by the BEA in 2016. These tables have a consistent treatment of investment (in particular, for IPP investment) which is critical for our purposes.[7] We also use the import ratios from the BEA's annual Use tables in constructing the domestic total requirements (Leontief inverse) tables for each year. As noted in Horowitz and Planting (2009), the BEA's source data does not allow the BEA to determine how imported commodities are distributed across using industries, and so we impute these as described below. Finally, a shortcoming of the annual BEA tables is that labor payments in each production sector are reported only beginning in 1997.[8] Therefore we rely on Jorgenson, Ho, and Samuels (2017) for consistent estimates of payroll shares by industry from 1963-2014.

2.3.2 Changes in the Composition of Investment from 1947-2020

Throughout this paper, we include consumer durables in our definition of investment, as both the commodity composition of durable goods and their responsiveness to monetary policy closely resemble that of equipment investment. The commodity composition can be seen in Table 2.1. A similar argument applies to residential investment, and we accordingly include residential investment in our measure of investment.[9] Therefore, we will often refer to equipment and durables together, and we will similarly group residential investment with investment in non-residential structures, which we collectively refer to as structures. When including durables and housing, gross nominal investment has maintained a relatively constant share of nominal GDP around 25%[10].

Figure 2.1 shows the changing gross nominal spending on three aggregated subcomponents of private investment and durables as a share of GDP: equipment and durable goods, structures (residential and nonresidential), and intellectual property products. Equipment and durables

[7]For details on the history and construction of these tables, see Lyndaker et al. (2016); for a summary, see section 11.1 of Eldridge et al. (2020).

[8]In the more detailed tables published every five years, labor's share of value is reported starting in 1982 as noted by Hubmer (2020).

[9]This is also partly driven by data limitations in the annual I-O Use tables, which do not distinguish between durables and non-durables consumption or residential and non-residential structures investment. We thus consider "structures" investment as a single category, and create a conservative estimate of durables consumption by assuming that commodities from certain categories (such as autos) allocated to consumption are durables. This leaves us with a conservative underestimate for the durables share of consumption relative to e.g. the NIPA estimates. See Table 2.1 for details.

[10]Many authors have documented that *net* investment has been declining as a share of GDP. Part of this discrepancy is due to the rising depreciation share of nominal GDP, which has risen from 12% to approximately 16% in the time period studied here, as measured by the BEA consumption of fixed capital. The BEA account code is A262RC.

Table 2.1: Selected Commodity Composition of Final Demand Components, 1997-2018

Commodity	Durables	Equipment	Structures	IPP
Computers	13.2	18.5	0.0	0
Motor vehicles, bodies and trailers, and parts	35.4	20.0	0	0
Machinery	1.4	19.1	0	0
Other transportation equipment	3.3	4.7	0	0
Electrical equipment, appliances, components	5.8	2.2	0.4	0
Furniture and related products	7.2	2.9	0.0	0
Miscellaneous manufacturing	15.4	3.9	0	0
Motion picture and sound recording industries	5.4	0	0	6.3
Wholesale	0	13.3	0.2	2.1
Construction	0	0	76.4	0
Real estate	0	0	10.1	0
Support activities for mining	0	0.1	8.3	0
Misc. Prof., Sci., Tech. services	0	3.9	0.2	49.6
Computer systems design & related services	0	1.5	0	23.6
Publishing industries, includes software	12.7	0	0	11.7
Other Commodities	0	9.8	4.4	6.8
Total	100	100	100	100

Notes: Each entry is the share of expenditures of final demand component i (column) spent on commodity j (row), averaged from 1997-2018. We assign commodities used in personal consumption expenditures to durable goods according to whether similar goods are reported in the BEA fixed asset tables, and all commodities considered that we classify as durable are included in this table. The commodity composition of consumer durable goods thus classified most closely resembles that of equipment, though there is some commonality with IPP in the use of "Publishing Industries" and "Motion picture and sound recording industries" commodities.

had been roughly 15% of GDP for most of the post-war period until a discrete drop following the Great Recession, to about 13-14% of GDP. Structures, combining both residential and non-residential decline from approximately 9-10% of GDP to near 7% of GDP. Gross nominal private investment in intellectual products have risen from below 1% of GDP to nearly 5% of GDP. A more detailed figure further decomposing the components can be found in the Appendix (see Figure 3.7.1).

2.3.3 Responsiveness of Investment Components to Monetary Policy Shocks

Investment spending is typically one of the most responsive components of GDP to monetary policy shocks. However, not all components of investment respond similarly. To the extent

Figure 2.1: Compositional Shifts among Investment and Durables, 1947-2020

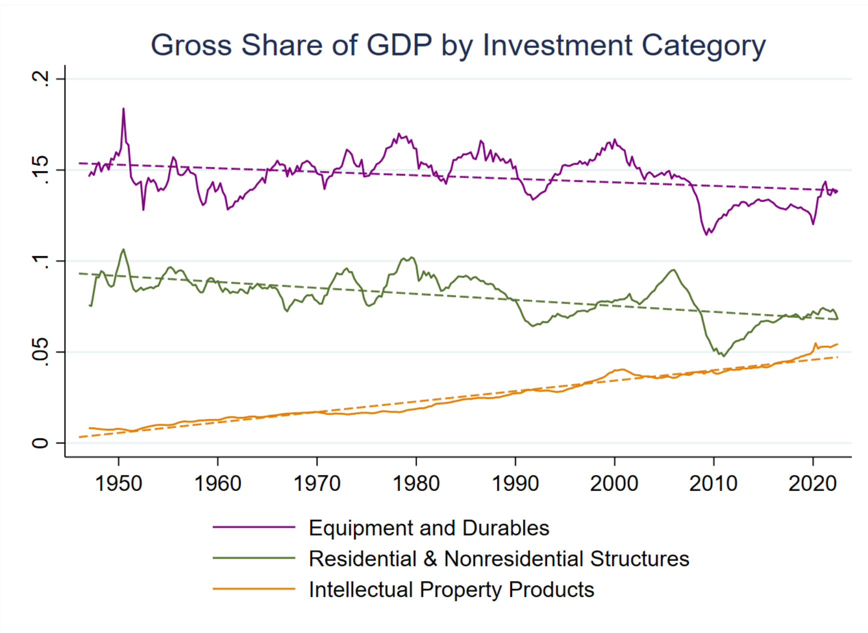

(a) Nominal Investment Shares with Linear Trend Lines

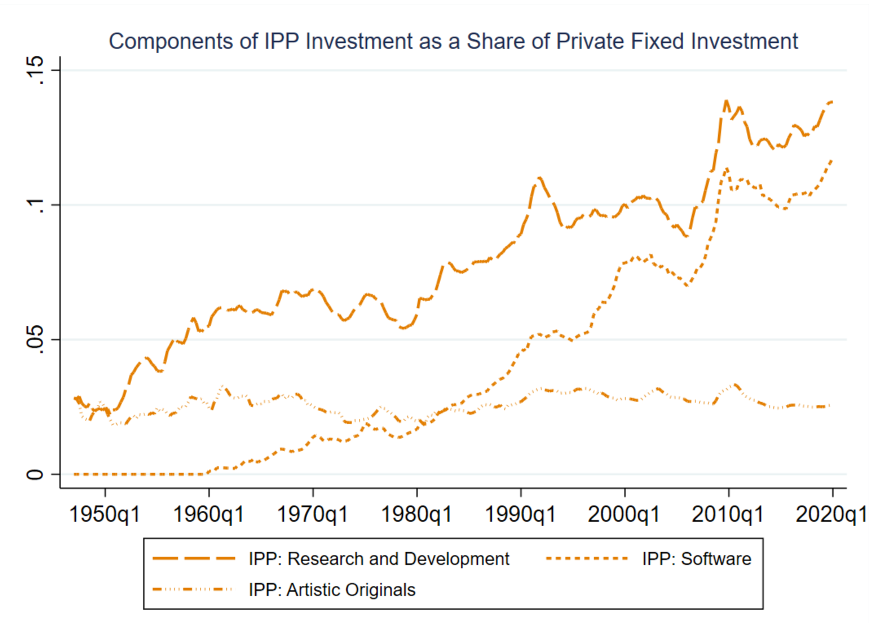

(b) IPP Shares of Investment

Notes: Equipment includes computers in addition to industrial equipment, transportation equipment, etc. Intellectual Property Products (IPP) has risen sharply as a share of private fixed investment, largely due to increases in software and research and development spending. All Data from NIPA Table 5.3.5. except consumer durables which is from Table 2.3.5.

that interest rates affect investment, one would predict that long-lived investments with slow depreciation may be more interest-sensitive. To the extent that collateral values are important for financing investment (Cloyne et al., 2018), types of investment that are difficult to collateralize may respond less to changes in collateral prices due to monetary policy. This suggests that the composition of investment matters for the overall response of investment.

We estimate impulse response functions of different components of investment to the narrative monetary policy shocks constructed by Romer and Romer (2004) and updated by Wieland and Yang (2020). To calculate these, we use the following single-equation regression framework as in Romer and Romer (2004): for endogenous variable Y_t (e.g. the level of real software investment) and exogenous shock variable x_t,

$$y_t = \alpha_0 + \alpha_1 \times trend_t + \sum_{l=1}^{L} \beta_{y,l} y_{t-l} + \sum_{j=1}^{J} \beta_{x,l} x_{t-j} + \sum_{s=1}^{4} \gamma_s D_{s,t} + \mu_t \tag{2.5}$$

where y_t is the log difference of Y_t and the $D_{s,t}$ are quarterly dummies used to deseasonalize the data. Following the arguments in Baek and Lee (2020), we choose the maximum lag length of our exogenous shocks $J = 12$ quarters to be the maximum horizon for our impulse response functions and use the BIC to select L subject to the restriction $L \leq J$. Following e.g. Cloyne et al. (2020), our specification allows for a linear time trend variable $trend_t$ (though in practice the BIC rarely selects this specification). We calculate one and two standard error bands using Monte Carlo methods as in Romer and Romer (2004) and use the same sample period of years 1968 to 1998 for reasons we will soon discuss.

Figure 2.2 plots the impulse response functions of various components of investment. Panels (a), (b), and (c) show the response of residential investment, non-residential structures, and equipment which peak at -15, -5, and -10 percentage points, respectively. In contrast, the response of intellectual property products bottoms out at -1 percentage point. Panels (e) and (f) show the response of durable goods and non-durable goods, demonstrating that the response of durable goods is a similar order of magnitude as equipment and other structures, while non-durables exhibit a much weaker response. The response of PCE services, not shown here, is similarly small.

How much weaker is the aggregate investment response to monetary policy today, given observed shifts in the composition of investment spending? The most natural excercise, which we present in the Appendix Figure 3.7.2, involves splitting the sample and estimating the effect of monetary policy before and after some date. While the point estimates do suggest monetary policy has become weaker, we hesitate to emphasize these results for the following reasons: first, low power and large standard errors make it difficult to read too much into the large differences in point estimates, since it is well-known that much of the variation in our policy shock series comes from the early period. Second, qualitative differences between the two series raise some concerns that endogeneity in the monetary policy shock series we use may be a greater problem in the later subsample.[11] In short, it is not clear that the narrative shock series consistently captures exogenous variation in the federal funds rate over the entire

[11]The identification in Romer and Romer (2004) rests on the assumption that no new information is incorporated into the FOMC's policy decisions between the creation of the Greenbook forecasts and the policy

Figure 2.2: Effects of a 1% Increase in the Federal Funds Rate on Real Investment and Consumption (in Ppt.)

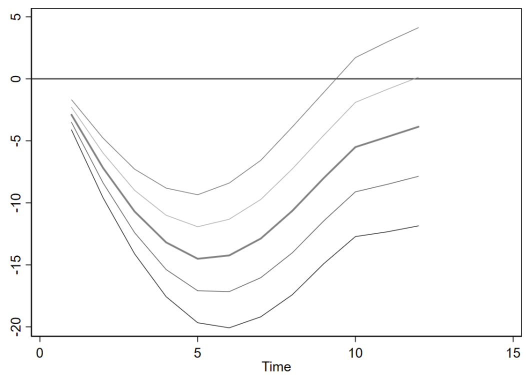

(a) Residential Investment

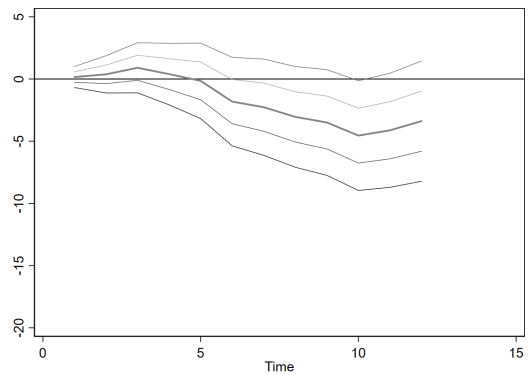

(b) Non-Residential Structures

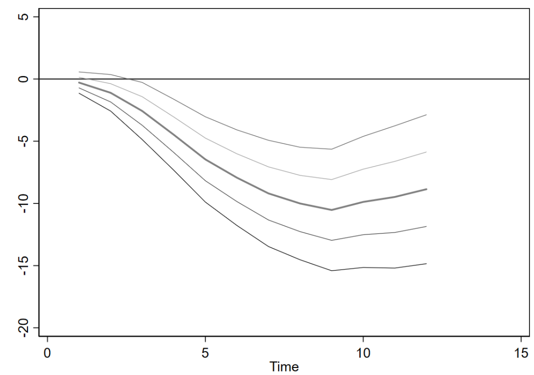

(c) Equipment

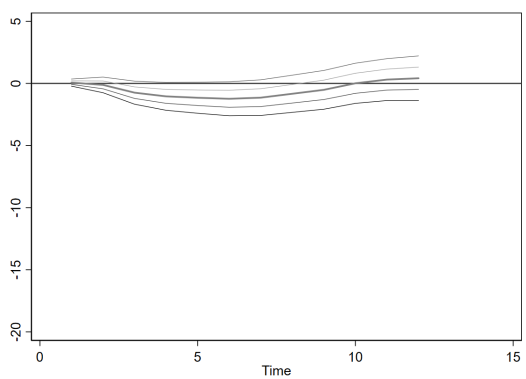

(d) Intellectual Property Products

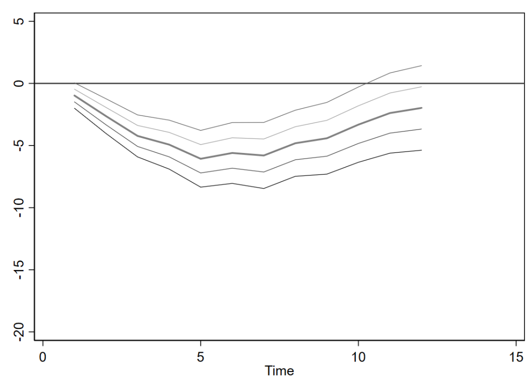

(e) Consumer Durable Goods

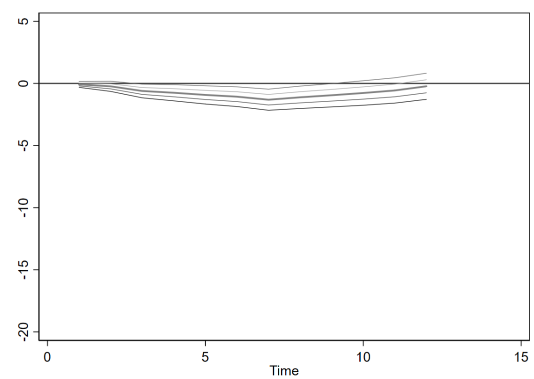

(f) Consumer Non-Durable Goods

Notes: "Time" is quarters after the shock. One and two standard error bands are plotted, calculated as in Romer and Romer (2004). The sample period is 1969-1996 inclusive, as in that paper. See text for estimation details.

post-war period, complicating any interpretation Figure 3.7.2 as evidence that the power of monetary policy has diminished.

These concerns motivate the following exercise, which uses impulse response functions estimated on the original time period from Romer and Romer (2004): note that aggregate, real investment and durables spending, $INV_t \equiv P_t^k I_t / P^c$, is the sum of L components D_t^l,

$$INV_t \equiv \sum_{l=1}^{L} D_t^l,$$

where in practice the D^l include several categories of spending on equipment and IPP investment, non-residential structures, durables, and residential investment all deflated by P_t^c.[12] We may always write percent changes as

$$\frac{INV_{t+\tau} - INV_t}{INV_t} = \sum_{l=1}^{L} \left(\frac{D_{t+\tau}^l - D_t^l}{D_t^l} \right) \frac{D_t^l}{INV_t},$$

where the right hand side demonstrates a relationship between changes in the components of investment and aggregate investment spending. Thus, estimated responses for each component of investment (inclusive of durables) to monetary policy imply a response for aggregate investment which depends upon the initial share of each component at the time of the shock.

Approximating percentage changes in D^l with the difference in logs, we can estimate the effect of monetary policy on each component using the procedure outlined above.[13] Assuming that these responses have not changed over time, Figure 2.3 plots implied impulse response functions for a 1960s economy (with the $\frac{D_t^l}{INV_t}$ set to the values observed in $t = $ 1960q1) and compares it to the implied impulse response function in a modern economy, with $\frac{D_t^l}{INV_t}$ set to the values observed in $t = $ 2020q1. This simple shift-share exercise suggests that the peak effect of monetary policy on investment has fallen from -7.5% to -6.0%, a 20% decline. This reflects the rising share of IPP investment, from zero to roughly a fifth of nominal investment spending, which is nearly unresponsive to monetary policy.

2.3.4 Falling Domestic Labor Content of Equipment and Durables

The previous section investigated how the changing composition of investment may affect the response of aggregate investment to monetary policy shocks, $\frac{\partial \ln INV_t}{\partial \ln R_t}$. This section explores

meeting itself. This assumption may be more grossly violated in the later sample; thus, Figure 3.7.2 also considers a specification of equation (2.5) which replaces the narrative shock series x_t with the high-frequency shock series from Nakamura and Steinsson (2018) when available, with little effect on the results.

[12] Specifically, we use the finest-available categories for equipment and IPP investment in NIPA table 5.3.5, in addition to non-residential structures, residential investment, and finally consumer durables from NIPA table 2.3.5, and then deflate by core PCE (Fred series PCEPILFE).

[13] Note that we use a different definition of real investment spending than in the Figure 2.2, using core PCE as a common deflator accross investment types, both to align more closely with the right hand side of equation (2.4) and to permit aggregation using nominal shares.

Figure 2.3: Effect of a 1% Hike in the Federal Funds Rate on Investment and Consumer Durables in 1960 vs. 2020

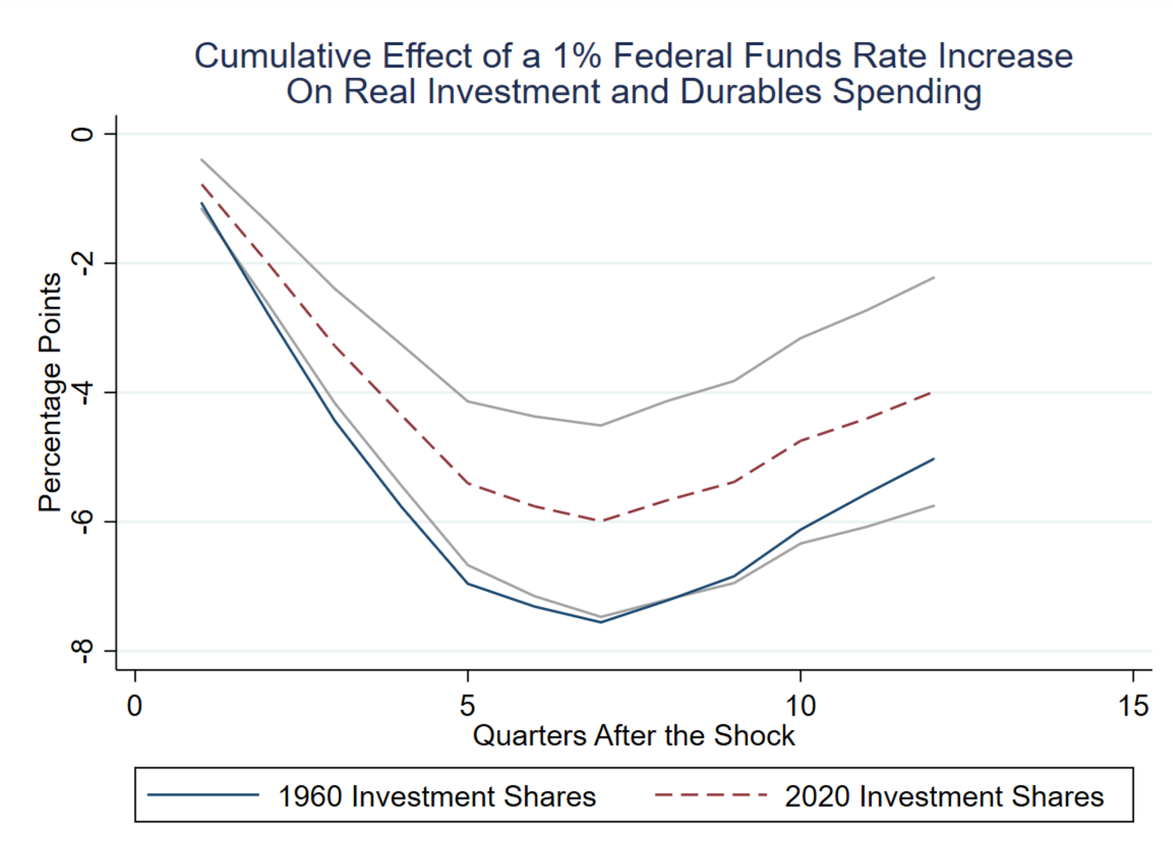

Notes: Two standard error bands are plotted, calculated as in Romer and Romer (2004). See text for estimation details.

changes to the transmission of investment to hand-to-mouth consumption, $\frac{\chi dlc_t^I}{1-\chi dlc_t^c}$, focusing on the numerator: the domestic labor content of investment dlc_t^I.

We calculate the domestic labor content for each final use component i (equipment, NDS, etc.) using the BEA Input-Output tables, performing a procedure similar to that of Hubmer (2020). The domestic labor content can be computed as:

$$dlc_t^i = \sum_k \omega_{ikt} \theta_{kt}^L,$$

where ω_{ikt} is the quantity of gross industry output from industry k demanded from a dollar of purchases of final demand component i, and θ_k^L is industry's k ratio of wage payments to gross output. ω_{ikt} is computed taking into account the full input-output structure of production, as well as the use of imported intermediates. The full derivation of ω_{ikt} is in

Appendix 3.7.2. Using a similar formula, we can compute the domestic value added share of final expenditure i as

$$v_t^i = \sum_k \omega_{ikt}\theta_{kt}^v,$$

where θ_{kt}^v is the value added share of gross output in industry k. To prelude the Cobb-Douglas production functions we will impose shortly, define the labor share of domestic value added $1-\alpha^i$ of final demand component i to be:

$$(1-\alpha_t^i) = dlc_t^i/v_t^i.$$

Denote the import share for final demand component i to be m^i. With simple rearranging, and noting that the import share $m^i = 1-v^i$, we can rewrite the domestic labor content as:

$$dlc_t^i = (1-\alpha_t^i)(1-m_t^i).$$

As demonstrated in Hubmer (2020), these empirical shares can be derived as an equilibrium outcome in an economy where industries produce with a CES production technology and take the output of other industries as intermediates. For our purposes, Cobb-Douglas production suffices.[14]

The top panel of Figure 2.4 plots the domestic labor content for investment (using the broad definition that includes durable goods and residential investment) and consumption (nondurables and services, henceforth NDS) from 1963-2014. Recall that the interpretation of this value is for every dollar spent on final investment or final consumption, how many cents are paid in domestic payroll. At the beginning of this period, every dollar of final expenditure on investment generated 59 cents of domestic payroll. By the end of the sample, the domestic labor content was only 46 cents on the dollar, representing a 22% decline. Over the same period, the domestic labor content of NDS fell only from 52 cents to 48-49 cents.

The bottom panel of Figure 2.4 performs a similar calculation for each component of investment separately. This plot shows that the decline in the DLC of investment is entirely due to the decline of the DLC for equipment and durable goods. Since equipment and durable goods account for nearly half of all gross nominal investment expenditure, the decline of the DLC for equipment and durables brings the investment-wide average down.

In Figure 2.5, we plot the decomposition of the DLC into the labor share of domestic production and the share of domestic expenditure: $dlc^i = (1-\alpha^i)(1-m^i)$. Figure 2.5(a) shows that the decline in the domestic expenditure share is driving the decline in the DLC for investment goods, and the domestic share of expenditure for NDS has changed very little since 1963. Figure 2.5(b) shows that the labor share in both investment and NDS was flat to slightly increasing from 1963 to around 2001, after which point the labor share declines in both components, though more dramatically in investment.

[14]See Jones (2011) for a proof that in a competitive input-output economy where the producing industries operate with a Cobb-Douglas production technology in capital, labor, and intermediates, then production aggregates into an aggregate Cobb-Douglas production, where the factor proportions are a function of the industry-specific factor proportions.

Figure 2.4: The Domestic Labor Content of Non-Durables and Services and Components of Investment

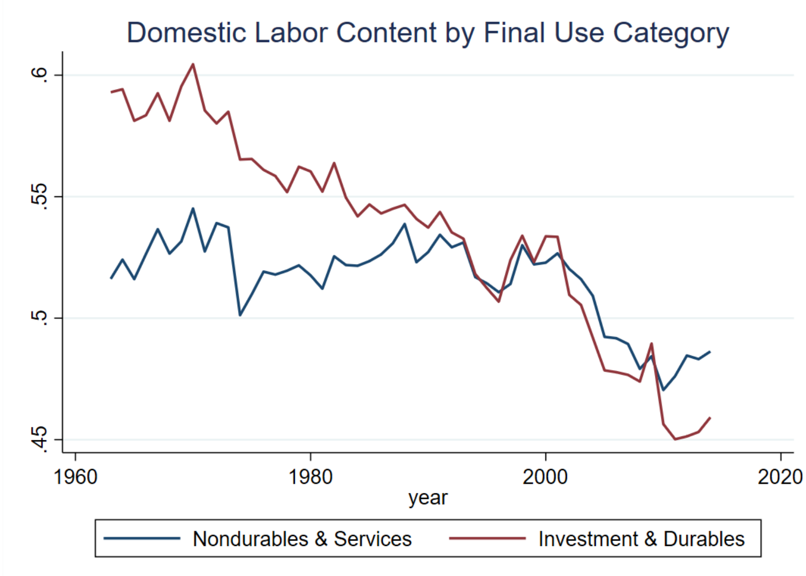

(a) Domestic Labor Content for NDS and Investment & Durables

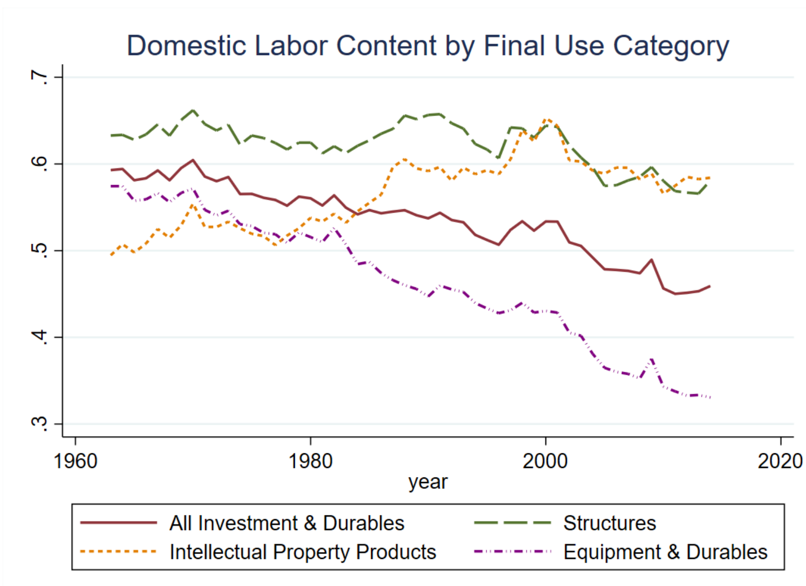

(b) Domestic Labor Content by Investment Category

Notes: The top panel shows that the domestic labor content of non-durables and services has declined by 4 percentage points, from 0.53 to 0.49, while the domestic labor content of investment and durables has fallen by 13 percentage points, from 0.59 to 0.46, a 22% decline. The bottom panel plots the domestic labor content of each component of investment, revealing that the decline is driven by the falling domestic labor content of equipment and durables, which falls by over 40% from 1963-2014.

Figure 2.5: Breakdown of Domestic Labor Content into "Domestic Share of Expenditure" and "Labor Share of Domestic Production" by Sector, [1963-2014]

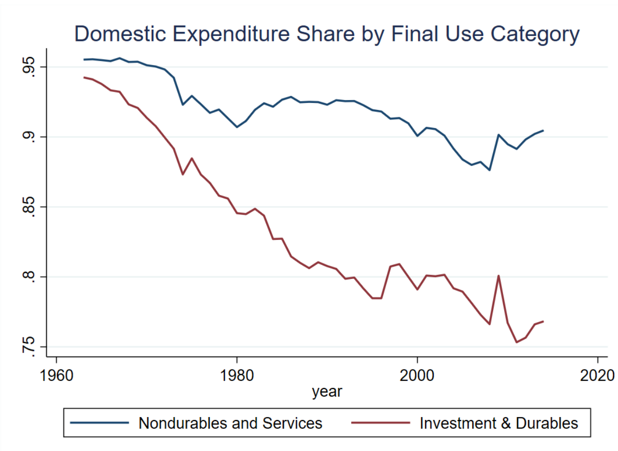

(a) Domestic Expenditure Share

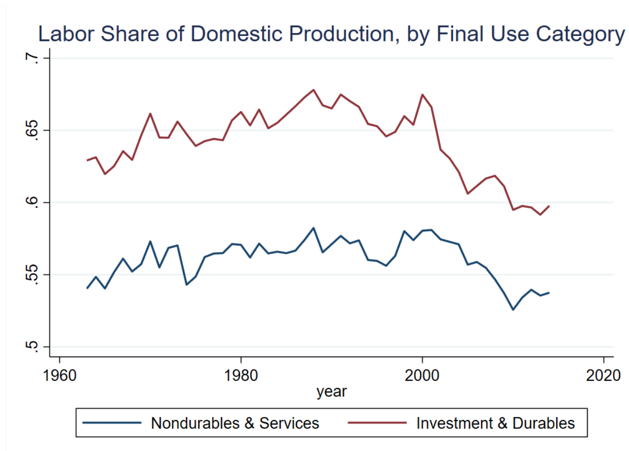

(b) Labor Share of Domestic Production

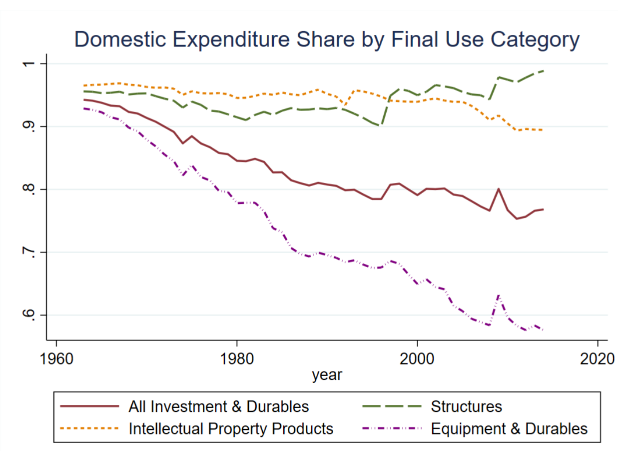

(c) Domestic Expenditure Share, Detailed

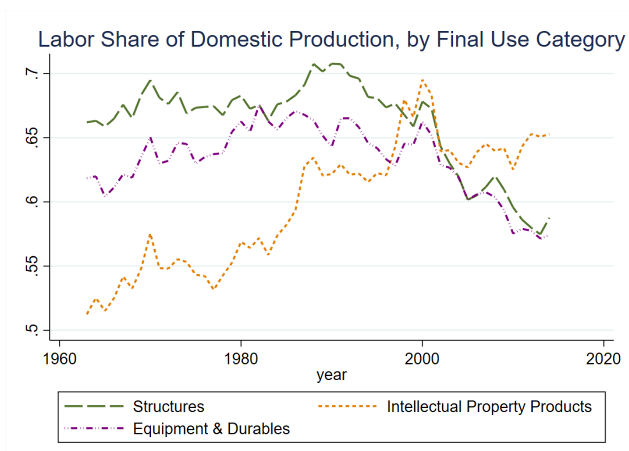

(d) Labor Share of Domestic Production, Detailed

Notes: These figures demonstrate that the majority of the decline in the domestic labor content for both investment and consumption has resulted from a rising share of imports, particularly for the production of investment goods.

Looking into the components of investment, Figure 2.5(c) shows that all of the decline in the domestic expenditure share is accounted for by equipment and durable goods. Structures and intellectual property products show almost no increase in use of imports. Finally, figure 2.5(d) shows the domestic labor share of each component. The labor share in the domestic IPP production has been rising over time, while it has fallen in both structures and equipment and durables.

The following section derives a small-open-economy, three-sector, two-agent new Keynesian model calibrated to match these observed declines in the domestic labor content. We will use this model to evaluate how the responsiveness of labor income and consumption to real interest rate shocks changes.

2.4 General Equilibrium Response to Monetary Policy Shocks

This section describes a simple small open economy new Keynesian model with two agents. We log-linearize the model around its non-stochastic steady state and study the perfect-foresight response of the economy to real interest rate shocks to demonstrate how a decline in the domestic labor content of investment reduces the effects of monetary policy on labor incomes and consumption in general equilibrium.[15] We follow Auclert, Rognlie, and Straub (2020) in placing nominal rigidities in wage setting rather than in prices, as this allows simple calculations of relative prices of goods produced in different sectors and abstracts from issues regarding the cyclicality of monopoly profits.

We then use this general equilibrium model to quantify the effects of three structural changes on monetary policy: (i) a falling labor share of domestic value added, (ii) the increased import content of investment goods, and (iii) the composition shift towards less responsive components of investment. The first two changes will take the form of changing parameters in the production of investment. The composition change will be modeled as a change in investment adjustment costs.

2.4.1 Households

There are two types of households, as described in Section 2.2. All households $i \in [0, 1]$ have an objective function

$$E_t\left[\sum_{\tau=0}^{\infty} \beta^{t+\tau}\left(\frac{C_{i,t+\tau}^{1-\sigma}}{1-\sigma} - \varphi\frac{N_{i,t+\tau}^{1+\eta}}{1+\eta}\right)\right],$$

where $N_{i,t} = \int_0^1 N_{i,k,t}dk$ are hours (or workers) supplied in k tasks from household i. However, only Ricardian households, of which there are measure $1-\chi$, can borrow and save in several assets; $B_{i,t}$ denotes the quantity of riskless, home bonds in zero net supply held by (Ricardian) household i at time t, and similarly $A_{i,t}$ denotes the quantity of foreign bonds denominated in foreign currency, each of which matures at $t+1$. S_t is the nominal exchange rate defined as home/foreign currency. The nominal, riskless return on the home bond is R_t (recycling

[15] See Auclert, Rognlie, and Straub (2020) and McKay, Nakamura, and Steinsson (2016) for examples of this approach; since the linearized model features multiple equilibria given a real interest rate path, uniqueness is achieved by requiring that the economy returns to steady state at some point in the distant future. Strictly speaking it is not necessary to log linearize the model first, but when solving the nonlinear version of the model the results for the impulse response functions are nearly unchanged.

notation from Section 2.2) and on the foreign bond R_t^f. Households can also acquire shares $\nu_{i,t}$ at price Q_t in a representative firm which accumulates capital, K_t, and pays out profits as dividends each period, D_t; we normalize total shares $\int_0^1 \nu_{i,t} di = 1$. Maximization is thus subject to the budget constraint:

$$\begin{aligned} C_{i,t} + \frac{B_{i,t}}{P_t^c} + \frac{S_t A_{i,t}}{P_t^c} + \frac{Q_t \nu_{i,t}}{P_t^c} &= \frac{S_t R_{t-1}^f \Phi_{t-1} A_{i,t-1}}{P_t^c} + \frac{R_{t-1} B_{i,t-1}}{P_t^c} + \frac{(Q_t + D_t)\nu_{i,t-1}}{P_t^c} \\ &+ \int_0^1 W_{k,t} N_{i,k,t} dk, \end{aligned}$$

where the price of the final consumption good P_t^c is implicitly used as the numeraire for the real wages $W_{k,t}$. The variable Φ_t is the "premium" which foreign assets pay over home assets, and depends on aggregate borrowing: letting P_t denote the price of exports and $a_t^f \equiv \int_0^1 S_t A_{i,t}/P_t di$, with

$$\Phi_t = \exp\left(-\phi_a (a^f{}_t - a_{ss})\right),$$

where the parameter a_{ss} determines the steady-state of a_t; in practice this is zero and $\phi_a > 0$ is calibrated to be small so that the risk premium term is small in equilibrium, as this term is included only to solve some well-known technical issues that arise in linearized, small open economy models (Schmitt-Grohé and Uribe, 2003).

The representative capital firm maximizes the present value of dividends subject to convex investment adjustment costs. Letting $\Lambda_t \equiv \Pi_{j=0}^{t-1} R_j^{-1}$, P_t^k be the nominal price of investment, and $\hat{R}_t^k$ be the nominal rental rate, the firm chooses a path for investment I_t to maximize:

$$E_t\left[\sum_{\tau=1}^{\infty} \Lambda_{t+\tau} D_{t+\tau}\right] = E_t\left[\sum_{\tau=1}^{\infty} \Lambda_{t+\tau}\left(\hat{R}_{t+\tau}^k K_{t+\tau-1} - \frac{\phi}{2}\left(\frac{I_{t+\tau}}{I_{t+\tau-1}} - \delta\right)^2 P_{t+\tau}^k I_{t+\tau-1} - P_{t+\tau}^k I_{t+\tau}\right)\right]$$

subject to the capital acccumulation constraint

$$K_t = I_t + (1-\delta) K_{t-1},$$

where ϕ governs the level of investment adjustment costs. While Ricardians have FOCs for consumption, bond holdings, and shares in the capital goods firm as shown above, the measure χ of hand-to-mouth (or "Keynesian") agents only have the condition that they consume all available labor income:

$$C_{i,t} = \int_0^1 W_{k,t} N_{i,k,t} dk.$$

Since all Keynesian agents obtain the same labor income in equilibrium, we write their consumption $C_{i,t} \equiv C_{k,t}$.

2.4.2 Union Wage Setting

Unions for each task k set nominal wages each period $\hat{W}_{k,t}$ subject to convex adjustment costs and downward-sloping demand from a labor packer who bundles labor in tasks k into aggregate labor N_t, and sells at nominal wage $\hat{W}_t$ to all final goods firms. Unions call upon their members to supply equal amounts of labor to meet demand, so $N_{i,k,t} = N_{k,t}$. Further, since each union faces the same problem, they always set the same wage and face the same labor demand, so

$$\begin{aligned} N_{k,t} &= N_t \\ \hat{W}_{k,t} &= \hat{W}_t. \end{aligned}$$

Unions set wages to maximize average expected utility of their members, putting equal weight on each household. This yields the following nonlinear wage Phillips curve: denoting gross nominal wage inflation as $\pi_t^w \equiv \hat{W}_t/\hat{W}_{t-1}$,[16]

$$\pi_t^w(\pi_t^w - 1) = \frac{\epsilon}{\psi}\left(\varphi N_t^{1+\eta} - \frac{\epsilon - 1}{\epsilon} W_t N_t \left((1-\chi)C_{r,t}^{-\sigma} + \chi(N_t W_t)^{-\sigma}\right)\right) + \beta \pi_{t+1}^w(\pi_{t+1}^w - 1),$$

where ϵ is the elasticity of substitution across tasks and ψ parameterizes the costliness of changing wages. In practice, we choose φ to normalize steady state labor supply $N = 1$ and choose ψ and ϵ to imply a slope of the linearized wage Phillips curve equal to 0.1, as in Auclert et al. (2018). Since labor income is the same across types, aggregate consumption demand takes the form given in equation (2.1) above.

2.4.3 Investment, Consumption and Export Goods Producers

All final goods production is Cobb-Douglas: letting K_t^j, L_t^j and M_t^j be the capital, labor and imported intermediates used in sector j at time t,

$$\begin{aligned} I_t &= \left(Z_t (K_t^i)^{\alpha_i} (N_t^i)^{1-\alpha_i}\right)^{1-m_i} (M_t^i)^{m_i} \\ C_t &= \left(Z_t (K_t^c)^{\alpha_c} (N_t^c)^{1-\alpha_c}\right)^{1-m_c} (M_t^c)^{m_c} \\ X_t &= Z_t (K_t^x)^{\alpha_x} (N_t^x)^{1-\alpha_x}, \end{aligned}$$

where we assume the export sector does not use imports; see the diagram in Figure 2.6. Each sector is perfectly competitive and sets prices flexibly, so that the labor share of value added and import shares of gross output in each sector j are constant and given by $1 - \alpha_j$ and m_j, respectively. These are calibrated to match the estimates in section 2.3.4.

We assume that imports are priced in the home currency and fixed at price P^m. We further assume that foreign demand for the home country's exports is a constant elasticity function of the price of home exports denominated in the foreign currency:

$$X_t = \left(\frac{P_t}{S_t}\right)^{-\tau} m_c^f Y^f,$$

where τ is the elasticity of foreign demand for home exports, and P_t is the price of exports in home currency. Dividing P_t by the exchange rate S_t converts the price of exports into foreign

[16]For a derivation, see Appendix C1 of Auclert et al. (2018).

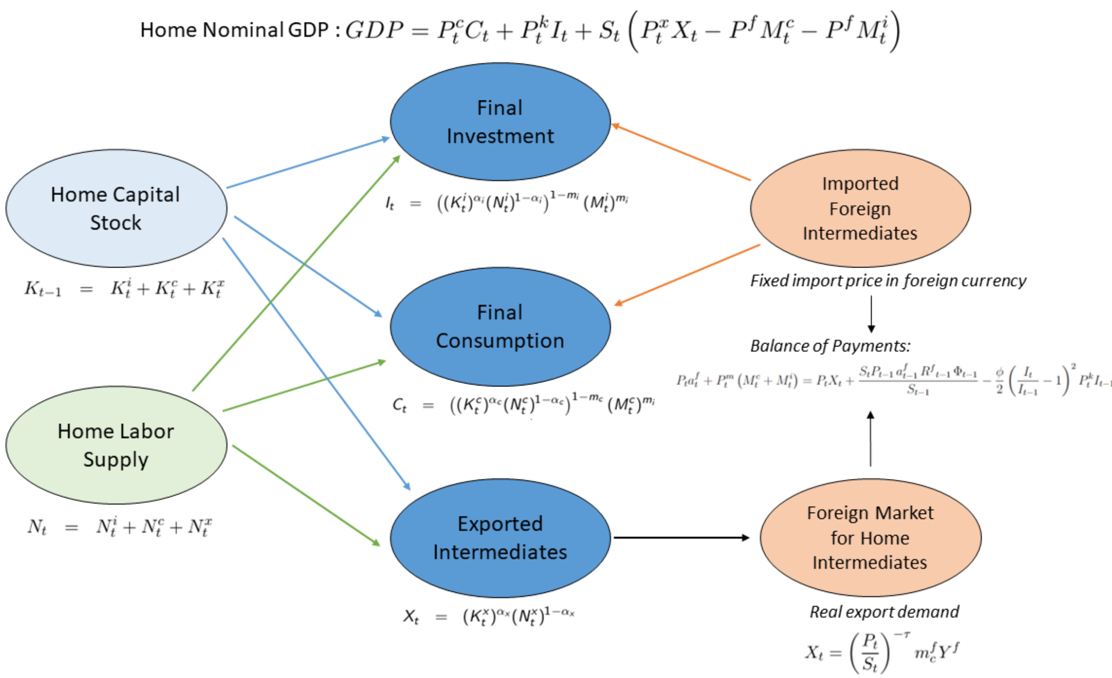

Figure 2.6: The Production Structure of Section 2.4's Small Open Economy Model

currency. Y_f is the level of foreign nominal income, and m_c^f is a scaling factor that affects how much export is demanded.[17]

All markets for labor, capital, and imports M_t clear:

$$\begin{aligned} N_t &= N_t^i + N_t^c + N_t^x \\ K_{t-1} &= K_t^i + K_t^c + K_t^x \\ M_t &= M_t^i + M_t^c. \end{aligned}$$

Finally, the balance of payments which equates income from home's foreign assets and exports with purchases of new assets and imports can be derived by aggregating across households' budget constraints:

$$P_t a_t^f + P_t^m \left(M_t^c + M_t^i \right) = P_t X_t + \frac{S_t P_{t-1} \, a_{t-1}^f \, R^f{}_{t-1} \, \Phi_{t-1}}{S_{t-1}} - \frac{\phi}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2 P_t^k I_{t-1},$$

and nominal GDP is defined as:

$$GDP = P_t^c C_t + P_t^k I_t + S_t \left(P_t^x X_t - P^f M_t^c - P^f M_t^i \right).$$

A Taylor rule closes the model and selects a zero-inflation steady state; however, going forward, we will consider shocks to the path of real interest rates as our policy experiment

[17] In the case of $\tau = 1$, demand is Cobb-Douglas, and m_c^f is the share of foreign income spent on the home country's exports.

of interest and so the specification is irrelevant. Table 2.2 summarizes the choices of the parameters used in the following exercises.

2.4.4 The Perfect-Foresight Response to a Transitory Real Interest Rate Shock

Log linearizing around a non-stochastic steady state yields a convenient, static expression decomposing any percentage deviation in hand-to-mouth consumption from steady state into a weighted average of three components: for any variable x_t, let x be the steady state of x_t and $\hat{x}_t \equiv \log x_t - \log x$,

$$\underbrace{\widehat{C_{k,t}}}_{\substack{\text{Hand-to-Mouth}\\ \text{Consumption}}} = \frac{N^x}{\Omega} \underbrace{\widehat{\left(\frac{P_t^x X_t}{P_t^c}\right)}}_{\text{Real Exports}} + \frac{N^i}{\Omega} \underbrace{\widehat{\left(\frac{P_t^k I_t}{P_t^c}\right)}}_{\text{Real Investment}} + \left(1 - \frac{N^x}{\Omega} - \frac{N^i}{\Omega}\right) \underbrace{\widehat{C_{r,t}}}_{\substack{\text{Ricardian}\\ \text{Consumption}}}, \quad (2.6)$$

where $\Omega \equiv 1 - \chi(1-\alpha_c)(1-m_c)$ and N^x and N^i are the steady state labor shares in the export and investment sectors, respectively.[18] The ratios N^i/Ω and N^x/Ω then function as weights on the percent deviations in exports and investment, with a complementary weight $(1 - N^i/\Omega - N^x/\Omega$ on Ricardian consumption. Equation (2.6) then clarifies how the response of hand-to-mouth consumption to a given path of real interest rates depends on both the interest sensitivity of investment and exports and the employment weights. Since investment in our calibration is the most responsive component of monetary policy, lowering the domestic labor content of investment weakens monetary policy's effect on hand-to-mouth consumption because less labor is used in the investment good sector, lowering N^i and thus putting less weight on investment in the right hand side of equation (2.6). The exercise we perform is a shock to the real interest rate of one hundred basis points (1%) at $t = 1$ that decays according to an AR(1) process with persistence ρ_m.

Calibration We set the share of hand-to-mouth agents to .3 to match the hand-to-mouth share in Kaplan et al. (2014) and the average MPC on nondurables and services of .2 as in Kaplan and Violante (2022). We set the intertemporal elasticity of substitution (IES) to 1/6 (so the inverse IES is 6), which is lower than is standard in the literature, but is closer to the estimate in Best et al. (2020) who find an IES of .1. We jointly calibrate the persistence in the real rate shock ρ_m and the investment adjustment costs ϕ to match the duration and amplitude of the empirical response of investment to a monetary policy shock in Figure 2.3, yielding $\rho_m = .9$, which is fairly standard, and $\phi = 14$. We set the sensititivity of export demand to the exchange rate τ to .1. While this is low relative to estimates of short-run trade elasticities, this value allows us to match our empirical estimates of changes in exports to Romer & Romer shocks. We set the import content of consumption goods $m_c = .05$.

[18] See Appendix 3.7.4 for a derivation.

Parameter	Value	Description	Notes
β	.99	Households' Discount Factor	Calibrated to steady state real annual interest rate of 4%
σ	6	Inverse Elasticity of Intertemporal Substitution	
η	1	Inverse Frisch Elasticity of Labor Supply	
φ	Varies	Scales level of disutility from labor	Chosen to normalize steady-state $N = 1$.
χ	0.3	Share of hand-to-mouth agents	Taken from Kaplan, Violante, and Weidner (2014).
ϕ_a	0.01	Responsiveness of the risk-premium to aggregate NFA	Chosen to be small as in Schmitt-Grohé and Uribe (2003).
a_{ss}	0	Steady-state NFA	
ϕ	14	Level of convex capital adjustment costs	Match empirical IRF of investment.
δ	0.03	Depreciation rate of capital	
τ	.1	Sensitivity of exports to exchange rates	Match empirical IRF amplitude of exports.
ϵ	10	Elasticity of substitution across tasks	Calibrated to a markup of 10% for unions
ψ	100	Level of convex wage adjustment costs	Sets the wage Phillips-curve's slope to 0.1 as Auclert et al. (2018).
α_j, m_j	Varies	Capital share and import share of sector j	Calibrated to estimates in Section 2.3.4.
$m_c^f Y^f / P^f$	Varies	Real exports in terms of foreign currency import price	Chosen to normalize steady-state terms of trade $P^m/P = 1$.

Table 2.2: Parameters for Section 2.4's Quarterly, Small Open Economy New Keynesian Model.

Results We simulate how consumption responds to the 100 basis point AR(1) shock to real interest rates before and after the structural change under three different calibrations. In the first calibration, we choose parameters for a "1960's economy," where the import share of investment goods m_i is .05 and the capital share of investment goods α_i is .35. Results are reported with a solid blue line. In the second calibration, with the dotted line, we show the results taking into account only the components of structural change affecting the domestic labor content: we increase the import content of investment goods to .25, and we increase their capital share to .4. In the third calibration using the dashed line, we keep the changes to the import content and capital shares, and in addition we raise the investment adjustment costs[19] from $\phi = 14$ to $\phi = 19.4$ to match the 20% dampening of investment in response to monetary policy shocks in our shift-share exercise in Figure 2.3.

Figure 2.7(a) plots the response of consumption in percent deviations from steady-state for these three different calibrations. We show both the aggregate consumption response and the response of hand-to-mouth agents, which has the identical reponse to shocks as aggregate labor income. Aggregate consumption rises on impact by about 1%, and hand-to-mouth consumption raises around 1.5% on impact. Not surprisingly, the effect of hand-to-mouth consumption is larger and differs more across calibrations. In the first calibration with $m_i = .05$ and $\alpha_i = .35$, the peak response of hand-to-mouth consumption is an increase of 2.76 percent, seven quarters into the shock. When the domestic labor content falls to 2010 levels in the second calibration with $m_i = .25$ and $\alpha_i = .4$, the response of hand-to-mouth consumption peaks at 2.28 percent, an approximate thirteen percent decline. Increasing the investment adjustment cost parameter ϕ in the third calibration lowers the peak response of hand-to-mouth consumption to 1.93 percent, a thirty percent decrease from the first calibration.

Figure 2.7(b) plots the response of each of the three right hand side components in equation (2.6): real Ricardian consumption, real investment, and real exports. The first panel shows the response of consumption by Ricardian agents. By construction, these impulse response functions are identical across calibrations, as the path of real Ricardian consumption is pinned down only by the path of real interest rates. The second panel shows the impulse response of real investment. The impulse response in the first two calibrations is nearly identical: changing the import content or the labor share of investment has very small effect on the path of desired investment, as the assumption of dollar invoicing of imports assures that the price of investment is mostly unchanged even as the import content of investment increases.[20] In the third calibration where adjustment costs increase, the response of investment is significantly muted and matches the 20% decline in amplitude that we found in our shift-share exercise in Figure 2.3. In the final panel is the response of real exports, which we will discuss shortly.

[19] Bloesch and Weber (2023) show that congestion in onboarding workers in software and research & development provides a microfoundation for convex adjustment costs in the production of intellectual property investments.

[20] In a previous version of the paper, we assumed that imports were priced in a foreign currency, so changes in exchange rates would have significant effects on the price of investment. This led monetary policy to have larger effects on investment. However, we now assume that imports are invoiced in dollars.

Figure 2.7: Response of Hand-to-Mouth Consumption and Determinants of Consumption

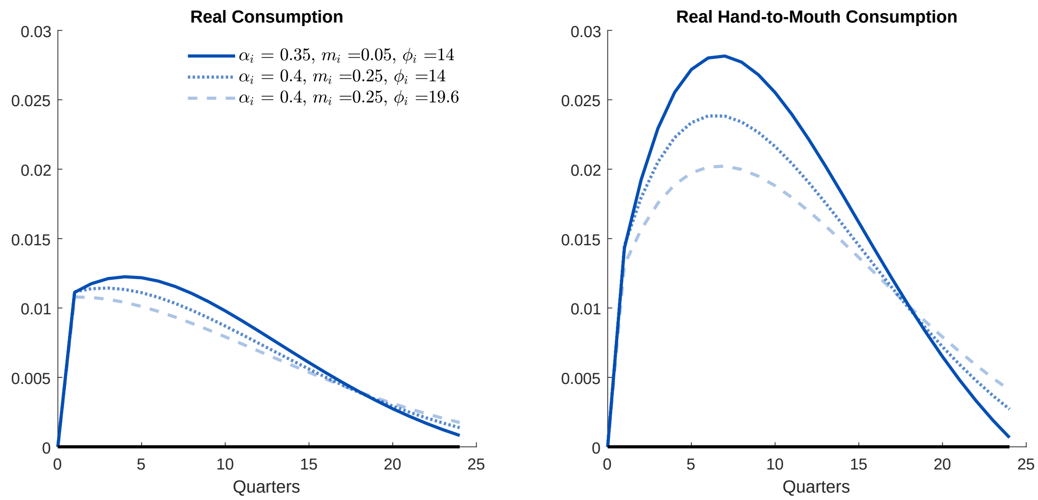

(a) Consumption Response

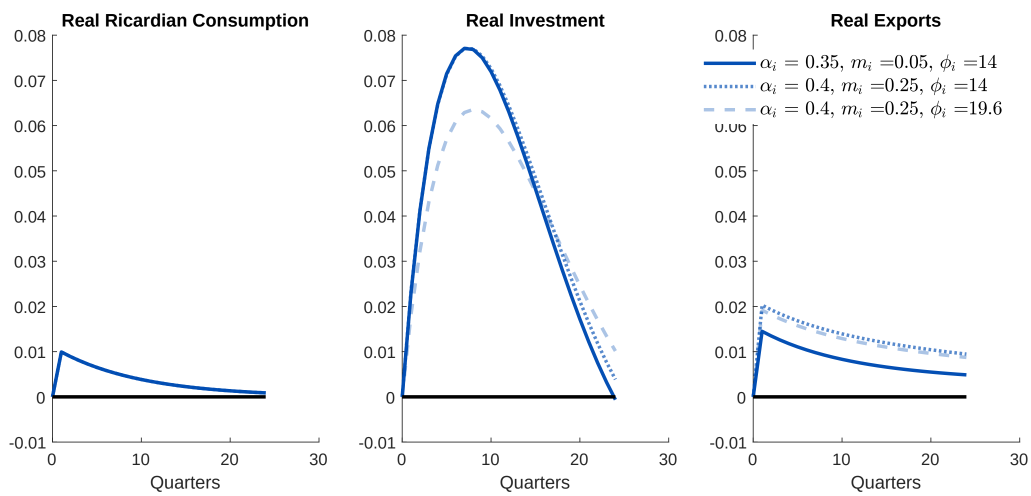

(b) Determinants of Hand-to-Mouth Consumption

Notes: Each figure displays the response of a particular variable to an expansionary 1% AR(1) shock to the real interest rate across three calibrations of Section 2.4's small open economy new Keynesian model; see text for details. Each line reflects a different calibration as indicated in the legend. The first, solid line is calibrated to be illustrative of a 1960s economy; the second dotted line changes the import share and capital share of domestic value added to its modern value, and finally a higher investment adjustment cost. The share of labor in the investment good sector N^i falls from 25 percent to 19 percent between the first and second calibration, while the share of labor in the export sector N^x grows from 5 percent to 11 percent.

How important is each component of structural change in understanding the total weakening of the effect of monetary policy on labor income and consumption? To estimate this, we compare the cumulative response of consumption and labor income over the first 12 quarters following the shock, but changing only one component of structural change at a time and keeping the other components the same as the 1960's calibration. Table 2.3 shows this decomposition. In the first row, we compare the cumulative 12 quarter response of consumption and labor income (summing up the percent devations from steady state) for the 1960's economy, and then only changing the capital share from .35 to .4. In the second row, we compare the cumulative response of the two outcome variables in the 1960's economy versus an economy that only changes the import content of investment goods from .05 to .25. The third row reports the same exercise, but only changing the adjustment cost parameter, keeping the other parameters at their 1960's values.

Table 2.3: Decomposition of Dampening of Effects of Real Rate Shocks

	Consumption		Labor Income	
Structural Change	Decline	Share of Total	Decline	Share of Total
Capital Share α_i: .35 to .4	.007	5%	.007	3%
Import Content m_i: .05 to .25	.074	48%	.122	48%
Adjustment Costs ϕ: 14 to 19.6	.093	61%	.154	61%
Sum of Individual Effects	.174	114%	.283	112%
Combined	.153	100%	.252	100%

Notes: This table decomposes the share of the dampening of the response of consumption and labor income to shocks to real interset rates that is accounted for by the three components of structural change studied here: (i) a fall in the labor share and rise in the capital share of domestic value added in investment, (ii) the rising import share of investment goods, and (iii) the shift towards Intellectual Propety Products, modeled by rising investment adjustment costs. For each row, we compute the cumulative response over 12 quarters of percent deviations of labor income and consumption from steady state for the 1960's calibration. Then holding all parameters at the levels of the 1960's calibration, we change one parameter reflecting structural change at a time (α_i, m_i, and ϕ, respectively), and we estimate how much smaller the cumulative 12 quarter response is for each change. The dampening effect from changing all three parameters one at a time is less than the sum of dampening from changing parameters individually.

The results in Table 2.3 show that, by far, the increase in the import share and the shift towards more intellectual property products investment (higher adjustment costs)[21] accounts for the largest shares of the declining responsiveness of labor income and consumption. Individually, each of these two pieces of structural change accounts for around half of the total dampening, whereas the increase in the capital share of investment goods accounts for a very small share. The effects of each structural change individually add up to more than the total because the interaction effects are negative in a model with only one capital

[21] In a previous version of this paper, we captured the dampening of investment due to the shift toward less-responsive intellectual property products by increasing the depreciation rate of capital.

Figure 2.8: Response of Net Exports and Net Foreign Assets

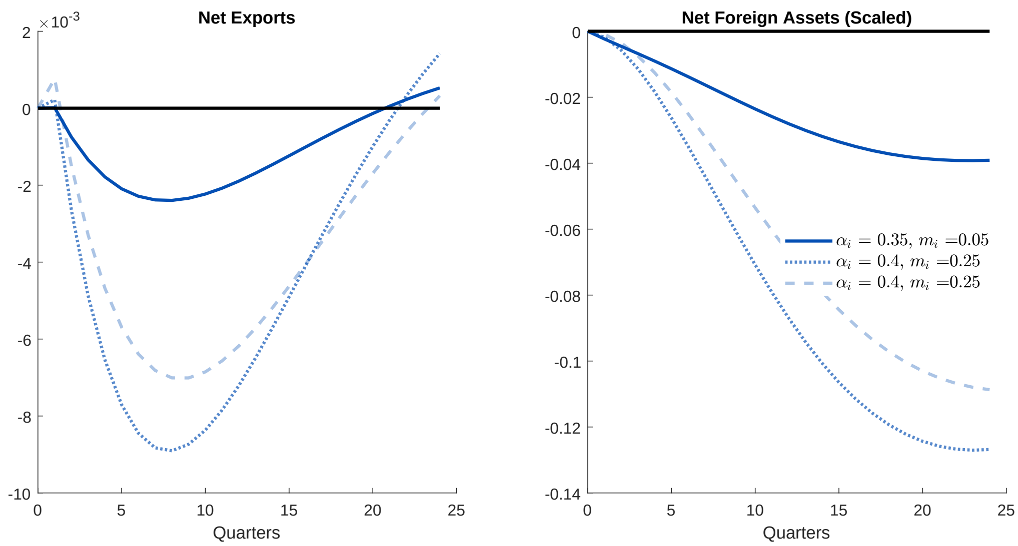

Notes: Each figure displays the response of a particular variable to an expansionary (negative) 1% AR(1) shock to the real interest rate in different calibrations of Section 2.4's small open economy new Keynesian model; see text for details. The calibrations are the same as in Figure 2.7.

good. For example, higher adjustment costs lower the response of labor income. However, if investment has lower domestic labor content, then a weakening of the response of investment has a smaller dampening effect on labor income in absolute terms, but a similar effect in percent terms.

Figure 2.8 reports responses for net exports and the level of net foreign assets in response to the real interest rate shock. In each calibration in the first panel, the response of net exports is negative, consistent with findings in Cloyne et al. (2020), Kim (2001), and Degasperi et al. (2023) that the US trade balance deteriorates in response to an expansionary US monetary policy shock. This is because the short run demand for imports, primarily as inputs in investment goods, swamps the effect on demand on exports from a weakened exchange rate. As the import content of investment goods grows, the demand for imports gets bigger in response to real rate shocks, amplifying the decline in net exports. As shown in the second panel, these declines in next exports are financed by a deterioration in net foreign assets. As the import share of investment goods increases, households increasingly respond to real interest rate shocks by borrowing more from the rest of the world.

Our model's result that net exports falls in response to an expansionary monetary policy shock is inconsistent with the standard Mundell-Fleming model but is consistent with recent empirical evidence, which we add to here. In Figure 2.9, we confirm the empirical results in Cloyne et al. (2020), Kim (2001), and Degasperi et al. (2023) that net exports responds positively (negatively) to a contractionary (expansionary shock). Figure 2.9(a) shows the percent deviation of exports to a contractionary monetary policy shock, an Figure 2.9(b) shows the percent deviation of imports. Imports respond negatively in the early periods and recover after 10 quarters, while exports respond negatively with a substantial lag. This indicates that the response of net exports is positive in the short run in response to a contractionary shock but negative in the long run. In the other direction, an expansionary shock worsens the trade balance in the short run, and the trade balance responds postively over longer horizons, as captured in our model. In total, these results support our model's prediction that the effect of investment on imports is larger than exchange rate effects on exports. In addition, the small magnitudes on the trade variables mean that investment is more important than trade in determining the response of domestic labor income to monetary policy shocks.

Figure 2.9: Impulse Reponse of Exports and Imports to Romer & Romer Shocks

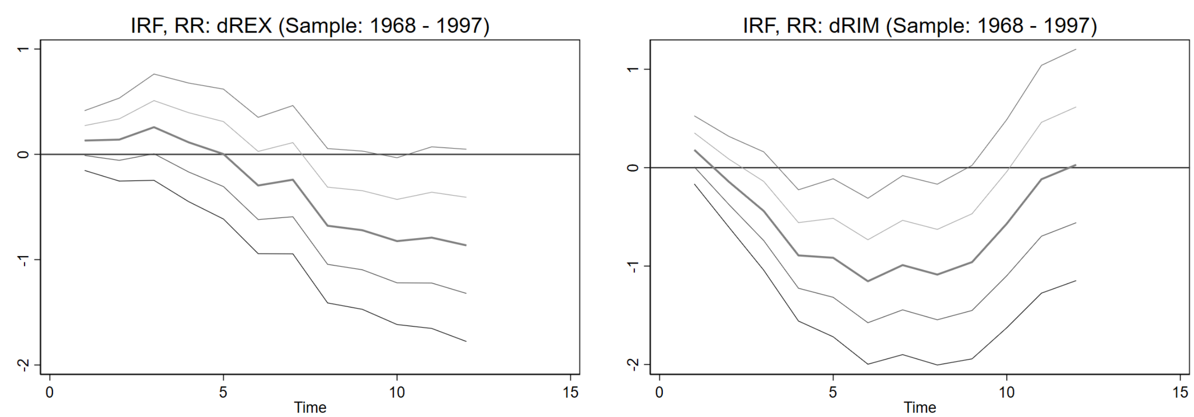

(a) Exports (b) Imports

Notes: These figures plot the percent response of exports and imports to a contractionary shock to the federal funds rate. Imports respond negatively in the early periods and recover after 10 quarters, while exports respond negatively with a substantial lag. At ten quarters, the effect on net exports is negative. This indicates that in response to a rise in real interest rates, net exports rises. In the other direction, an expansionary shock worsens the trade balance in the short run, and the trade balance responds postively over longer horizons, consistent with Cloyne et al. (2020)'s Figure 13, Kim (2001), and Degasperi et al. (2023).

2.5 Conclusion

This paper has documented that secular change in both the production and composition of investment goods has weakened private investment's role in the transmission of monetary policy to labor earnings and consumption. Due to a declining domestic labor content of investment, investment fluctuations no longer amplify the effects of monetary policy shocks on consumption as they once did. Moreover, the investment response to monetary policy is likely weaker as well due to an increasing share of "intangible" investment in final demand.

These results may have important implications for optimal monetary policy. When combating high inflation, central banks in countries which have experienced secular changes in investment like those documented here for the United States may find that it takes a larger increase in interest rates now to reduce aggregate labor income than in previous decades. In times of low demand, if the quantity of conventional monetary policy needed to stimulate output is greater in modern economies, then policy makers should expect to find themselves constrained by the effective lower bound (ELB) on interest rates more frequently, all else equal. Further study of the implications in a more realistic, quantitative model and for dynamics surrounding ELB episodes is on our agenda, and may shed additional light on the observed tendency for open, increasingly "service based" economies to find themselves constrained by the ELB in recent decades.

Chapter 3

How Powerful is Unannounced, Sterilized Foreign Exchange Intervention?[1]

3.1 Introduction

Most developing and many advanced economies intervene in foreign exchange markets to manage volatility and unwanted exchange rate movements. Many central banks both sterilize their interventions and conduct them secretly, meaning without any announcement beforehand or acknowledgement after the fact.[2] Using recently-collected data on the Bank of England's daily operations during the Bretton Woods era, we conduct a new analysis of such unannounced, never-acknowledged sterilized intervention and find statistically significant effects on the level of the exchange rate.

Our analysis informs an ongoing debate regarding whether sterilized foreign exchange intervention could meaningfully impact the exchange rate. In this context, "sterilized intervention" means market operations undertaken to influence the exchange rate while leaving the monetary base unchanged. This usually takes the form of a paired transaction in which the central bank buys (sells) domestic currency in foreign exchange markets while simultaneously selling (buying) domestic currency bonds. Policymakers generally believe that *unsterilized* intervention could work through impacting relative interest rates, and that public, *sterilized* intervention may work through a signaling channel. With sterilized, secret intervention these channels will be absent or muted, and policymakers exhibit less agreement on whether intervention may be effective in this context.[3]

Despite this, many policymakers intervene in secret, including e.g. developing economies in Asia (Fratzscher et al., 2019). Additionally Chamon et al. (2019) report that despite

[1] I thank Alain Naef for allowing me to use our joint work in this chapter (Naef and Weber, 2023).

[2] From a recent survey of 22 "Emerging Market Economy" central banks, 17 reported "Never" or "Rarely" pre-announcing their intervention; when asked if they published data after the fact, 7 reported never publishing data at all and only 9 reported publishing data at a greater-than-monthly frequency (Mohanty and Berger, 2013); separately, most reported routinely sterilizing their interventions.

[3] When surveyed, most central bankers agree that the signaling channel of sterilized foreign exchange intervention is "effective" while exhibiting less agreement on other channels (Mohanty and Berger, 2013).

officially committing to floating regimes and inflation targeting, many countries in Latin America use foreign exchange intervention as part of their policy mix. The IMF has also included foreign exchange intervention as part of its new institutional policy framework, ensuring that the debate on efficacy will remain relevant.

The academic literature on intervention surveyed in e.g. Sarno and Taylor (2001) and Neely (2008) focuses disproportionately on the few central banks that intervene publicly and publish their intervention data. Even when circumstances have allowed for the study of other central banks, it is not always clear whether the operations were really secret or promptly sterilized.[4] This renders disentangling the channels at play challenging. Finally, even with access to quality data, all studies of the effects of intervention must grapple with the issue of endogeneity, as intervention is far from randomly assigned. See Fratzscher et al. (2019) for a recent discussion of both issues. Recent work has used high-frequency, intra-day data to overcome some of these issues and identifies the effects of intervention with event study or regression discontinuity approaches (Fatum and King, 2005; Menkhoff, 2010; Kuersteiner et al., 2018). In our setting, such data are not available; we therefore rely on alternative methods and daily data.

By studying the Bank of England, we contribute a case study of a central bank that intervened frequently (on over 80% of trading days in sample), sterilized immediately, and operated with a high degree of secrecy. Relative to previous studies on the Bank of England (Naef, 2019, 2021) reporting mixed evidence for effectiveness from correlations and event-studies, this paper presents causal point estimates of the effect of sterilized foreign exchange intervention on the exchange rate. We leverage the institutional setting of the Bretton Woods era to motivate two approaches to identification: an instrumental variables (IV) approach which forms the benchmark analysis in the paper, and a "Policy Rule" approach presented in Appendix 3.8.3. The latter proceeds by estimating a rule for determining the quantity of intervention normally conducted by the Bank of England, and treating deviation from it as a shock to intervention on days when the Bank of England was closed for a UK-specific Banking Holiday. The point estimates of these two independent approaches yield reassuringly similar results.

To understand how both approaches deal with endogeneity, we appeal to a reduced-form model of the portfolio balance channel, as this theory traditionally garnered the most attention as an explanation for the efficacy of sterilized intervention absent information effects.[5] The model informs our regression specifications and disciplines our discussion of identification issues. We note that ordinary least squares (OLS) estimates of the effects of intervention

[4]Examples include Fratzscher et al. (2019), which includes 23 non-public datasets on daily intervention and likely includes many secret interventions (the authors found news coverage for only half of their intervention episodes) and which the contributing central banks affirmed were sterilized. Some recent high-frequency studies on secret, intraday data on intervention in the market for the Czech koruna include Dominguez et al. (2013), who look at sales of reserves, and Scalia (2008), who studies sterilized interventions without observing quantities, making inference instead based on intervention dates.

[5]We avoid framing our results as showing a portfolio balance channel in recognition of e.g. a market microstructure channel or other channels which might operate independently of the signaling channel, and which our empirical analysis will not rule out.

will be biased if central banks "lean against the wind." For example, if the Bank of England attempted to strengthen the pound whenever it was weakening due to some fundamental shock, this would bias OLS estimates towards finding intervention unproductive (or even counterproductive). Both approaches yield results consistent with such a bias, despite relying on completely different identification assumptions.

Our IV approach takes advantage of the Bank of England's explicit exchange rate target during Bretton Woods, and uses the lagged distance of the exchange rate from target as an instrument for intervention. The motivation for this instrument is that the dealers working for the Bank of England, who were charged with intervening, may have been quicker to act if the exchange rate closed further from target the day before. The identifying assumption is that the level of the exchange rate the day before does not impact the growth rate of the exchange rate, except through the (secret) actions of the central bank. We motivate this by noting that if our assumption did *not* hold, and the level of the exchange rate was useful for forecasting its growth rate, then market participants consistently "left money on the table" in a large and liquid market. Point estimates obtained from this approach are precisely estimated and robust to variations in specification of the instrument, set of controls, and time period of estimation. This identification assumption may not hold if there is mean reversion in fundamental shocks to the exchange rate, and the hope is that such mean reversion is small.[6]

We are not the first to estimate the effects of sterilized foreign exchange intervention. Rieth, Menkhoff, and Stöhr (2019) use a structural vector autoregressive model with external instruments to identify the effects of the Bank of Japan's pre-announced interventions. We use data on interventions that were never made public and take different approaches to identification. Several recent papers present estimates using readily-available, low-frequency proxies for intervention.[7] A large event study literature evaluates the effect of intervention on a number of explicit success criteria, such as the direction of the exchange rate, reporting mixed results. For example, Fratzscher et al. (2019) finds that intervention is effective in up to 80% of cases, while Bordo et al. (2012) argue that the success rate for US interventions was historically no better than random.

This paper contributes to this literature by establishing the presence of nontrivial effects of sterilized intervention absent a significant signaling channel. This finding at least partially rationalizes the choice of many central banks to conduct intervention secretly.

[6] Appendix 3.8.4 presents a simple theoretical model where the exclusion restriction nearly holds if unobserved, fundamental shocks to the level of the exchange rate are almost a random walk; Appendix 3.8.5 explores this issue quantitatively, establishing through simulation that small amounts of mean reversion are unlikely to be driving our results.

[7] See e.g. Blanchard, Adler, and Filho (2015), who use changes in foreign exchange reserves observed at a quarterly frequency; Daude, Yeyati, and Nagengast (2016) who use changes in the ratio of reserves to M2 at a monthly frequency; and Adler, Lisack, and Mano (2019), who use changes in the central bank balance sheet at a monthly frequency.

3.2 Historical and Institutional Context

The Bretton Woods system of pegged exchange rates lasted from 1944 to 1971, but our analysis begins in January of 1952 when the Bank of England reopened the foreign exchange market. This setting has three important features for our analysis: the Bank of England was charged with managing a clear exchange rate target; the sterilization of the foreign exchange operations we study was automatic; and interventions were secret.

In the system, the dollar was fixed to gold at $35 an ounce and all other currencies were pegged to the dollar with a band of plus or minus 2%. The pound was fixed at the official price of $2.80 per pound between 1949 and 1967, and at $2.40 per pound between 1967 and the collapse of the system when President Nixon closed the "gold window" on August 15th, 1971. We use the pound/dollar exchange rate, and Figure 3.1 plots the exchange rate and target over time in these units.

Another key feature of our setting is that the interventions we study were systematically sterilized, as they were conducted through the Exchange Equalisation Account (EEA); Howson (1980) and more recently Allen (2019) establish that sterilization was a built-in feature of the EEA. By design, any operation of the EEA had a counterparty in UK Treasury bills, leading to automatic sterilization (note that the EEA is a government body technically independent from the Bank of England, which only manages the EEA). This makes us confident that we are indeed estimating effects of sterilized intervention and not simply picking up effects resulting from changes in the money supply.

Our setting's last notable feature is that all interventions studied were conducted in secret, meaning that the Bank of England did not communicate their daily intervention operations or make public the data at any point.[8] A natural question is whether these operations were secret in practice, as the Bank of England's counterparties (a small number of private banks) knew when it engaged in foreign exchange transactions. We argue that intervention was likely secret for three reasons. First, not all of the central bank's foreign exchange transactions were associated with intervention, making it difficult in practice for counterparties to determine the bank's intentions at short horizons. The Bank of England, nationalized in 1946, retained many private customers (including other central banks) and frequently engaged in "customer operations" in addition to the intervention we study. These operations were substantial: on approximately 40% of our observed intervention days, the bank was also engaged in customer operations. This would have made it difficult to disentangle the Bank's intention from its observed purchases and sales.

Second, the Bank of England outsourced some of its intervention to other central banks, reflecting the global nature of the foreign exchange market. Although most intervention was conducted in the London spot market, the Bank could request that the Federal Reserve, for example, intervene in the New York market and frequently did so (on just over 15% of all trading days).

Finally, anecdotal evidence confirms that secrecy was a goal from an early date, and that the Bank of England's dealers believed themselves to have been successful in this goal. In

[8] Public announcements accompanied interventions beginning in the 1980s.

1936, the senior official at the Bank of England in charge of foreign exchange matters, Harry Arthur Siepmann, noted that:

> Experience has enabled some progress to be made in concealing or advertising the presence of the 'control', and this has led to masked intervention being resorted to more frequently and successfully. It is sometimes surprising to find how wide of the mark are the press reports of the EEA activity, as when on 6 April we bought nearly Fr.200 million but were reported by the financial news next morning as having "retired from the market soon after opening"... It is equally clear that, whatever precautions are taken, the presence of the 'control' cannot, as a general rule, escape observation, though guesses may be good or bad about the extent of its daily operations.[9]

Thus while markets understood the Bank intervened, day-to-day variations were difficult to discern. By the 1950s, "masked" intervention was the rule, and secrecy an established policy goal. In a 1956 communication with the New York Fed, Roy Bridge, head dealer at the Bank of England, explained his strategy: 'I shall ask you to go into the market after lunch. ... Don't go before lunch. I thought it wise to change tactics a bit. It is a good thing.'[10] In short, the Bank took pains to conceal its intervention, and believed these efforts to have been successful.[11]

3.3 New Archival Data

We analyze a new, daily time series on the foreign exchange operations of the Bank of England taken from confidential reports sent from the Bank of England to the Treasury, which discriminates between "customer" operations and intervention meant to influence exchange rates (Naef, 2019, 2021). Figure 3.2 presents the daily series on intervention, deflated by UK M0. While the majority of this intervention was conducted directly by the Bank of England in the London pound/dollar spot market, the measure also includes intervention conducted by other central banks on the Bank of England's behalf in offshore markets. Note the sheer frequency of intervention: on approximately 83.3% of the trading days the Bank of England intervenes in the spot market. For comparison, Fratzscher et al. (2019) find in a sample of 33 central banks observed from 1995-2011 an average frequency of intervention of 19.1% of trading days. Data on exchange rates comes from Accominotti, Cen, Chambers, and Marsh (2019). See Table 3.3 for descriptive statistics. We also include various interest rate controls in some specifications, documented in Appendix 3.8.6.[12]

[9]Archives of the Bank of England, Harry Arthur Siepmann papers, reference C14/1, 1936.

[10]Telephone conversation with Mr. Bridge, Bank of England at 11:15 am, H. L. Sanford to file, 10 August 1956, New York, Archives of the Federal Reserve, box 617015.

[11]For more on the Bank of England intervention strategies, see Naef (2021).

[12]We also rely on gold reserve data from the Bank of England which we potentially allow to influence our forecasting and policy rules in Section 3.8.3. An earlier draft explored lagged gold reserves as an instrument but the results were uninformative due to low power in the first stage; see Naef and Weber (2021).

3.4 IV Estimates

Let e_t be the pound/dollar exchange rate at the end of day t, plotted in Figure 3.1, and let Q_t be intervention undertaken to appreciate the pound defined as dollar sales as a share of UK M0, plotted in Figure 3.2. A naive attempt to estimate the marginal effect of intervention on the exchange rate would be to estimate the following via OLS:

$$\ln e_t - \ln e_{t-1} = \beta_0 + \beta_1 Q_t + \epsilon_t, \tag{3.1}$$

where ϵ_t is an (unobserved) shock to growth in the exchange rate. Conventional wisdom suggests β_1 should be negative, as a sale of dollars (and purchase of pounds) should decrease the value of the dollar relative to the pound. But if Q_t is not randomly assigned, and the central bank "leans against the wind," estimates of β_1 will be biased upward.

To surmount this, we instrument for endogenous Q_t with the square of the distance from the exchange rate target at time $t-1$, which we allow to take on negative values when below target. Formally, the distance instrument $Z_t \equiv (\ln e_{t-1} - \ln e_{t-1}^{target})^2 \times \text{sign}(\ln e_{t-1} - \ln e_{t-1}^{target})$, where the target is time-varying only because of the devaluation in sample (note changes in the lagged, log level of the exchange rate drive the variation in this instrument so using $Z_t = \ln e_{t-1}$ yields very similar results; see the tables in Appendix 3.8.1.2). We motivate this instrument by arguing that if yesterday's market closed far from the target, then regardless of today's developments the dealers may intervene more aggressively.

Formally, exclusion requires that our instrument only impacts the growth of the exchange rate through its effect on the actions of the Bank of England, which were not observed directly by market participants at this time. To evaluate this, consider what it would mean for our exclusion assumption to *not* hold. This would imply that the level of the exchange rate was useful for forecasting its growth rate at very short horizons, and thus that traders consistently left money on the table in a large and liquid market.

A concern is that there may be predictable mean reversion at short horizons independent of the actions of the central bank, which would violate the exclusion restriction. Appendix 3.8.4 illustrates this concern by analyzing a simple first-order approximation of a reduced-form portfolio balance channel model, noting that mean reversion in (unobserved) fundamental shocks to the exchange rate violates the exclusion restriction. However, we note that exclusion will nearly hold if the shocks to the exchange rate follow a nearly random walk (i.e. a near unit root, stationary first-order autoregressive process).

Is this assumption valid? In the wake of Meese and Rogoff (1983), subsequent work documented that while exchange rates appeared well-modelled by a random walk, some mean reversion is detectable at longer horizons (Rossi, 2013). However, at the daily frequency considered here, the quantity of mean reversion consistent with long run estimates is quite small. Thus, while any deviations from a unit root process would violate our identification assumption and bias our estimates of β_1 downward, plausibly-small amounts of mean reversion are unlikely to be driving our results (Appendix 3.8.5 establishes this through simulations and discusses this issue in greater detail).

Table 3.1 presents estimates of β_1 in equation (3.1) comparing OLS to IV estimates. The results accord well with theory. When we estimate (3.1) by OLS, we find results consistent with severe upward bias as the coefficient is *positive* and significant, the opposite of what theory and the intuition of generations of central bankers suggests. However, when using the distance instrument, the sign flips, becoming negative and of reasonable magnitude. The interpretation of each coefficient is the increase (in percentage points) in the exchange rate that would result from a sale of dollars/purchase of pounds equivalent to 1% of British M0: the IV result suggests an effect of negative 10 basis points. Table 3.2 presents the first stage to verify the economic intuition. The first stage is strong and signs are as expected: when the pound is "too strong" relative to target yesterday, the Bank of England moved to weaken it (and vice versa when the pound was too weak).

We also consider changes to time period of estimation, estimating regressions for the pre-devaluation and post-devaluation periods separately, as well as a sample dropping the entire month of the devaluation, and finally a sample dropping data prior to 1959 as an important liberalization in UK capital markets occurred in late December 1958 when current account convertibility was restored as a consequence of the European Monetary Agreement. While the capital account was still not completely liberated, the policy meant much larger capital flows in and out of the UK. This last exercise is particularly useful for understanding whether the presence of greater capital controls in the 1950s is driving our results, which does not appear to be the case.

3.4.1 Robustness Checks and Comparision to the Literature

Appendix 3.8.1 considers several robustness checks. Appendix 3.8.1.1 adds controls to equation (3.1), guided by the model in Appendix 3.8.4, while Appendix 3.8.1.2 considers changing the instrument to use the lagged, log level of the exchange rate.[13] The results are consistent. Appendix 3.8.2 reports results measuring intervention as dollar sales, which also eases comparision with the literature. Those results suggest a sale of \$1bn USD appreciates the pound by 1.2 percentage points, which compares favorably with other estimates; Arango-Lozano, Menkhoff, Rodríguez-Novoa, and Villamizar-Villegas (2020) report an average effect of one percentage point in a meta study of 74 empirical studies.

Finally, Appendix 3.8.3 estimates the effect of intervention by first estimating a policy rule for intervention and treating deviations driven by UK-specific bank holiday's (when the Bank of England was closed and did not intervene much) as shocks to intervention. The point estimates are similar to those in the IV regressions above.

While the results in Table 3.1 range from estimated effects of 6-17 basis points, our preferred conclusion of an effect of 4-5 basis points stems from consideration of the large standard errors in that table and the estimates in Appendices 3.8.1.1 and 3.8.3, which are generally more precise (see e.g. Table 3.7, Column 4).

[13]The lagged, signed, squared distance from target instrument is our headline result as it has a stronger first-stage.

3.5 Conclusion

This paper established the presence of nontrivial effects of sterilized intervention even absent a significant signaling channel. Our results suggest that a sterilized, unannounced sale of dollars equivalent to 1% of UK M0 causes a 4-5 basis point appreciation of the pound. Given that the median absolute change in the exchange rate was 2.2 basis points during this period, our estimates are consistent with the view that the Bank of England's interventions were useful in offsetting day-to-day fluctuations in the exchange rate.

We took two approaches to identification to obtain these results. The first used the lagged, squared distance from target as an instrument for intervention, which requires the exclusion restriction that the level of the exchange rate is not useful for forecasting it's own growth rate except through its effects on the (secret) actions of the central bank. Since this would be violated by even a small quantity of mean reversion in the exchange rate, which would bias our results downward, we also took a second approach which treated deviations from an estimated policy rule for intervention as shocks. We showed that it was important to restrict our attention to shocks which occurred during holidays, when we know that deviation from normal behavior was due to the central bank being closed, and that doing so resulted in point estimates consistent with our IV regressions. Jointly, these results suggest the modest but economically significant effects on the level of the exchange rate which we report above.

While the context of the Bank of England during Bretton Woods is different from many central banks today, our results remain relevant to modern policymakers. Many modern central banks manage exchange rate pegs, even if their bands are often broader than those studied here, and intervene in foreign exchange markets. Our results at least partially rationalize the decision of many central banks to both intervene in secret and sterilize their interventions by demonstrating that such interventions can still have economically significant effects on the exchange rate. Our use of historical data affords us a relatively rare opportunity to study the track record of systematic, secret, sterilized intervention over a long period of time, confident in the knowledge that the data has not been manipulated prior to publication or selectively provided.

Figure 3.1: The Pound/Dollar Exchange Rate Over Time

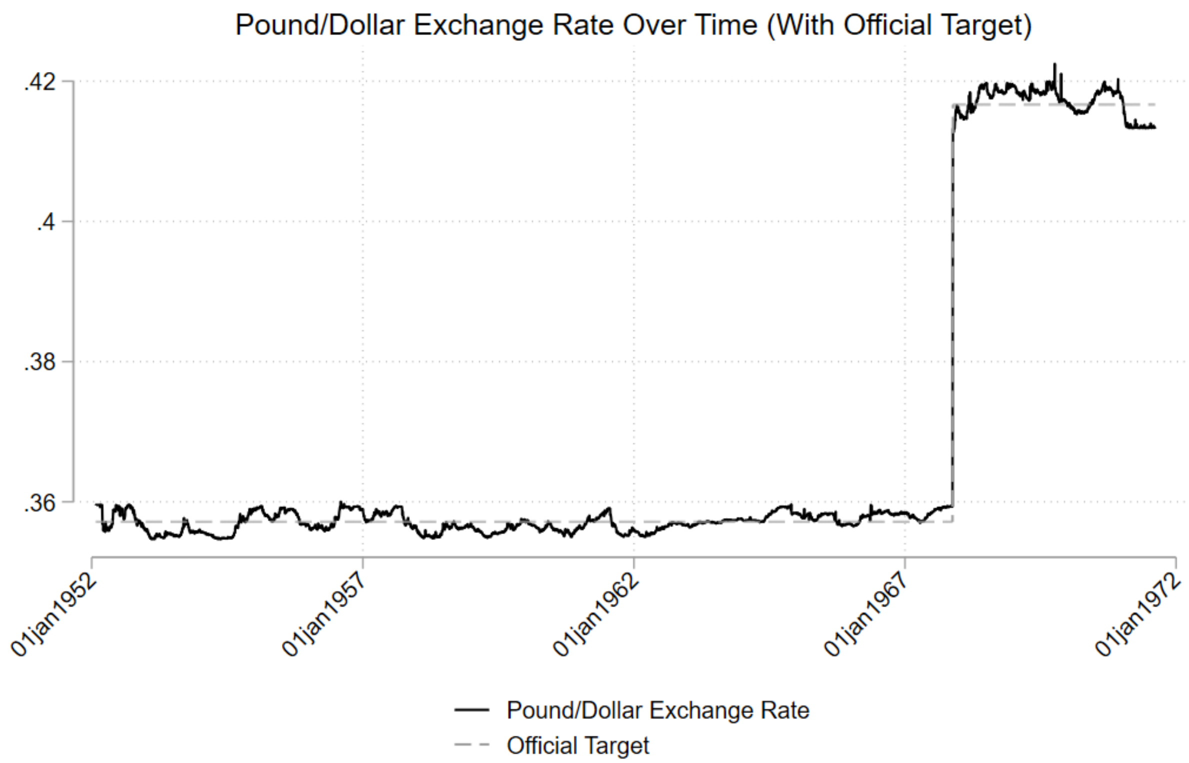

Notes: Note the devaluation of the pound in late November of 1967. See text for source.

Figure 3.2: Intervention Over Time

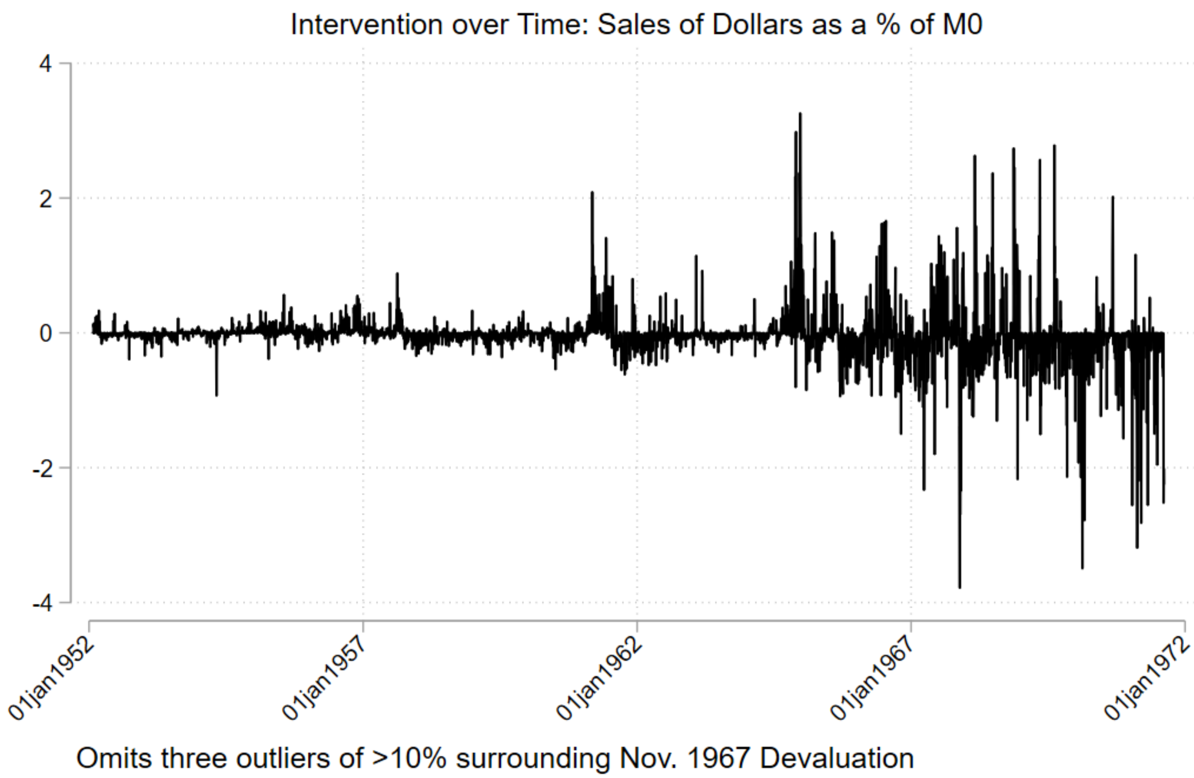

Notes: Intervention over time, deflated by UK M0. Positive values indicate sales of dollars by the Bank of England, understood as attempts to appreciate the pound vis-a-vis the dollar. This is the key right-hand-side variable in all regressions below. See text.

Table 3.1: Effect of Intervention on the Change in the Exchange Rate by Subsample [1952-1971]

	OLS	IV				
	(1) Full Sample	(2) Full Sample	(3) Pre-Devaluation	(4) Post-Devaluation	(5) Drop Nov. '67	(6) After 1958
Intervention	0.02^{***}	-0.10^{***}	-0.06^{***}	-0.17^{*}	-0.13^{***}	-0.10^{**}
	(0.01)	(0.03)	(0.02)	(0.07)	(0.04)	(0.03)
Observations	5244	5244	4278	966	5224	3277

Notes: Heteroskedasticity-robust standard errors in parentheses. The dependent variable is the log growth in the value of the dollar relative to the pound, and "Intervention" is the daily quantity of dollar sales by the Bank of England (divided by UK M0) undertaken to appreciate the pound. Columns (1) and (2) present OLS and IV estimates of the effects of intervention, demonstrating the bias in OLS and suggesting that an intervention equivalent to 1% of UK M0 appreciates the pound by 10 basis points. Columns (3)-(6) present IV estimates for subsamples, where (5) drops the entire month containing the devaluation (November of 1967) and (6) keeps only the period after an important liberalization in UK capital markets in 1958 when current account convertibility was restored as a consequence of the European Monetary Agreement. While the capital account was not completely liberated, the policy meant much larger capital flows in and out of the UK. All regressions include day-of-week, month and year dummies and drop the first trading day after the November 1967 devaluation.
Stars indicate: * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$

Table 3.2: First Stage Regressions: Effect on Intervention (Dollar Sales) by Subsample

	(1) Full Sample	(2) Pre-Devaluation	(3) Post-Devaluation	(4) Drop Nov. '67	(5) After 1958
Lagged, Squared Distance from Target	0.33^{***}	0.26^{***}	0.49^{***}	0.29^{***}	0.48^{***}
	(0.04)	(0.03)	(0.07)	(0.03)	(0.06)
Constant	-0.03^{***}	-0.00	-0.18^{***}	-0.03^{***}	-0.06^{***}
	(0.01)	(0.00)	(0.02)	(0.00)	(0.01)
Observations	5244	4278	966	5224	3277
R^2	0.041	0.028	0.084	0.040	0.051
F	90.47	59.84	45.29	103.50	67.39

Notes: Heteroskedasticity-robust standard errors in parentheses. The dependent variable is the daily quantity of intervention (measured as dollar sales divided by UK M0). The signs confirm the economic intuition underlying the relevance assumption: when the lagged distance from target instrument is positive, the pound is "too weak" relative to target and the Bank of England acts to strengthen it. All regressions drop the first trading day after the November 1967 devaluation.
Stars indicate: * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$

Table 3.3: Descriptive Statistics for Key Variables [1952-1971]

	Observations	Mean	Standard Deviation	Minimum	Maximum
Percent Growth in the Exchange Rate: $(\ln e_t - \ln e_{t-1}) \times 100$	5,244	-3.64e-06	0.0510	-0.891	0.865
Percent Distance from Target: $\left(\ln e_{t-1} - \ln e_{t-1}^{target}\right) \times 100$	5,244	0.00133	0.399	-0.959	1.39
Dollar Sales (in millions USD)	5,245	-3.09	33.6	-316	1,232
Dollar Sales as a Percentage of UK M0: $Q_t \times 100$	5,245	-0.0345	0.373	-3.78	12.8

Notes: Summary statistics for daily observations from 1952 to 1971 which exclude the first trading day after the devaluation (i.e. the same sample as Columns 1 and 2 of Table 3.1). Note dollar sales takes on negative values when the Bank of England purchases dollars. The variable Q_t comes from converting dollar sales into pounds (using the exchange rate target) and dividing by the previous month's value for UK M0.

Conclusion

This dissertation improves our understanding of how the actions of the Federal Reserve and other central banks affect real activity in the economy. Chapters 1 and 2 focused on understanding how the transmission of monetary policy shocks (changes in the federal funds rate, in the United States) has weakened over time due to the secular changes in the composition of investment spending.

Chapter 1 developed a model in which R&D and other intangible investment responds little to monetary policy shocks because of high adjustment costs stemming from congestion in onboarding new workers. Having tested the model's predictions against novel data on the productivity of individual software developers collaborating on GitHub, we conclude that investment adjustment costs are a deep, structural feature of the production of R&D and other intangible investment goods.

Chapter 2 takes this result seriously as the explanation for the low observed responsiveness in the cross section for intangible investment as compared to tangible investment. Given this, and the shift towards more and more intangible investment in the U.S. economy, this implies that investment spending overall is becoming less sensitive to changes in interest rates. Combined with two other secular changes documented in Chapter 2—a rising import share in investment spending and a decline in the labor share of domestically produced investment—this implies that consumption and labor incomes for hand-to-mouth agents (and thus consumption and labor incomes overall) respond less to monetary policy shocks in general equilibrium. A small open economy new Keynesian model calibrated to match these facts implies a 25% and 15% weaker response of labor income and aggregate consumption, respectively, to real interest rate shocks in a 2010's economy relative to a 1960's economy.

This evidence for the potential growing weakness of conventional monetary policy, even when the central bank is not constrained by the effective lower bound, motivates the final Chapter's interest in less conventional monetary policy tools. Chapter 3 explores the effects of sterilized foreign exchange intervention, and uses novel historical data from the Bank of England during Bretton Woods to show that even unannounced, sterilized interventions can meaningfully affect exchange rates. While the effects uncovered are modest, they suggest that unannounced, sterilized foreign exchange intervention can play a role in managing modest day-to-day fluctuations in a country's exchange rate. This finding at least partially rationalizes the choice of many modern central banks to intervene in secret.

However, Chapter 3 has little to say on whether or not central banks *should* engage in unannounced sterilized foreign exchange intervention. Similarly, Chapter 1 and Chapter

2 fall short of making normative claims about what policymakers should do when facing an economy that grows increasingly less interest rate sensitive by the year. In short, this dissertation explored *positive* questions about monetary policy, increasing our understanding of how e.g. a 1% change in the policy rate, or a sale of $1 million dollars in the foreign exchange market, affects key quantities and market prices.

One important normative question is how the waning power of monetary policy might bear on the costs and benefits of various proposals for dealing with the effective lower bound on interest rates, which range from raising the Fed's inflation target to abolishing physical currency. On the one hand, the U.S. experience of being at the zero lower bound from 2009 to 2015 following the Great Recession suggests that long periods at the zero lower bound may not be as costly as policymakers once thought (see again Figure 0.1), especially given central banks' increased willingness to use unconventional monetary policy, thus making dramatic action less attractive. On the other hand, if the Federal Reserve's main policy instrument is less effective in stimulating demand, it may need more space to maneuver now than in the past, making zero lower bound episodes more common. Future work should explore which of these effects dominates, and re-evaluate the desirability of implementing policy changes geared at mitigating the zero lower bound.

Bibliography

Accominotti, O., J. Cen, D. Chambers, and I. W. Marsh (2019). Currency regimes and the carry trade.

Adler, G., N. Lisack, and R. C. Mano (2019). Unveiling the effects of foreign exchange intervention: A panel approach. *Emerging Markets Review 40*, 100620.

Ahrens, A., C. B. Hansen, and M. E. Schaffer (2020). lassopack: Model selection and prediction with regularized regression in stata. *The Stata Journal 20*(1), 176–235.

Allen, W. A. (2019, February). *The Bank of England and the Government Debt: Operations in the Gilt-Edged Market, 1928-1972*. New York: Cambridge University Press.

Alves, F., G. Violante, G. Kaplan, and B. Moll (2020). A further look at the propagation of monetary policy shocks in hank. *Journal of Money, Credit and Banking*.

Anzoategui, D., D. Comin, M. Gertler, and J. Martinez (2019). Endogenous technology adoption and r&d as sources of business cycle persistence. *American Economic Journal: Macroeconomics 11*(3), 67–110.

Arango-Lozano, L., L. Menkhoff, D. Rodríguez-Novoa, and M. Villamizar-Villegas (2020). The effectiveness of fx interventions: A meta-analysis. *Journal of Financial Stability*, 100794.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica 89*(5), 2375–2408.

Auclert, A., M. Rognlie, and L. Straub (2018). The intertemporal keynesian cross. Technical report, National Bureau of Economic Research.

Auclert, A., M. Rognlie, and L. Straub (2020). Micro jumps, macro humps: monetary policy and business cycles in an estimated hank model. Technical report, National Bureau of Economic Research.

Baek, C. and B. Lee (2020). A guide to single equation regressions for impulse response estimations. Technical report.

Baldi, G. and A. Lange (2019). The interest rate sensitivity of investment. *Credit and Capital Markets 52*(2), 173–190.

BEA (2021, December). Concepts and methods of the u.s. national income and product accounts. Technical report, Bureau of Economic Analysis, https://www.bea.gov/resources/methodologies/nipa-handbook/pdf/all-chapters.pdf.

Bessec, M. (2003). Mean-reversion vs. adjustment to ppp: the two regimes of exchange rate dynamics under the ems, 1979–1998. *Economic Modelling 20*(1), 141–164.

Best, M. C., J. S. Cloyne, E. Ilzetzki, and H. J. Kleven (2020). Estimating the Elasticity of Intertemporal Substitution using Mortgage Notches. *The Review of Economic Studies 87*(2), 656–690.

Bianchi, F., H. Kung, and G. Morales (2019). Growth, slowdowns, and recoveries. *Journal of Monetary Economics 101*, 47–63.

Bilbiie, F. O., D. Känzig, and P. Surico (2020). Capital and income inequality: An aggregate-demand complementarity.

Blanchard, O., G. Adler, and I. d. C. Filho (2015, July). Can Foreign Exchange Intervention Stem Exchange Rate Pressures from Global Capital Flow Shocks? Working Paper 21427, National Bureau of Economic Research.

Bloesch, J. and J. Weber (2021). Structural changes in investment and the waning power of monetary policy. *Available at SSRN 3809439*.

Bloesch, J. and J. Weber (2023). Congestion in Onboarding Workers and Sticky R&D. Working Paper.

Boivin, J., M. T. Kiley, and F. S. Mishkin (2010). How has the monetary transmission mechanism evolved over time? In *Handbook of monetary economics*, Volume 3, pp. 369–422. Elsevier.

Bordo, M. D., O. F. Humpage, and A. J. Schwartz (2012). The Federal Reserve as an Informed Foreign Exchange Trader: 1973–1995. *International Journal of Central Banking*.

Caggese, A. and A. Pérez-Orive (2020). How stimulative are low real interest rates for intangible capital? Technical report, UPF working paper.

Caplin, A., M. Lee, S. Leth-Petersen, J. Sæverud, and M. D. Shapiro (2022, August). How worker productivity and wages grow with tenure and experience: The firm perspective. Working Paper 30342, National Bureau of Economic Research.

Casares, M. (2006). Time-to-build, monetary shocks, and aggregate fluctuations. *Journal of Monetary Economics 53*(6), 1161–1176.

Cavallo, M. and A. Landry (2010). The quantitative role of capital goods imports in us growth. *American Economic Review 100*(2), 78–82.

Cavallo, M. and A. Landry (2018). Capital-goods imports and us growth. Technical report, Bank of Canada Staff Working Paper.

Chamon, M., D. J. Hofman, N. E. Magud, and A. M. Werner (2019). *Foreign exchange intervention in inflation targeters in Latin America*. International Monetary Fund.

Chang, D. and C. Zimmermann (2019, 07). 20/19 hindsight: Checking policymakers' economic predictions against the data. https://fredblog.stlouisfed.org/2019/07/20-19-hindsight/. Accessed May 6th 2023.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy 113*(1), 1–45.

Christiano, L. J., M. S. Eichenbaum, and M. Trabandt (2018). On dsge models. *Journal of Economic Perspectives 32*(3), 113–40.

Cloyne, J., C. Ferreira, M. Froemel, and P. Surico (2018). Monetary policy, corporate finance and investment. Technical report, National Bureau of Economic Research.

Cloyne, J., C. Ferreira, and P. Surico (2020). Monetary policy when households have debt: new evidence on the transmission mechanism. *The Review of Economic Studies 87*(1), 102–129.

Cloyne, J., J. Martinez, H. Mumtaz, and P. Surico (2022). Short-term tax cuts, long-term stimulus. Technical report, National Bureau of Economic Research.

Comin, D. and M. Gertler (2006). Medium-term business cycles. *American Economic Review 96*(3), 523–551.

Cosentino, V., J. Luis, and J. Cabot (2016). Findings from github: Methods, datasets and limitations. In *Proceedings of the 13th International Conference on Mining Software Repositories*, pp. 137–141.

Cushman, D. O. (2007). A portfolio balance approach to the canadian–us exchange rate. *Review of Financial Economics 16*(3), 305–320.

Daude, C., E. L. Yeyati, and A. J. Nagengast (2016). On the effectiveness of exchange rate interventions in emerging markets. *Journal of International Money and Finance 64*, 239–261.

De Ridder, M. (2019). Market power and innovation in the intangible economy.

Degasperi, R., S. Hong, and G. Ricco (2023). The Global Transmission of US Monetary Policy. Working Paper.

Del Negro, M., M. Giannoni, and C. Patterson (2015). The Forward Guidance Puzzle. *Mimeo*.

Dominguez, K. M., R. Fatum, and P. Vacek (2013). Do sales of foreign exchange reserves lead to currency appreciation? *Journal of Money, Credit and Banking 45*(5), 867–890.

Döttling, R. and L. Ratnoski (2020). Monetary policy and intangible investment.

Edge, R. M. (2007). Time-to-build, time-to-plan, habit-persistence, and the liquidity effect. *Journal of Monetary Economics 54*(6), 1644–1669.

Eldridge, L. P., C. Garner, T. F. Howells, B. C. Moyer, M. Russell, J. D. Samuels, E. H. Strassner, and D. B. Wasshausen (2020). Toward a bea-bls integrated industry-level production account for 1947–2016. In *Measuring Economic Growth and Productivity*, pp. 221–249. Elsevier.

Erceg, C. J., D. W. Henderson, and A. T. Levin (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics 46*(2), 281–313.

Fatum, R. and M. R. King (2005). Rules versus discretion in foreign exchange intervention: evidence from official bank of canada high-frequency data.

Flynn, J. P., C. Patterson, and J. Sturm (2020). Shock propagation and the fiscal multiplier: the role of heterogeneity. *Available at SSRN*.

Forrest, C. (2010). *The Bank of England 1950s to 1979.* Cambridge University Press, Cambridge.

Forsgren, N., M.-A. Storey, C. Maddila, T. Zimmermann, B. Houck, and J. Butler (2021). The space of developer productivity: There's more to it than you think. *Queue 19*(1), 20–48.

Frankel, J. (1984). Tests of monetary and portfolio balance models of exchange rate determination. *Exchange rate theory and practice (University of Chicago Press, Chicago)*.

Fratzscher, M., O. Gloede, L. Menkhoff, L. Sarno, and T. Stöhr (2019). When Is Foreign Exchange Intervention Effective? Evidence from 33 Countries. *American Economic Journal: Macroeconomics*.

Gabaix, X. and M. Maggiori (2015, August). International Liquidity and Exchange Rate Dynamics. *The Quarterly Journal of Economics 130*(3), 1369–1420.

Gousios, G. (2013). The ghtorrent dataset and tool suite. In *Proceedings of the 10th Working Conference on Mining Software Repositories*, MSR '13, Piscataway, NJ, USA, pp. 233–236. IEEE Press.

Gousios, G., M. Pinzger, and A. v. Deursen (2014). An exploratory study of the pull-based software development model. In *Proceedings of the 36th international conference on software engineering*, pp. 345–355.

Gousios, G. and D. Spinellis (2012). Ghtorrent: Github's data from a firehose. In *2012 9th IEEE Working Conference on Mining Software Repositories (MSR)*, pp. 12–21. IEEE.

Griliches, Z. (1990). Patent statistics as economic indicators: A survey. *Journal of Economic Literature 28*(4), 1661–1707.

Hall, B. H. and J. Lerner (2010). The financing of r&d and innovation. In *Handbook of the Economics of Innovation*, Volume 1, pp. 609–639. Elsevier.

Hann, I., J. Roberts, S. Slaughter, and R. Fielding (2004). An empirical analysis of economic returns to open source participation. *Unpublished working paper, Carnegie-Mellon University*.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica*, 213–224.

Herrendorf, B., R. Rogerson, and Á. Valentinyi (2020). Structural change in investment and consumption: A unified analysis. *The Review of Economic Studies*.

Hertel, G., S. Niedner, and S. Herrmann (2003). Motivation of software developers in open source projects: An internet-based survey of contributors to the linux kernel. *Research policy 32*(7), 1159–1177.

Horowitz, K. and M. Planting (2009). Concepts and methods of the input-output accounts. bureau of economic analysis (bea), us department of commerce. 2006. Technical report, updated 2009—1-206 pp.

House, C. L., A.-M. Mocanu, and M. D. Shapiro (2017). Stimulus effects of investment tax incentives: Production versus purchases. Technical report, National Bureau of Economic Research.

Howes, C. and A. von Ende-Becker (2022). Monetary policy and intangible investment. *Economic Review 107*(2).

Howson, S. (1980). *Sterling's managed float: the operations of the exchange equalisation account, 1932-39*. International Finance Section, Department of Economics, Princeton University.

Hubmer, J. (2020). The race between preferences and technology. *Unpublished Working Paper*.

Jaravel, X., N. Petkova, and A. Bell (2018). Team-specific capital and innovation. *American Economic Review 108*(4-5), 1034–73.

JetBrains (2021). The state of developer ecosystem 2021. Technical report, https://www.jetbrains.com/lp/devecosystem-2021/.

Jones, C. I. (2011). Misallocation, economic growth, and input-output economics. Technical report, National bureau of economic research.

Jorgenson, D. W., M. S. Ho, and J. D. Samuels (2017). Educational attainment and the revival of us economic growth. In *Education, Skills, and Technical Change: Implications for Future US GDP Growth*. University of Chicago Press Chicago.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2010). Investment shocks and business cycles. *Journal of Monetary Economics 57*(2), 132–145.

Kalliamvakou, E., D. Damian, K. Blincoe, L. Singer, and D. M. German (2015). Open source-style collaborative development practices in commercial projects using githhub. In *2015 IEEE/ACM 37th IEEE international conference on software engineering*, Volume 1, pp. 574–585. IEEE.

Kalliamvakou, E., G. Gousios, K. Blincoe, L. Singer, D. M. German, and D. Damian (2014). The promises and perils of mining github. In *Proceedings of the 11th working conference on mining software repositories*, pp. 92–101.

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to hank. *American Economic Review 108*(3), 697–743.

Kaplan, G. and G. L. Violante (2022). The Marginal Propensity to Consume in Heterogeneous Agent Models. *Annual Review of Economics 14*, 747–775.

Kaplan, G., G. L. Violante, and J. Weidner (2014). The Wealthy Hand-to-Mouth. *Brookings Papers on Economic Activity 1*.

Kerr, W. R. and R. Nanda (2015). Financing innovation. *Annual Review of Financial Economics 7*, 445–462.

Kim, S. (2001). International transmission of us monetary policy shocks: Evidence from var's. *Journal of monetary Economics 48*(2), 339–372.

Kline, P., N. Petkova, H. Williams, and O. Zidar (2019). Who profits from patents? rent-sharing at innovative firms. *Quarterly Journal of Economics 134*(3), 1343–1404.

Koh, D., R. Santaeulàlia-Llopis, and Y. Zheng (2020). Labor share decline and intellectual property products capital. *Econometrica 88*(6), 2609–2628.

Krugman, P., M. Obstfeld, and M. Melitz (2015). International economics: Theories and policy.

Kuersteiner, G. M., D. C. Phillips, and M. Villamizar-Villegas (2018, July). Effective sterilized foreign exchange intervention? Evidence from a rule-based policy. *Journal of International Economics 113*, 118–138.

Kydland, F. E. and E. C. Prescott (1982). Time to build and aggregate fluctuations. *Econometrica*, 1345–1370.

Lakhani, K. R. and R. G. Wolf (2003). Why hackers do what they do: Understanding motivation and effort in free/open source software projects. *Open Source Software Projects (September 2003)*.

Lazear, E. P. (2009). Firm-specific human capital: A skill-weights approach. *Journal of Political Economy 117*(5), 914–940.

Lerner, J. and J. Tirole (2005). The economics of technology sharing: Open source and beyond. *Journal of Economic Perspectives 19*(2), 99–120.

Li, W. C. and B. H. Hall (2020). Depreciation of business r&d capital. *Review of Income and Wealth 66*(1), 161–180.

Lucca, D. (2007). Resuscitating time-to-build. *Manuscript, Federal Reserve Board*.

Lyndaker, A. S., T. F. Howells III, E. H. Strassner, and D. B. Wasshausen (2016). Integrated historical input-output and gdp-by-industry accounts, 1947-1996. *Survey of Current Business 96*(2).

McDermott, G. R. and B. Hansen (2021). Labor reallocation and remote work during covid-19: Real-time evidence from github. Technical report, National Bureau of Economic Research.

McKay, A., E. Nakamura, and J. Steinsson (2016). The power of forward guidance revisited. *The American Economic Review 106*(10), 3133–3158.

Medeiros, M. C. and E. F. Mendes (2016). $\ell 1$-regularization of high-dimensional time-series models with non-gaussian and heteroskedastic errors. *Journal of Econometrics 191*(1), 255–271.

Meese, R. A. and K. Rogoff (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of international economics 14*(1-2), 3–24.

Mehran, H. and J. S. Tracy (2001). The effect of employee stock options on the evolution of compensation in the 1990s. *Economic Policy Review 7*(3).

Menkhoff, L. (2010). High-frequency analysis of foreign exchange interventions: what do we learn? *Journal of Economic Surveys 24*(1), 85–112.

Mercan, Y., B. Schoefer, and P. Sedláček (2021). A congestion theory of unemployment fluctuations. Technical report, National Bureau of Economic Research.

Mergel, I. (2015). Open collaboration in the public sector: The case of social coding on github. *Government Information Quarterly 32*(4), 464–472.

Mohanty, M. S. and B.-e. Berger (2013, October). Central Bank Views on Foreign Exchange Intervention. *No. 73*.

Mombach, T. O. (2019). A comparative study of apis for querying github data.

Moran, P. and A. Queralto (2018). Innovation, productivity, and monetary policy. *Journal of Monetary Economics 93*, 24–41.

Moris, F. (2019, April). Software r&d: Revised treatment in u.s. national accounts and related trends in business r&d expenditures. Technical report, National Science Foundation, https://www.nsf.gov/statistics/2019/nsf19315/nsf19315.pdf.

Moylan, C. E. and S. Okubo (2020, March). The evolving treatment of r&d in the u.s. national economic accounts. Technical report, Bureau of Economic Analysis, https://www.bea.gov/system/files/2020-04/the-evolving-treatment-of-rd-in-the-us-national-economic-accounts.pdf.

Murray, C. J. and D. H. Papell (2002). The purchasing power parity persistence paradigm. *Journal of International Economics 56*(1), 1–19.

Naef, A. (2019). Blowing against the Wind? A Narrative Approach to Central Bank Foreign Exchange Intervention. *University of California, Berkeley Mimeo*.

Naef, A. (2021). Dirty float or clean intervention? the bank of england in the foreign exchange market. *European Review of Economic History 25*(1), 180–201.

Naef, A. and J. P. Weber (2021). How powerful is unannounced, sterilized foreign exchange intervention? Technical report.

Naef, A. and J. P. Weber (2023). How powerful is unannounced, sterilized foreign exchange intervention? *Journal of Money, Credit and Banking*.

Nakamura, E. and J. Steinsson (2018). Identification in macroeconomics. *Journal of Economic Perspectives 32*(3), 59–86.

Neely, C. J. (2008, February). Central bank authorities' beliefs about foreign exchange intervention. *Journal of International Money and Finance 27*(1), 1–25.

Peters, R. H. and L. A. Taylor (2017). Intangible capital and the investment-q relation. *Journal of Financial Economics 123*(2), 251–272.

Rieth, M., L. Menkhoff, and T. Stöhr (2019). The dynamic impact of FX interventions on financial markets. Technical report.

Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review 94*(4), 1055–1084.

Rossi, B. (2013). Exchange rate predictability. *Journal of economic literature 51*(4), 1063–1119.

Sarno, L. and M. P. Taylor (2001). Official Intervention in the Foreign Exchange Market: is it Effective and, if so, How Does it Work? *Journal of Economic Literature 39*(3), 839–868.

Scalia, A. (2008). Is foreign exchange intervention effective? some microanalytical evidence from the czech republic. *Journal of International Money and Finance 27*(4), 529–546.

Schmitt-Grohé, S. and M. Uribe (2003). Closing small open economy models. *Journal of international Economics 61*(1), 163–185.

Schmöller, M. E. and M. Spitzer (2021). Deep recessions, slowing productivity and missing (dis-) inflation in the euro area. *European Economic Review 134*, 103708.

Smets, F. and R. Wouters (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American Economic Review 97*(3), 586–606.

Stol, K.-J., P. Avgeriou, M. A. Babar, Y. Lucas, and B. Fitzgerald (2014). Key factors for adopting inner source. *ACM Transactions on Software Engineering and Methodology (TOSEM) 23*(2), 1–35.

Subramanian, V. N. (2020). An empirical study of the first contributions of developers to open source projects on github. In *2020 IEEE/ACM 42nd International Conference on Software Engineering: Companion Proceedings (ICSE-Companion)*, pp. 116–118. IEEE.

Summers, L. H. and A. Stansbury (2019). Whither central banking?

Sun, Q. and M. Z. Xiaolan (2019). Financing intangible capital. *Journal of Financial Economics 133*(3), 564–588.

Sweeney, R. J. (2006). Mean reversion in g-10 nominal exchange rates. *Journal of Financial and Quantitative Analysis 41*(3), 685–708.

Taylor, M. P. and D. A. Peel (2000). Nonlinear adjustment, long-run equilibrium and exchange rate fundamentals. *Journal of international Money and Finance 19*(1), 33–53.

Torkar, R., P. Minoves, and J. Garrigós (2011). Adopting free/libre/open source software practices, techniques and methods for industrial use. *Journal of the Association for Information Systems 12*(1), 1.

Van Zandweghe, W. and J. C. Braxton (2013). Has durable goods spending become less sensitive to interest rates? *Economic Review-Federal Reserve Bank of Kansas City*, 5.

Vasilescu, B., D. Posnett, B. Ray, M. G. van den Brand, A. Serebrenik, P. Devanbu, and V. Filkov (2015). Gender and tenure diversity in github teams. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, pp. 3789–3798.

Wieland, J. F. and M.-J. Yang (2020). Financial dampening. *Journal of Money, Credit and Banking 52*(1), 79–113.

Willis, J. L. and G. Cao (2015). Has the us economy become less interest rate sensitive? *Economic Review (01612387) 100*(2).

Wyrich, M., R. Ghit, T. Haller, and C. Müller (2021). Bots don't mind waiting, do they? comparing the interaction with automatically and manually created pull requests. In *2021 IEEE/ACM Third International Workshop on Bots in Software Engineering (BotSE)*, pp. 6–10. IEEE.

Appendices

3.6 Appendix to Chapter 1

3.6.1 BEA's Treatment of Software and R&D Spending in Intellectual Property Products (IPP) Investment

This appendix presents summary statistics illustrating the growing importance of software and R&D in US investment spending, and elaborates on the various ways software spending appears in the NIPAs.

The BEA began capitalizing expenditures on software as investment in 1999, and other R&D expenses as investment in 2013, reflecting their growing importance. These are generally measured at cost, including e.g. the wages and salaries of workers involved in development; see the NIPA handbook Ch. 6 (BEA, 2021) for details. Non-residential, fixed investment (i.e. not counting inventories) thus now consists of structures, equipment, and a new category called "Intellectual Property Products" (IPP). IPP contains both software expenditures, R&D (including software R&D), and a small share of "literary arts and originals" investment, e.g. the production of films, books, etc. The BEA's definition of intangible investment (IPP) is narrow in the sense that the NIPAs do not capitalize e.g. marketing or advertising expenses, finance and insurance costs of new product development, training costs, or organizational capital as investment; see Koh, Santaeulàlia-Llopis, and Zheng (2020) for a discussion. Figure 3.3 illustrates this new breakdown for fixed investment quantitatively for the year 2021.[14] Ignoring the "Literary and Artistic Originals" component, which has remained stable as a share of investment, the remaining components of IPP have risen steadily as a share of US investment, as shown in Figure 3.4.

The "Software" category of IPP includes purchases of prepackaged software and of customized software from companies that are "primarily engaged in software development," as well as expenditures for the own-account production of new or "significantly enhanced" software that a firm develops in-house.[15] Own-account software does not include the development

[14]This figure replicates Figure 1 in Howes and von Ende-Becker (2022) but for the year 2021 instead of 2020. Note also that their Figure 1's exact dollar amounts reflect outdated GDP statistics: as of the September 29th, 2022 revision IPP was larger than equipment in 2020 as claimed in the introduction (see NIPA table 1.1.5).

[15]Prepackaged software excludes software embedded, or bundled, in computers and other equipment (Moris, 2019).

Figure 3.3: The Components of US Fixed Investment

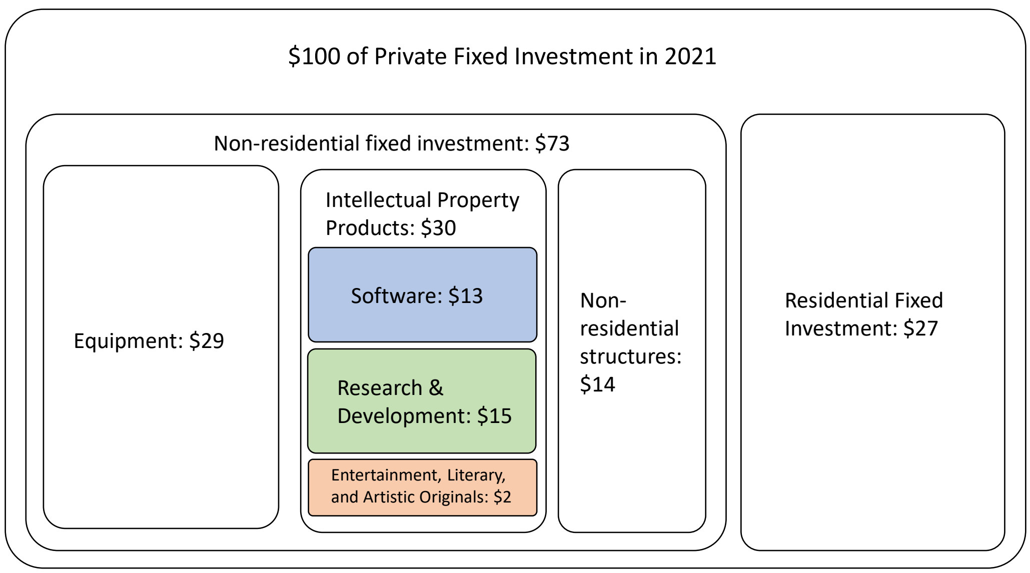

Notes: The components of fixed investment (i.e. excluding inventories) in the US national accounts. Note that a large share of R&D is software R&D. Source: BEA and Authors' calculations.

Figure 3.4: Secular Rise in Software and R&D Investment

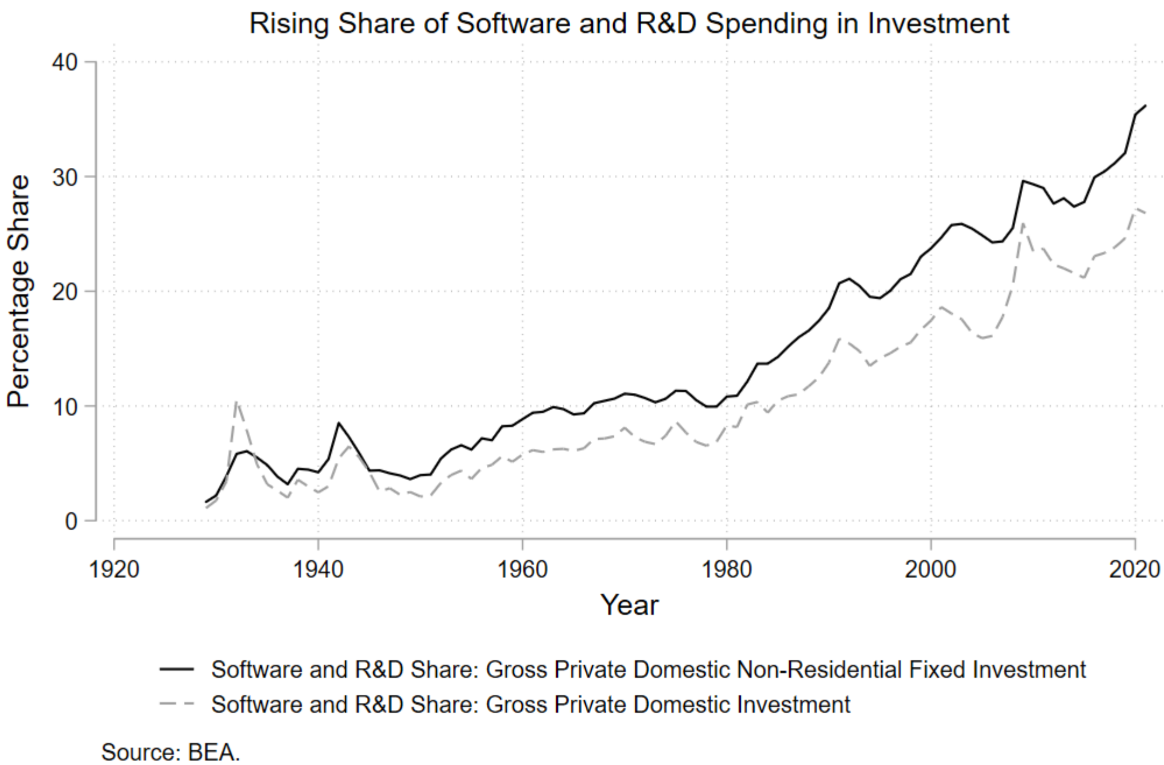

Notes: Software and R&D – the two largest components of IPP investment – have risen steadily as a share of US investment. The excluded category "Literary and Artistic Originals" is a small share of U.S. investment and has been stable over time. Source: BEA.

of software originals from which copies are made for sale (i.e. product development) or incorporated into other products (such as vehicles or appliances); these expenses are instead included in the R&D category of IPP (BEA, 2021) reflecting recent changes in 2018 (Moylan and Okubo, 2020). Roughly 1/3 of R&D is software R&D (32.2% in 2016), reflecting a secular increase over the past two decades (Moris, 2019).

While the BEA does not currently publish the components of R&D separately by type, underlying NSF survey data permits separating software R&D from other kinds of R&D for specific years. Table 3.4 presents a breakdown of IPP for 2016 using data from the NSF in Moris (2019) and the BEA to break out software R&D from other R&D, showing that software expenditures make up a majority of IPP.

Table 3.4: Composition of Non-Residential Investment in 2016

Category	**Investment Share (Ppt.)**
Software R&D and Other Software Investment	18.8
Software R&D	5.3
Other Software Investment	13.6
Non-software R&D	11.1
Literary and Artistic Originals	3.3
Equipment and Structures	66.7

Notes: Shares of U.S. gross private domestic fixed, non-residential investment. Source: authors' calculations from BEA data and NSF data in Moris (2019).

3.6.2 Proof of Proposition 1 and Discussion

Statement: Consider the problem of a firm choosing paths $\{I_{t+1}, J_t, S_t\}_{t=0}^{\infty}$ subject to the law of motion (1.1) and the production function $I_t = S_{t-1}^{\nu}$ to maximize the present discounted value of current and future profits (1.2). In a solution where (1.1) binds and $J_t > 0$ always, then the firm's problem can be written as:

$$\max_{\{I_{t+1}, J_t, S_t\}_{t=0}^{\infty}} \mathbf{E}\left[\sum_{t=0}^{\infty} \Lambda_{0,t}\left[P_t^k I_t - W_t(S_{t-1} + J_t)\right]\right]$$

subject to

$$\begin{aligned} I_t &= S_{t-1}^{\nu} \\ S_t &= (1-d)S_{t-1} + \rho\left(\frac{J_t}{S_{t-1}}\right) J_t \end{aligned}$$

where $\rho(x) \in [0,1]$ on $x \in [0,\infty)$ and $\rho'(x) < 0$. Let $f(x) \equiv \rho(x)x$ be strictly increasing on some domain D that does not restrict the firm's optimal choice. Then there exists an equivalent maximization problem yielding the same solution for I_t:

$$\max_{\{I_{t+1}\}_{t=0}^{\infty}} \mathbf{E}\left[\sum_{t=0}^{\infty} \Lambda_{0,t}\left[P_t^k I_t - W_t\left(1 + \underbrace{\Phi\left(\frac{I_{t+1}}{I_t}\right)}_{\substack{\text{Convex Adjustment Costs} \\ \text{from Onboarding}}}\right) I_t^{\frac{1}{\nu}}\right]\right]$$

and a domain G which does not restrict firm's optimal choice and where $\Phi' > 0$ on G. Further, if $f''(x) < 0$ on D then $\Phi'' > 0$ on G.

Proof: Since the law of motion (1.1) from the main text binds and $J_t > 0$ by assumption (so that the complementary slackness condition on this constraint can be ignored), rewrite the binding law of motion for S_t, equation (1.1), as:

$$\frac{S_t}{S_{t-1}} = (1-d) + \rho\left(\frac{J_t}{S_{t-1}}\right)\frac{J_t}{S_{t-1}} \equiv (1-d) + f\left(\frac{J_t}{S_{t-1}}\right) \tag{3.2}$$

Assume the function $f\left(\frac{J_t}{S_{t-1}}\right) = \rho\left(\frac{J_t}{S_{t-1}}\right)\frac{J_t}{S_{t-1}}$ is strictly increasing (and therefore invertible) and concave in $\frac{J_t}{S_{t-1}}$ on D. Note D is a subset of $[0,\infty)$ since $J_t \geq 0$ implies $J_t/S_{t-1} \geq 0$. Then (3.2) implies that $\frac{S_t}{S_{t-1}}$ is a concave, strictly increasing function of the term $(1-d)$ and the ratio $\frac{J_t}{S_{t-1}}$. Define this function as $F(\frac{J_t}{S_{t-1}})$, suppressing dependence on d, such that $F^{-1}(\frac{S_t}{S_{t-1}}) = \frac{J_t}{S_{t-1}}$ is the inverse of $F(\cdot)$.[16] Then F^{-1} is convex and strictly increasing on $G \equiv F(x) \forall x \in D$. Pulling an S_{t-1} out of the final term in per-period profits in (1.2) and plugging this in for the resulting J_t/S_{t-1} term yields:

$$P_t^k I_t - W_t\left(S_{t-1} + F^{-1}\left(\frac{S_t}{S_{t-1}}\right)S_{t-1}\right)$$

Now note since $\frac{S_t}{S_{t-1}}$ is an increasing, convex function of $\frac{I_{t+1}}{I_t}$, i.e. $\frac{S_t}{S_{t-1}} = (\frac{I_{t+1}}{I_t})^{\frac{1}{\nu}}$, it follows that F^{-1} is an increasing, convex function of $\frac{I_{t+1}}{I_t}$. Substituting in, we obtain the following for profits in period t:

$$P_t^k I_t - W_t\left(1 + \underbrace{\Phi\left(\frac{I_{t+1}}{I_t}\right)}_{\substack{\text{Convex Adjustment Costs}\\ \text{from Onboarding}}}\right)I_t^{\frac{1}{\nu}}$$

which yields the result.

Discussion: for $\Phi(\cdot)$ to be increasing and convex on G, we need $f(x) = \rho(x)x$ to be increasing and concave on D. Neither follows easily from the assumptions $\rho'(x) < 0$ and $\rho(x) \in [0,1]$. To see this consider the expressions for f' and f'' in terms of ρ,

$$f'(x) = \rho'(x)x + \rho(x)$$
$$f''(x) = \rho''(x)x + 2\rho'(x)$$

Note that for x small enough, we can always find a neighborhood where $f'(x) > 0$ and $f''(x) < 0$ under the assumption that $\rho'(x) < 0$ *and the additional assumptions* that $\rho'(x)$ and $\rho''(x)$ are bounded as $x \to 0$, since under these added assumptions

$$\lim_{x\to 0} f'(x) = \rho(0) > 0$$
$$\lim_{x\to 0} f''(x) = 2\rho'(0) < 0$$

[16]This function was defined as $F^{-1}(\cdot) \equiv \boldsymbol{\mathcal{F}}(\cdot)$ in the text's Section 1.2.

The neighborhood with $x \approx 0$ is of interest because this corresponds to a steady state of the model: when $f(x) = d$, the model is in steady state where $S_t/S_{t-1} = 1$. When d is small, x may also be small (or zero if $d = 0$). So we can have local concavity of f (and convexity of Φ) under very mild assumptions about the boundedness of the first and second derivative of ρ at zero. This is relevant since many aggregate models log-linearize around a steady state, and only require that the adjustment cost function be convex when evaluated at that point. Of course when d is large, the model's steady state may be far from $x \approx 0$. Thus, we may remain concerned that our assumptions on f may not hold for large x.

To alleviate these concerns, we note that the requirement that $f'(x) > 0$ and $f''(x) < 0$ does not put overly restrictive requirements on ρ in light of our empirical results. In practice, those results suggest that ρ:

- is a function of J_t/S_{t-1}
- satisfies $\rho(x) \in [0, 1]$ on $x \in [0, \infty)$ and $\rho'(x) < 0$
- is convex, i.e. $\rho''(x) > 0$.

One function that satisfies all these properties is $\rho(x) \equiv \frac{1}{ax+b} + c$, given appropriate choices of $a, c \geq 0$ and $b > 0$ so that $\rho(0)$ is well-defined. To see this, note ρ is decreasing and convex:

$$\rho'(x) = -a \left(\frac{1}{ax+b}\right)^2 < 0 \quad \forall x \geq 0$$
$$\rho''(x) = 2a^2 \left(\frac{1}{ax+b}\right)^3 > 0 \quad \forall x \geq 0$$

while f is increasing and concave:

$$\begin{aligned}
f'(x) &= \rho'(x)x + \rho(x) \\
&= -a \left(\frac{1}{ax+b}\right)^2 x + \frac{1}{ax+b} + c \\
&= \left(\frac{1}{ax+b}\right) \left(\frac{-ax}{ax+b} + 1\right) + c > 0 \quad \forall x \geq 0 \\
f''(x) &= \rho''(x)x + 2\rho'(x) \\
&= 2a^2 \left(\frac{1}{ax+b}\right)^3 x - 2a \left(\frac{1}{ax+b}\right)^2 \\
&= 2a \left(\frac{1}{ax+b}\right)^2 \left(\frac{ax}{ax+b} - 1\right) < 0 \quad \forall x \geq 0.
\end{aligned}$$

Additionally, a linear function or linear approximation will also work: if $\rho(x) = b - ax$ then $f(x) = bx - ax^2$ is quadratic, and no cost-minimizing firm will every choose a point where $x > \frac{b}{2a}$. So for the domain $x \in D \equiv [0, \frac{b}{2a})$ which does not restrict the choices of a cost-minimizing firm:[17]

$$\begin{aligned} f'(x) &= b - 2ax > 0 \quad && \forall x \in \left[0, \frac{b}{2a}\right) \\ f''(x) &= -2a < 0 && \forall x. \end{aligned}$$

3.6.3 Narrative Evidence on "Project-Specific Capital"

The returns to project-specific tenure that we document reflects a combination of skill-acquisition and earned trust or reputation within a team, which our model in Section 1.2 is general enough to encompass. We emphasize the acquisition of project-specific skills, as this frequently arises in interviews with practitioners. While human capital acquired while working on a specific project may in principle be portable (and imperfectly captured by our controls for overall programming experience) in general the evidence that knowledge gained by working on OSS projects is applicable elsewhere is weak: in a longitudinal study of contributors to the Apache project, Hann et al. (2004) found that increases in human capital as measured by total contributions to the project, did not lead to increased wages.[18]

What is this non-portable, "project-specific" human capital acquired in the first few months? Even for experienced developers, joining a new project or a new team entails acquiring knowledge specific to how that team operates and how existing code is structured ("software architecture"). Interviews with developers reveal that project-specific knowledge such as learning about the needs and requirements of end users, the "dos and don'ts" of design for a particular project or company, and – for tacit knowledge – "knowing who knows what" all plays a role in making a newcomer productive on a software development team (Stol et al., 2014).

Consistent with this, early studies recommending the adoption of OSS software development practices within private firms highlighted the advantages of adapting OSS development methods for private industry because of their ability to reduce onboarding times: these methods include making the entire history of design decisions and code base accessible to newcomers, and also to assign them relatively easier tasks that they can use to build skills and demonstrate their newly-acquired competence. As noted by Torkar et al. (2011),

> "It is important to have a predefined path that allows new developers to learn while doing productive activities... If this issue is left unattended, there is a risk of placing newcomers in positions for which they are unqualified or making

[17] There are two cases for any $x > 0$ not in D. For any choice $x \in [\frac{b}{2a}, \frac{b}{a}]$, there is a choice $x \in D$ that weakly dominates because it achieves the same growth in S_t at a smaller cost. Choosing very large $x > \frac{b}{a}$ means paying $W_t J_t$ to *reduce* the stock of S_t given S_{t-1} – assuming that the law of motion for S binds rules out the possibility that this is optimal (in our quantitative exercise, we allow for "free disposal" of S which guarantees firms would never choose $x > \frac{b}{a}$).

[18] The authors took advantage of the Apache project's unusual hierarchy, which includes five rankings, to show that instead these earned credentials explained wage growth, which they interpret as consistent with a signaling theory of the benefits to contributing to an OSS project.

> their learning curve unnecessarily long. With proper support from experienced developers, bug fixing and technical debt reducing activities are a good entry point for new developers. Such tasks allow new developers to familiarize themselves with the software architecture... Following this strategy, they would be ready to be incorporated sooner in regular development project activities. Additionally, resourceful developers would have a greater chance to stand out sooner, reducing employee frustration…"

Moreover, having all changes and discussions publicly logged, as GitHub and the Merge/Pull model enable, would both serve to improve onboarding and mitigate the damage done when senior workers left:

> "This archive [of past design and implementation decisions] would form a useful knowledge base that can be used to lower the learning curve for newcomers and ground further decision-making for experienced developers. Moreover, this knowledge would be permanent and independent of key employees leaving a project."

Torkar et al. (2011) contrasted these OSS practices with existing practices at Ericcson, a global telecommunications company, which had typical problems acclimating new workers to their in-house software development methodology, "Streamline:" it took 38% of newcomers over a month just to acclimate to the in-house methodology, and a majority never graduated from the initial "software testing" tasks that Ericsson commonly assigned to newcomers as part of the onboarding process.[19]

OSS software projects also use simple initial tasks to build and assess competency, as Torkar et al. (2011) pointed out. In the words of one developer and OSS project founder surveyed and quoted by Kalliamvakou et al. (2015):

> "Even if you are not sure if the other dev[eloper] is capable of contributing good code, you can review pull requests and if the fifth pull request is good you give him/her commit bit [the power to make direct changes in the OSS repository, i.e. merge others' pull requests]."

Consistent with this, empirical studies find that initial OSS contributions are often more trivial tasks (Subramanian, 2020) and that a developer's track record with a project is the single most important predictor for time-to-merge in OSS projects on GitHub (Gousios et al., 2014).[20] In short, both OSS projects and the private sector use simple initial tasks to build

[19] Ericsson was not alone in this: in a series of workshops, Torkar et al. (2011) asked representatives from three large, multinational software companies to rank a list of their suggested benefits of adopting OSS methods by desirability, and found that both "Define an entry path for newcomers" and "Increase information availability and visibility" were consistently prioritized as the most important potential benefits.

[20] Secondary factors were project-specific or measures of pull request size and complexity. Gousios et al. (2014) also investigate the determinants of pull request *acceptance*, finding that almost all pull requests are eventually merged and suggest the number may be as high as 90% once one corrects for merges occurring outside of GitHub. They report that the single most relevant factor for eventual merging is whether the files

competence and evaluate the performance of newcomers, with private industry adopting OSS development practices in no small part because they facilitated this process. The increase in productivity we document, via a decrease in approval times, reflects this process by which juniors take on tasks, learn from seniors during code review, and eventually both write better code (which is merged faster) and build trust, allowing them to take on more serious tasks and have their changes merged with less scrutiny.

All this suggests that attention from seniors, both for education and evaluation, is critical for onboarding juniors.

touched have been modified recently (i.e. are relevant to ongoing development). We thus do not consider this as an outcome variable for the purposes to determining whether and when a developer achieves competency.

3.6.4 GitHub Data Appendix

This section describes how we worked with GHTorrent data to create the datasets necessary to estimate the main text's regressions. Before doing so, we first discuss the structure of GitHub's data, which Gousios (2013) accesses through GitHub's API to build GHTorrent.

In GitHub, every repository can be concieved of as a collection of events on a timeline that keeps track of changes to the repository's files and the discussion surrounding those changes. Figure 3.5 presents an example of workflow on a GitHub project being developed using pull requests. Each grey circle represents a discrete event, or action, taken on the part of some user. These events are timestamped and stored as JSON files with associated characteristics that depend on the action in question; see Gousios and Spinellis (2012) for a more technical description. These files can be accessed directly through GitHub's API, subject to request limits, and some researchers, studying a small number of repositories, obtain data from GitHub's API directly.

To obtain data on all public repositories, GHTorrent surmounts the request limit by crowdsourcing API keys from multiple donors and processes the raw, timestamped event data into a MySQL database; see Figure 1 in Gousios (2013) for the database schema. The result is a "snapshot" of the timeline for all public repositories on GitHub, taken at different points in time.[21] Various snapshots are currently publically available on Google BigQuery, which allows a user to query GHTorrent's MySQL dataset using SQL. The main text uses the 2019Q2 snapshot, and Appendix 3.6.5.4 replicates our results on the 2016Q3 snapshot as a robustness check.[22] This broad coverage, ease of access, and preprocessing explains the popularity of GHTorrent with researchers (Cosentino et al., 2016).[23]

The first step in processing the GHTorrent data is to merge data on pull request comments and actions (open, merge, etc.) to create a table of pull request activity on each repository. We then drop all repositories that have less than a specified number of "merge" events: 100 in the main text.[24] This effectively keeps repositories that are actually using GitHub to jointly develop code; as discussed in the main text, repositories that are actually personal projects or websites generally do not use pull requests, and some repositories on GitHub are "mirrors" of projects actually being developed elsewhere, and thus may have only open and closed pull

[21]For example, the GHTorrent sample, or "snapshot" used in the main text of this paper is from 2019, so that a project that began in 2008 but was deleted in 2018 would not appear in that dataset. However, the project would appear in the 2016 sample used in Appendix 3.6.5.4, and we could see the timeline of events from 2008 to 2016 there.

[22]For this project, we accessed the GHTorrent dataset through BigQuery's API in Python. To access the 2016 snapshot, the project id is "ghtorrent-bq" and the 2016 vintage has the dataset id "ght". To access the 2019 snapshot, the dataset id is "ghtorrentmysql1906" and the project id is "MySQL1906". For more documentation, see `https://github.com/ghtorrent/ghtorrent.org`.

[23]GH Archive (`https://www.gharchive.org/`) also records this public GitHub timeline data. For a comparison of the costs and benefits of various methods of accessing GitHub's data, see Mombach (2019).

[24]An early version of this paper used 120 merge events as the threshold, which corresponded to keeping exactly 25% of all pull requests from the largest projects in the 2016 GHTorrent snapshot. We have tried using both 120 and 200 merge events as the threshold in the 2019 snapshot and verified that the results in the main text and appendices are not sensitive to the precise choice.

requests, but no merge events.[25] Effectively, by conditioning on projects with many merge events, we select only projects that are being *developed* on GitHub.

We then drop all events initiated by bots using regex filters on logins, following Wyrich et al. (2021). We also drop events initiated by organizational accounts (that stand in for groups of users) and events for accounts that can't be linked to a user ("fake" accounts in GHTorrent parlance; these are real users who have not configured a GitHub account and whose commit activity is tracked via email address in GHTorrent). Finally, we drop events where the user is simply missing. These are minor issues. Collectively, this means dropping only 2.8% of the PR activity data, and the vast majority of this is "bot" activity (2.1% of all activity) .

There are some errors in GHTorrent: for example, we also drop duplicate events. However, having made all of the above cuts, there are a few instances of multiple recorded "open" or "merge" events for the same pull request at different times. We resolve this by keeping the earliest open event and last merge event, but these constitute a tiny fraction of the data (a few hundred events collectively in a final dataset of over 65 million events). As a last step, we remove some "test" repositories that have the word "test" in the repository name and an average time to merge measured in seconds (less than one minute) though this only removes 16 repositories.

Having made these cuts, we are left with a dataset of over 65 million pull request events on 36,537 large repositories that we consider joint attempts to develop code, which we now call *projects*.[26] We can use this dataset in turn to create two different panel datasets that are used to produce the figures in the main text.

First, we build a panel dataset on individual contributions (merged pull requests) from the events in this dataset. The goal is to estimate equation (1.4) in the main text, reproduced here: letting $y_{i,p,t}$ be either the approval time or total comments received for a contribution opened by user i on project p at time t, we estimate the following via OLS:

$$\begin{aligned} y_{i,p,t} = &\sum_{j=1}^{13} D(\text{Months Project Experience} = j)_{i,p,t} \\ &+ \sum_{k} D(\text{Months Programming Experience} = k)_{i,t} \\ &+ D_{i,p} + \beta_{PA,p} ProjectAge_{p,t} + \epsilon_{i,p,t}. \end{aligned}$$

The first sum consists of dummy variables for having between one and thirteen or more months of experience on project p at time t, and the second sum consists of dummy variables for overall programming experience measured by GitHub account age at time t, which is capped at 49 months, so that the dummy for having exactly 49 months includes all individuals with more than four years (48 months) of experience on GitHub measured using account

[25] Missing merge *events* in GitHub for pull requests that are actually merged is very common—see Gousios et al. (2014) for a discussion.

[26] We thus treat forked repositories as separate projects, if they have enough merge events; see Section 1.3.1.

Figure 3.5: Workflow on GitHub: Simplified Example

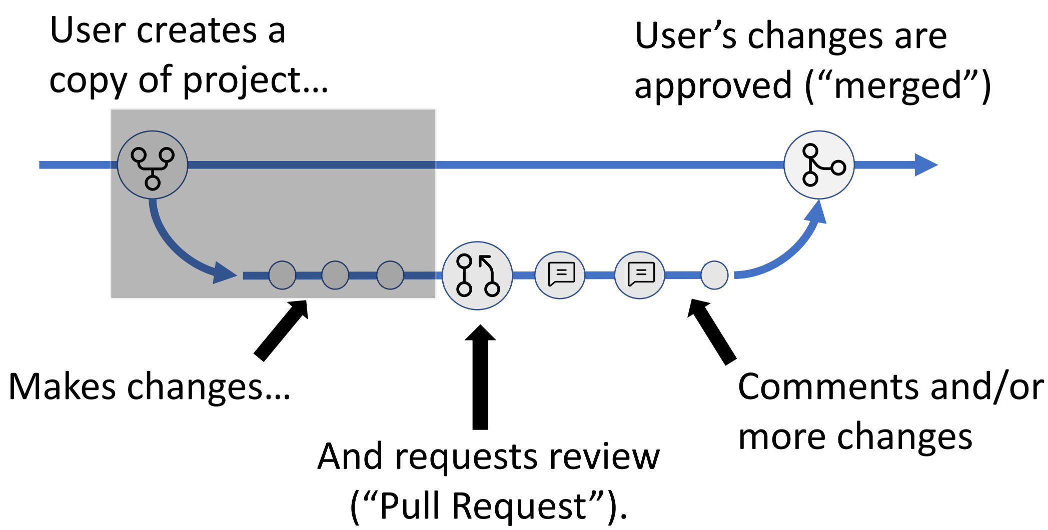

Notes: A GitHub repository can be concieved of as a collection of events on a timeline tracking changes to the repository's files and the discussion surrounding those changes. Each grey circle represents a discrete event, or action, taken by some user. These events are timestamped and stored as JSON files with additional information depending on the type of event. In this example, a user creates a copy of a repository, and then makes changes (grouped into individual "commit" events by the user) to the copy before opening a pull request to have their changes merged into the original. The grey box over this part of the workflow reflects that we may not observe this part of the contribution process: if the repository is copied as a new repository on GitHub, for example, then we will observe these changes. But if the project is e.g. downloaded by the user and then changed locally, we may not be able to observe this activity.

age.[27] We also allow for individual-by-project fixed effects ($D_{i,p}$) and project-specific linear time trends ($\beta_{PA,p} ProjectAge_{p,t}$). Table 3.5 presents example of this dataset, with fictitious data, used to estimate this regression. Recall that Figure 1.5 in the main text uses the marginal effects estimated from equation (1.4) to show that approval time and the number of comments recieve fall precipitously in the first six months of project-specific experience. This provides evidence of a nontrivial onboarding period for juniors, as discussed in the main text.

Table 3.5: Panel Data on Individual Contributions (Example)

n	Opened t	Contributor i	Project p	Project Experience	Approval Time	Total Comments
1	1/1/15	Jake	Project A	0 days	6 days	4
2	1/9/15	Jake	Project A	8 days	5 days	2
3	1/9/15	Jake	Project B	0 days	17 days	5
4	1/12/15	Jake	Project B	3 days	15 days	3
5	1/12/15	Justin	Project B	0 days	1 days	0

Notes: An illustrative example of the panel dataset used to estimate regression (1.4) in the main text. The actual dataset has N = 10,881,355 observations (approved contributions) on 36,537 large projects (with at least 100 approved contributions).

Next, we build a panel dataset of juniors joining new projects that is used to estimate equation (1.5) in the main text, yielding the estimate for ρ plotted in Figure 1.8. We begin by identifying junior J type and senior S type workers using their pull request activity. In each calendar month t and each project p, we assign each user with activity on at least one pull request in p at t into either category J or category S. We drop users who never contribute, and restrict attention to those who open at least one pull request that is merged. A J type transitions to an S type on a particular project either when they have reached a tenure of six months on that project, or when we observe them reviewing code (i.e., when we observe them merging/closing/commenting on pull requests opened by others). Project tenure is measured as the length of time between a user's first observed activity and their last observed activity on a project. This definition implies that some workers are S types from their first month on a project. Also note that once a J type worker transitions on a project, they are counted as an S type in any calendar month when they "re-appear" on that project in the pull request data.

We define the quantity of J types on project p at time t as $J_{p,t}$, tabulated as the number of users who have contributed to that project (i.e. authored at least one pull request that was eventually merged) at time t with less than six months of tenure and who have not reviewed code written by others (i.e. who have not been observed merging/closing/commenting on a

[27] A small number of user accounts and projects have pull request activity predating the reported date of account or project creation. We correct these using the time of earliest observed pull request activity.

pull request opened by someone else). The other active users are summed into $S_{p,t}$. Table 3.6 presents an example of this dataset, with fictitious data, used to estimate equation (1.5), reproduced here: letting $\mathbf{1}(i \text{ joining } p \text{ at } t \text{ onboards})$ denote an indicator function for whether a Junior i on project p (counted in the sum $J_{p,t}$) will eventually transition to being an S type on project p, we estimate:

$$\begin{aligned}\mathbf{1}(i \text{ joining } p \text{ at } t \text{ onboards}) &= \sum_b D\left(\frac{J_{p,t}}{S_{p,t}} \text{ in bin } b\right) \\ &+ D_p + \beta_{PA,p} ProjectAge_{p,t} + X_t + \gamma_{i,t} + \epsilon_{i,p,t}.\end{aligned}$$

We estimate the effect of $J_{p,t}/S_{p,t}$ non-parametrically via OLS by measuring it as a set of dummy variables representing equidistant bins for junior-senior ratios. Project specific dummies D_p control for unobservable project-specific features that may make some projects easier to join, while $ProjectAge$ is a project-specific time trend meant to capture the project life cycle, since some projects may become harder to join as they age; X_t are year fixed effects, and $\gamma_{i,t}$ captures Junior-specific controls, which include fixed effects for the year user i's account was created, and dummies for the age of user i's account.

We cannot include user-project specific fixed effects here because they are collinear with the outcome variable (we only observe one outcome per project for each individual: either they onboard, or they do not). Relatedly, we cannot well-estimate individual fixed effects because in practice most individuals join very few OSS projects in sample over time. Note if someone joins only one project in our sample of large OSS projects, we cannot estimate a fixed effect for them. Appendix 3.6.5.2 discusses this and shows that our results are qualitatively unchanged by adding individual fixed effects, though the sample size shrinks.

We linearly approximate the non-parametric estimates of congestion in Figure 1.8 using OLS for use in calibrating our DSGE model. This is the blue line in that figure. Note that the precise slope is not too sensitive to the choice of threshold for the number of pull request merge events we include in our sample: across the thresholds of 100, 120, and 200 that we have explored, the slope of the line varies from -0.064, -0.067, and -0.075, respectively while the intercept is stable at 0.47; we thus pick a slope of -0.07 for calibrating our DSGE model.

Table 3.6: Panel Data on Juniors Joining Projects (Example)

n	Month Joined t	Junior i	Project p	Cohort Size $J_{p,t}$	$S_{p,t}$	$\frac{J_{p,t}}{S_{p,t}}$	Onboards?
1	1/2015	Jake	Project A	2	4	.5	Y
2	1/2015	Justin	Project A	2	4	.5	N
3	1/2015	Justin	Project B	1	10	.1	Y

Notes: An illustrative example of the panel dataset used to estimate regression (1.5) in the main text. The actual dataset has N =1,044,039 observations, one for every junior on each project joined.

3.6.5 Robustness of Empirical Analysis with GitHub Data

3.6.5.1 Congestion Results with More Bins

Figure 3.6 extends the congestion analysis of Figure 1.8 to allow for J/S ratios greater than 2. This results in a slightly flatter calibration for our linear onboarding function, chosen to approximate the nonlinear relationship apparent in Figure 3.6: $\rho = .45 - .04\left(\frac{j_t}{s_{t-1}}\right)$. Figure 3.7 replicates the IRFs in Figure 1.10 using this new, flatter calibration to show that this implies slightly less stickiness relative to the benchmark model with convex investment adjustment costs. Put another way, the results are not terribly sensitive to the ρ calibration, provided there is sufficiently negative slope.

We prefer our headline calibration because we view it as a local linear approximation of the nonlinear function over a space which is more relevant for optimizing firms in our model: the steady-state ratio of J/S is generally much, much less than one, reflecting the fact that at any given point in time most workers are already onboarded in the investment sector.

Figure 3.6: Non-parametric Estimate of the Onboarding Function ρ

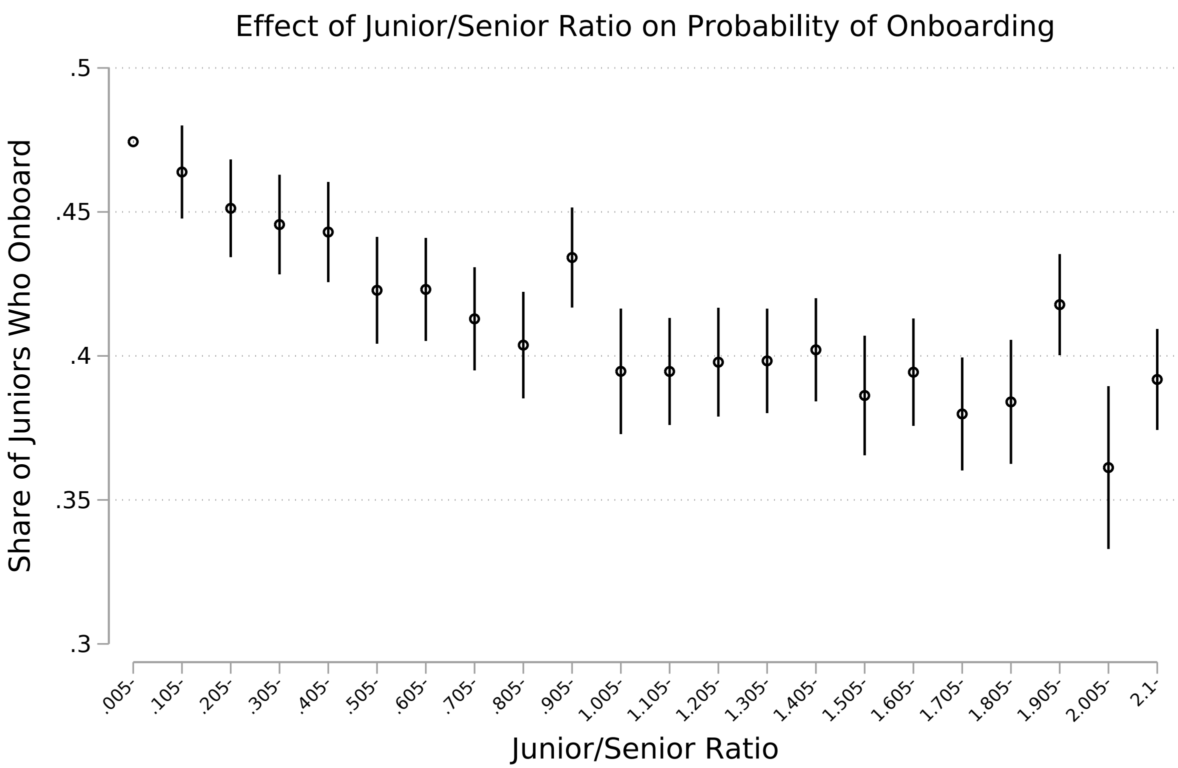

Notes: Estimates from: $\mathbf{1}(i \text{ joining } p \text{ at } t \text{ onboards}) = \sum_b D\left(\frac{J_{p,t}}{S_{p,t}} \text{ in bin} = b\right) + D_p + \beta_{PA,p} ProjectAge_{p,t} + X_t + \gamma_{i,t} + \epsilon_{i,p,t}$. The "spikes" in bins containing $J/S = 1$ or $J/S = 2$ reflects the fact that being "one-on-one" or "two-on-one" with an incumbent worker is particularly helpful for successful onboarding. Note over 75% of all project-month observations have $J/S \leq 1$ and over 90% have $J/S \leq 2$.

Figure 3.7: Model Responses to a Contractionary Monetary Policy Shock: Flatter ρ

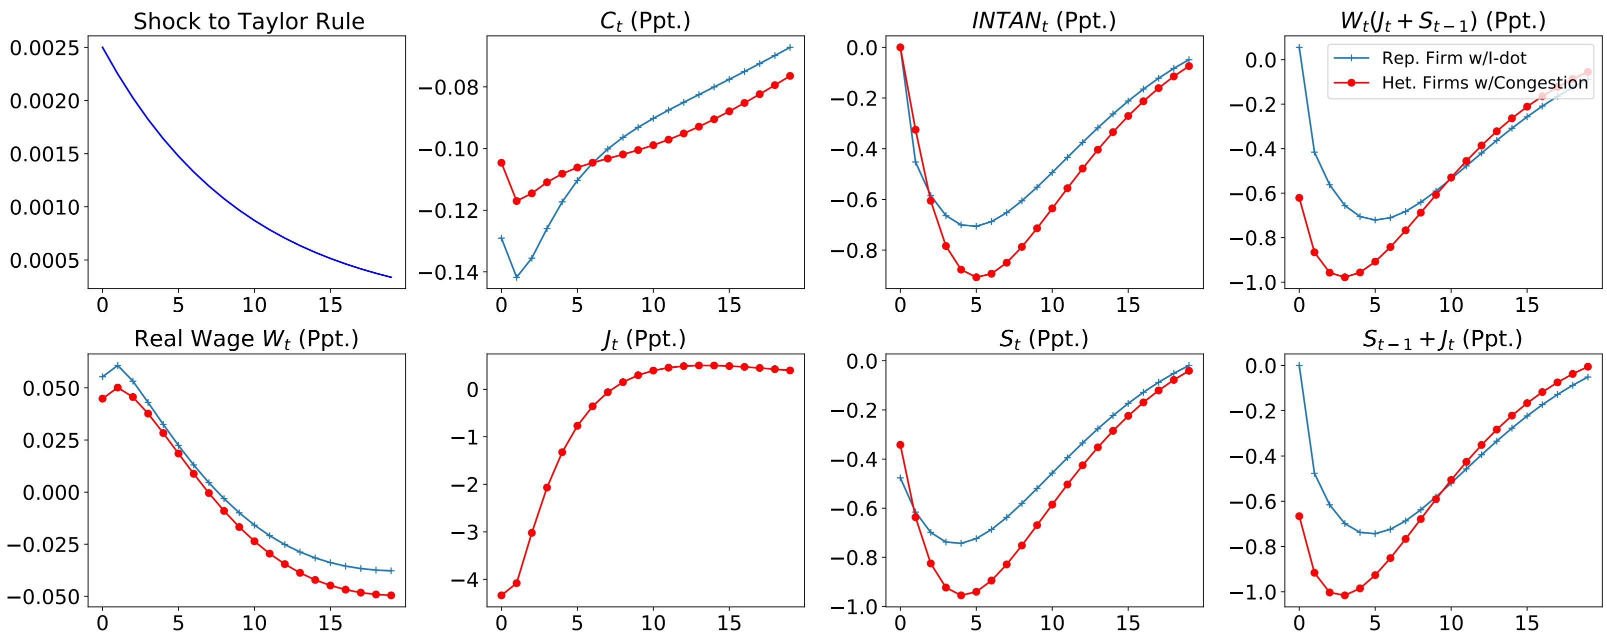

Notes: Quarterly impulse response functions in the model with flatter ρ calibrated to capture the nonlinear relationship in Figure 3.6: $\rho = .45 - .04\left(\frac{j_t}{s_{t-1}}\right)$, compared to a simple representative firm model with convex investment adjustment costs. All other parameters are the same as described in Section 1.4.

3.6.5.2 Adding Individual Fixed Effects to the Congestion Regression

Adding individual fixed effects to the congestion regressions in the main text shrinks the sample size by about a third. This is because we can only identify individual fixed effects for contributors who successfully merge pull requests on multiple large open source projects (recall we only consider projects with many pull requests; see Section 3.8.6). Including fixed effects drops individuals who e.g. only contribute and join one major open source project. We include fixed effects here and note that in spite of the diminished sample size (N = 646,843 as opposed to N =1,044,39 in the headline results in Figure 1.8), the results are qualitatively similar. See Figure 3.8.

Figure 3.8: Non-parametric Estimate of the Onboarding Function ρ with Individual Contributor Fixed Effects

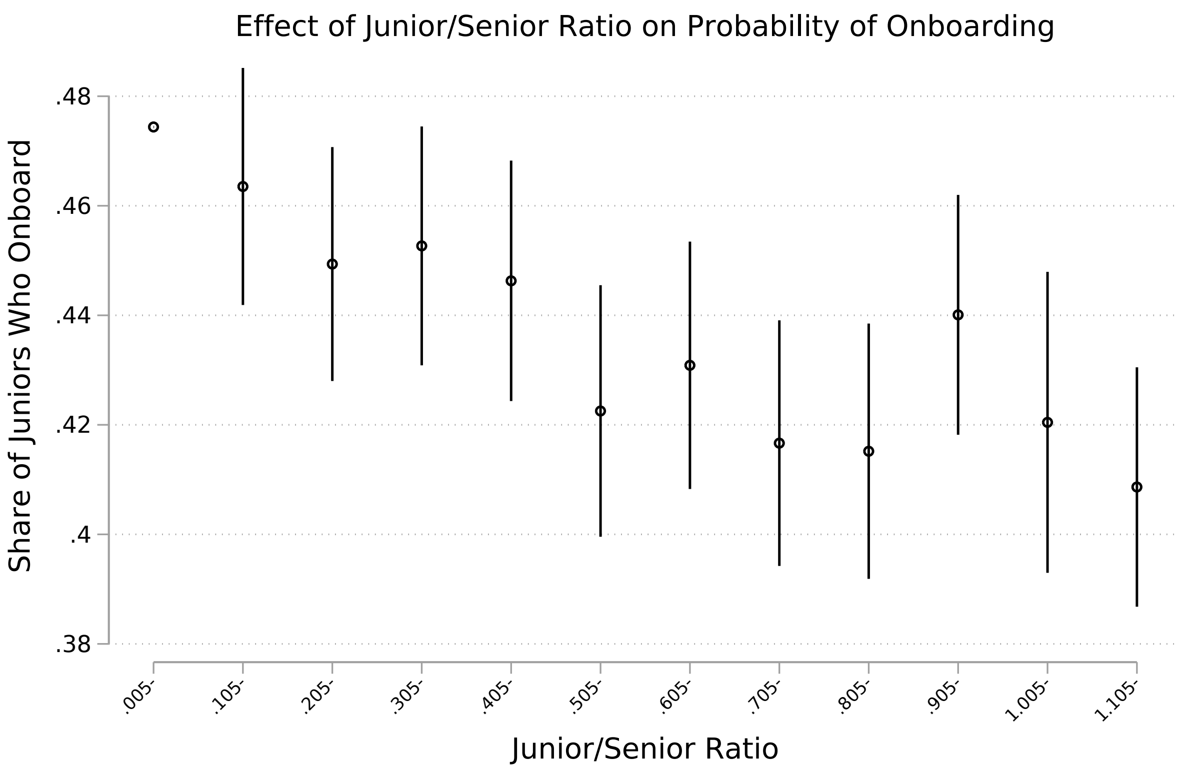

Notes: Estimates from: $\mathbf{1}(i \text{ joining } p \text{ at } t \text{ onboards}) = \sum_b D\left(\frac{J_{p,t}}{S_{p,t}} \text{ in bin} = b\right) + D_p + \beta_{PA,p} ProjectAge_{p,t} + X_t + \gamma_i + \epsilon_{i,p,t}$. The "spikes" in bins containing $J/S = 1$ or $J/S = 2$ reflects the fact that being "one-on-one" or "two-on-one" with an incumbent worker is particularly helpful for successful onboarding. Note over 75% of all project-month observations have $J/S \leq 1$ and over 90% have $J/S \leq 2$.

3.6.5.3 Congestion Results for Narrower Definitions of Onboarding

We count juniors as "onboarding" successfully in the main text using two observables:

- A junior goes on to remain with the project at least six months
- A junior eventually begins commenting on/merging/closing pull requests opened by others (i.e. code review).

If any junior eventually satisfies one of these two conditions, they get counted as "onboarded." As discussed, many do not and leave within a month of joining the project and without reviewing anyone else's code. Our main specification uses both of these definitions, but using just one or the other yields qualitatively similar results: i.e., an "onboarding function" that looks downward sloping, consistent with congestion.

Figure 3.9: Non-parametric Estimate of the Onboarding Function with Onboarding Success Determined by Junior's Activity Only

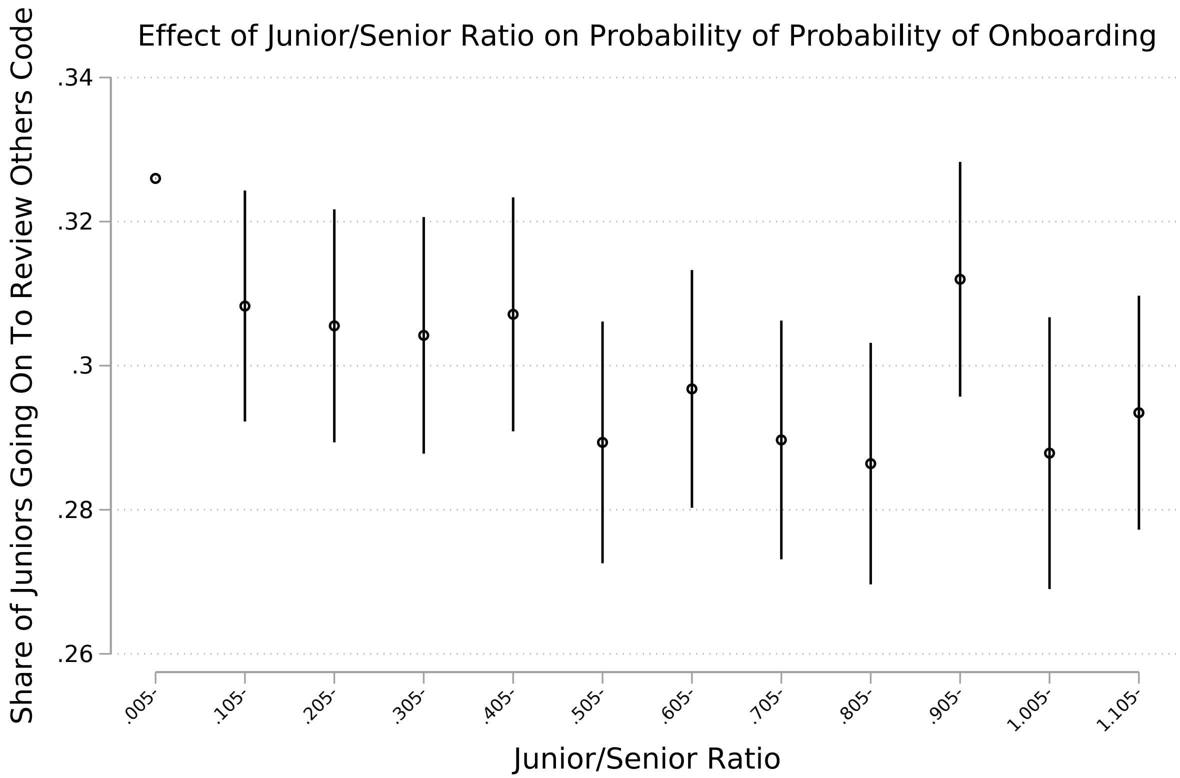

Notes: Estimates from: $\mathbf{1}(i \text{ joining } p \text{ at } t \text{ onboards}) = \sum_b D\left(\frac{J_{p,t}}{S_{p,t}} \text{ in bin} = b\right) + D_p + \beta_{PA,p} ProjectAge_{p,t} + X_t + \gamma_{i,t} + \epsilon_{i,p,t}$. The "spike" in the bin which contains $J/S = 1$ reflects the fact that being "one-on-one" with an incumbent worker is particularly helpful for successful onboarding. Note over 75% of all project-month observations have $J/S \leq 1$.

Figure 3.10: Non-parametric Estimate of the Onboarding Function with Onboarding Success Determined by Eventual Project Tenure Only

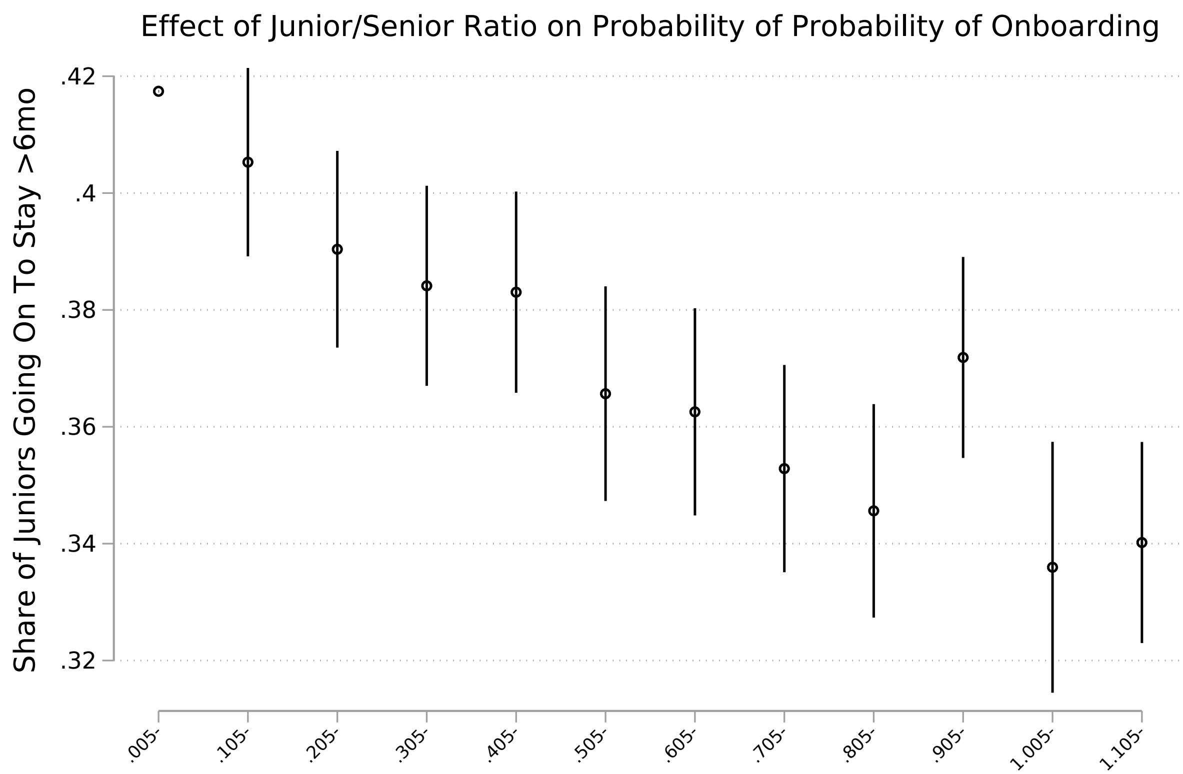

Notes: Estimates from: $\mathbf{1}(i \text{ joining } p \text{ at } t \text{ onboards}) = \sum_b D\left(\frac{J_{p,t}}{S_{p,t}} \text{ in bin} = b\right) + D_p + \beta_{PA,p} ProjectAge_{p,t} + X_t + \gamma_{i,t} + \epsilon_{i,p,t}$. The "spike" in the bin which contains $J/S = 1$ reflects the fact that being "one-on-one" with an incumbent worker is particularly helpful for successful onboarding. Note over 75% of all project-month observations have $J/S \leq 1$.

3.6.5.4 Replicating Main Results on an Earlier GitHub "Snapshot"

The results in the main body of the paper use a snapshot of GitHub from provided by Gousios (2013) on Google BigQuery from June 2019. This section demonstrates that the main results of this paper are qualitatively robust to using an early snapshot: the earliest available on Google BigQuery from September 2016. This predates the acquisition in 2018 by Microsoft, the addition of new features and changes to the API, etc. In short, this section demonstrates that the results of the paper are not sensitive to the "vintage" of data used in the analysis. Note that due to exponential growth in the popularity of GitHub, these three years of data make a large difference: the main text's 2019 sample is almost an order of magnitude larger, which partly explains why results for e.g. pull request comments and approval times in Figure 3.11 are noisier here.

The sample here is otherwise identical to the sample in the main paper, except for the fact that here we kept repositories with at least 120 merged pull request events, instead of 100 as in the main text, which was originally chosen to keep exactly 25% of all pull requests from the largest projects in the 2016 GHTorrent snapshot. The results in the main text and appendices using the 2019 snapshot are robust to using this threshold of 120, as well as higher thresholds (we checked results which kept only repositories with at least 200 merge events as well). See Appendix 3.6.4 (and footnote 24) for details.

Figure 3.11: Becoming Productive Requires Onboarding (Results with 2016 Data)

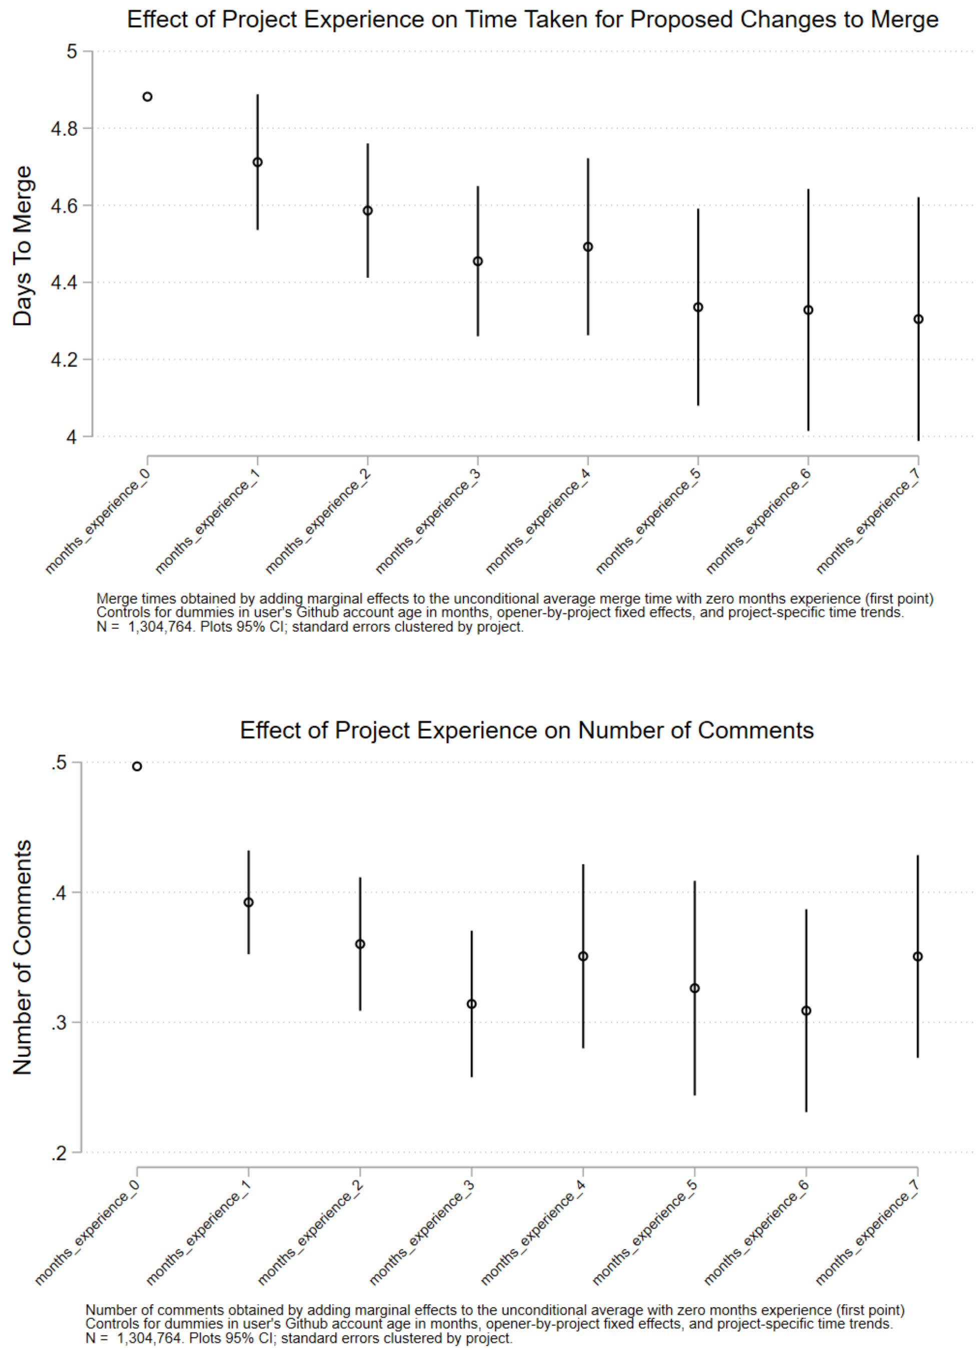

Notes: Over time, new contributors' proposed changes are approved faster, with less discussion. Estimates from $y_{i,p,t} = \sum_{j=1}^{13} D(\text{Months Experience} = j)_{i,p,t} + \sum_k D(\text{Months Ind. Exp.} = k)_{i,t} + D_{i,p} + \beta_{PA,p} ProjectAge_{p,t} + \epsilon_{i,p,t}$ where $y_{i,p,t}$ is either the total time to merge the proposed change in days or number of total number of comments during code review.

Figure 3.12: Most Work is Done by Experienced Team Members (2016Q3 Sample)

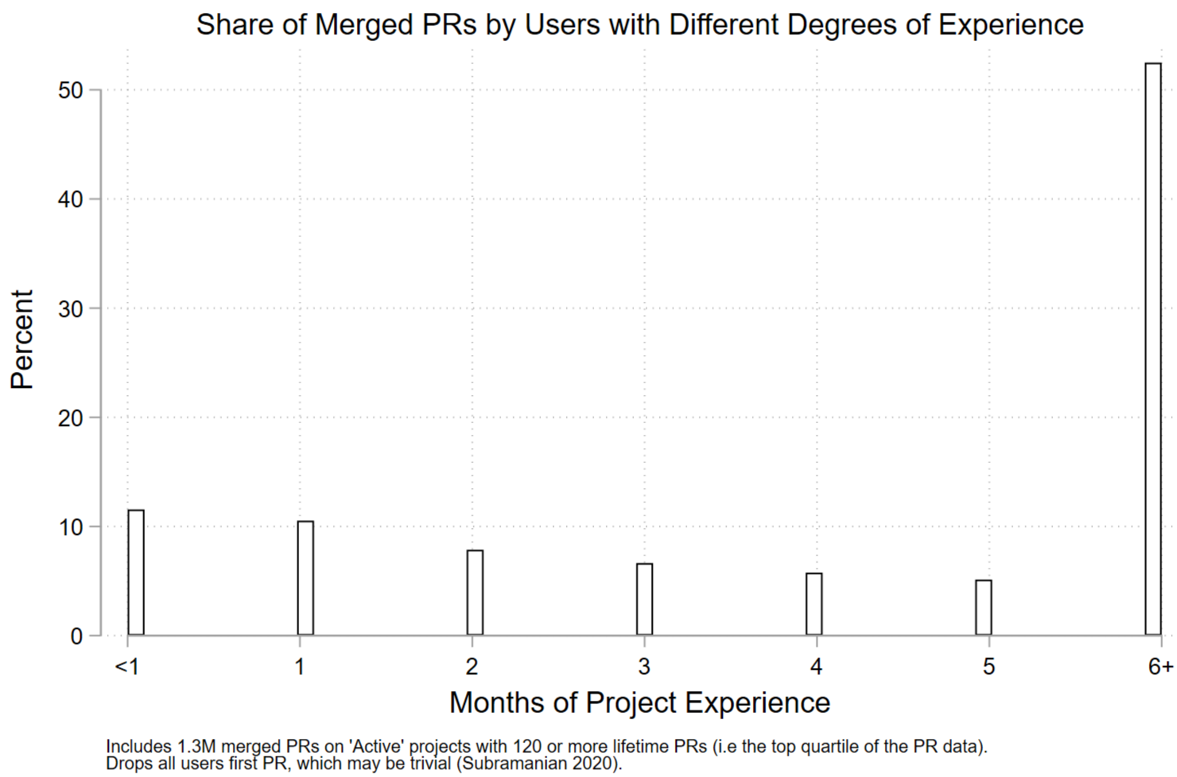

Notes: this figure plots the share of all merged, non-bot pull requests that are opened by users with different degrees of experience, showing that most work is done by those who have at least six months of project-specific experience. We exclude each user's first pull request, given evidence by Subramanian (2020) that these are more often trivial changes. Since we do not otherwise control for complexity or importance of the various tasks completed by these pull requests, and given that longer-tenure workers take on more complex and important tasks, this figure likely understates the importance of work done by senior workers. Source: GHTorrent.

Figure 3.13: Distribution of Programming Languages (2016Q3 Sample)

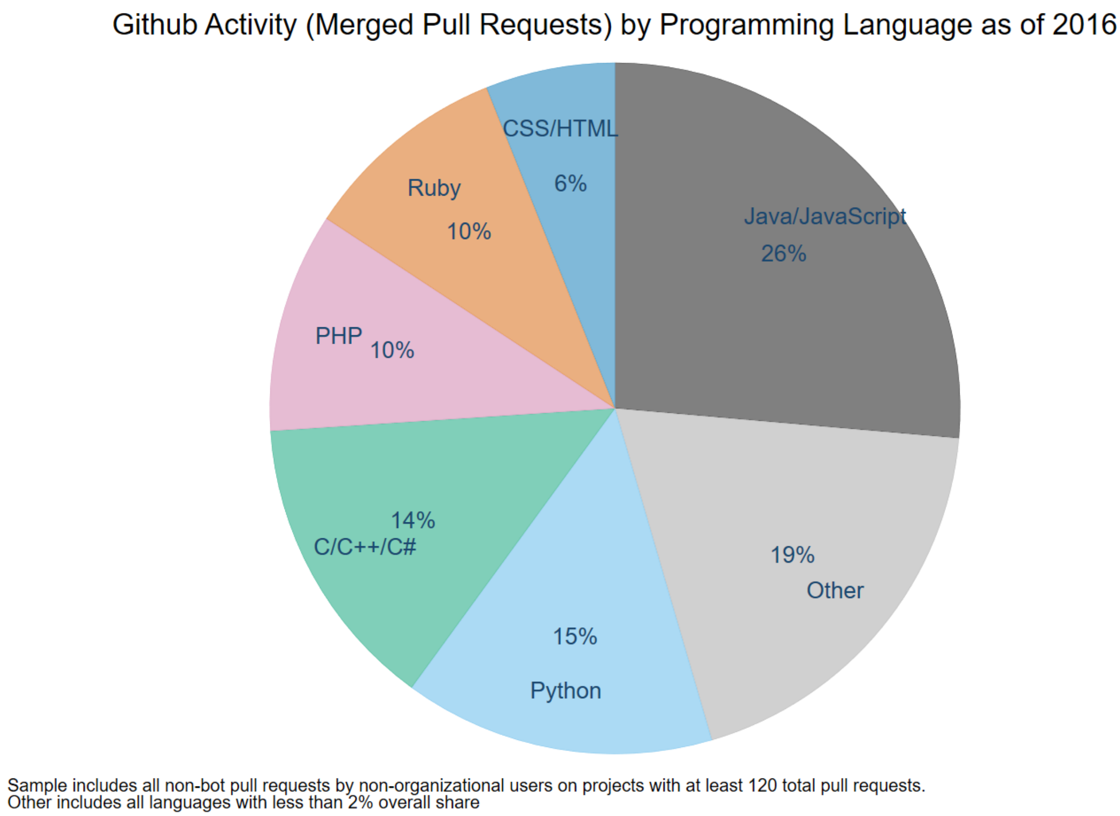

Notes: This does not weight each repository by size, other than dropping all small repositories with less than 120 total merged pull requests. Other includes languages like R, Matlab, and others which are a very small share of the projects in our sample. Source: GHTorrent.

Figure 3.14: Newcomers Either Contribute Once, or Stay a Long Time (2016Q3 Sample)

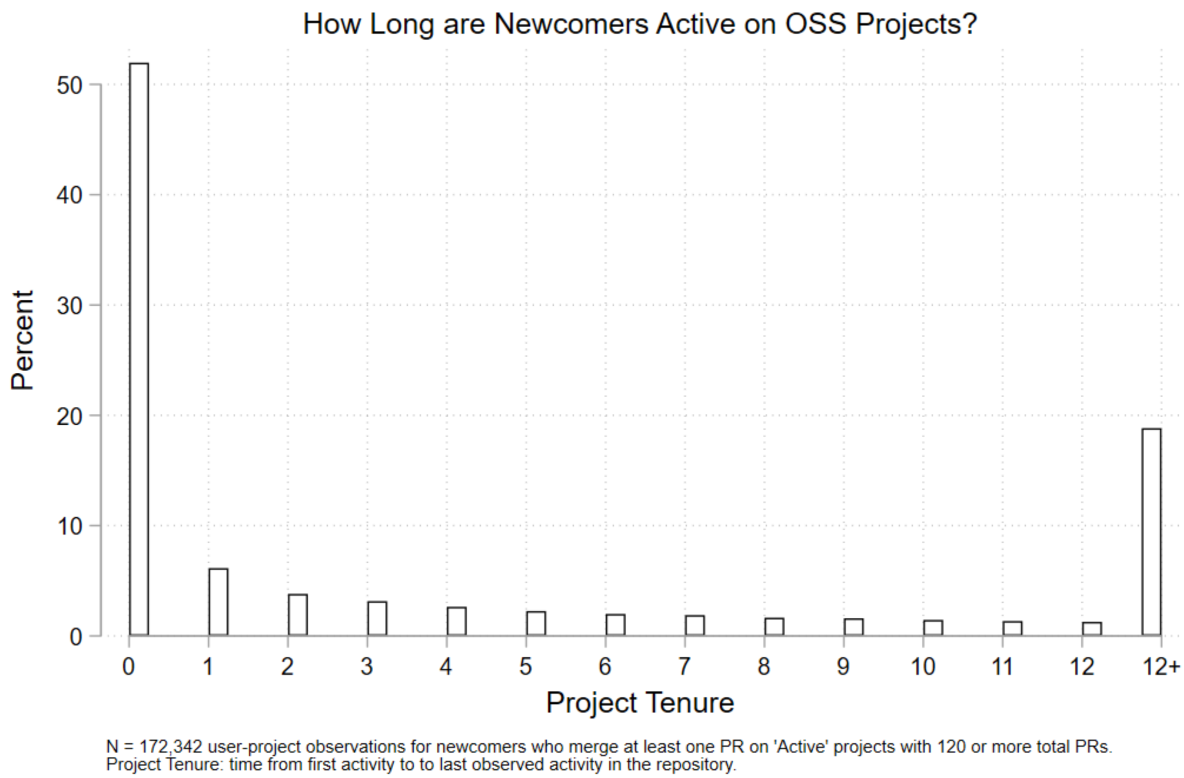

Notes: this figure plots the share of all newcomers J (non-bot users who join a project and successfully contribute at least one PR) by their subsequent observed tenure. Most newcomers will go on to have very short tenure (rounded to the nearest month) and contribute once, followed by a nontrivial second group who remain much longer. Source: GHTorrent.

Figure 3.15: Onboarding Requires Attention from Senior Workers (2016Q3 Sample)

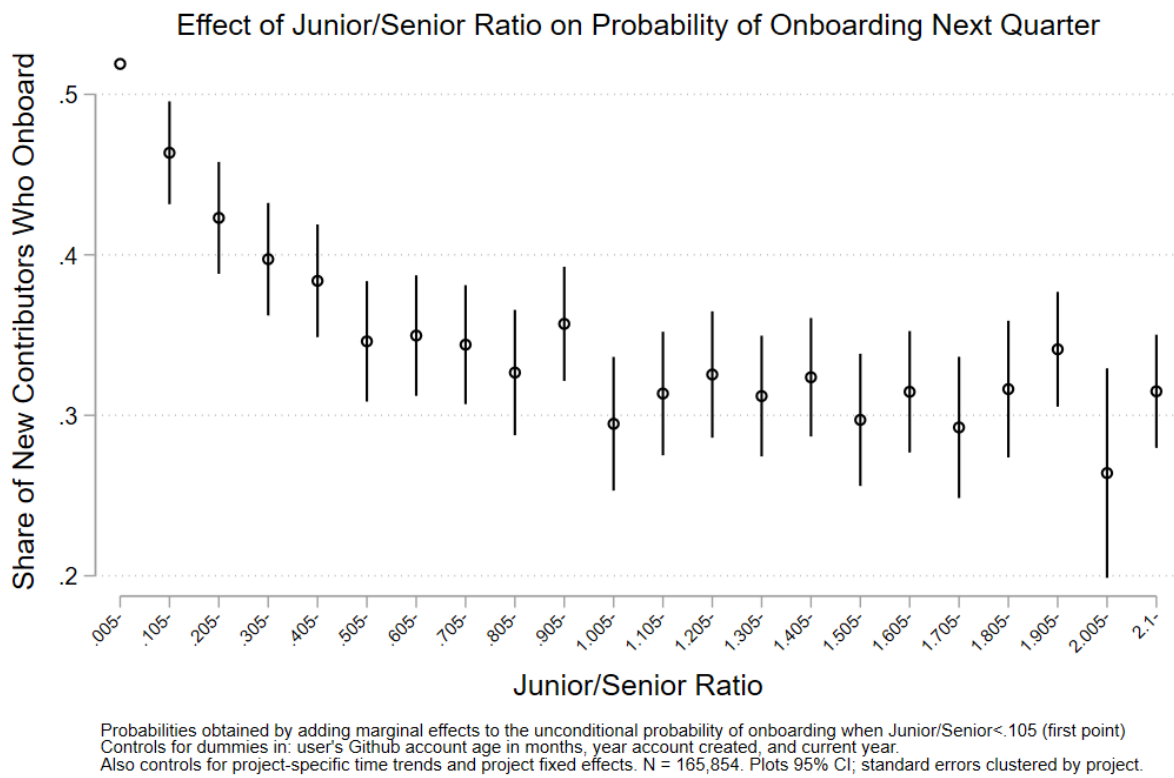

Notes: Over 75% of all project-month observations have $J/S \leq 1$. Estimates from:
$\mathbf{1}(i \text{ joining } p \text{ at } t \text{ onboards}) = \sum_b D\left(\frac{J_{p,t}}{S_{p,t}} \text{ in bin } b\right) + D_p + \beta_{PA,p} ProjectAge_{p,t} + X_t + \gamma_{i,t} + \epsilon_{i,p,t}$

3.7 Appendix to Chapter 2

3.7.1 Additional Figures

Figure 3.7.1: Compositional Shifts among Investment Components, 1960-2020

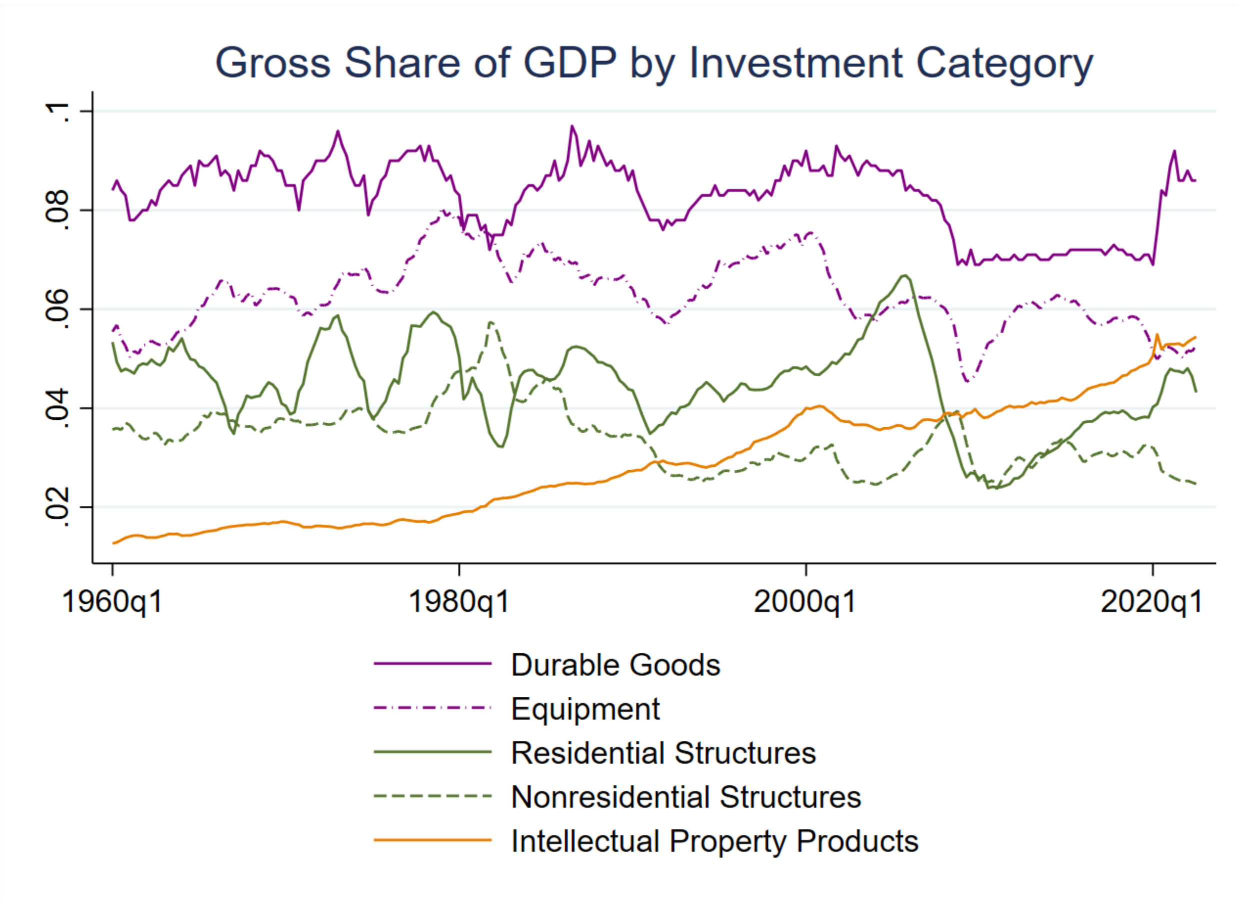

Figure 3.7.2: The Diminishing Effect of a 1% Federal Funds Rate Hike Over Time?

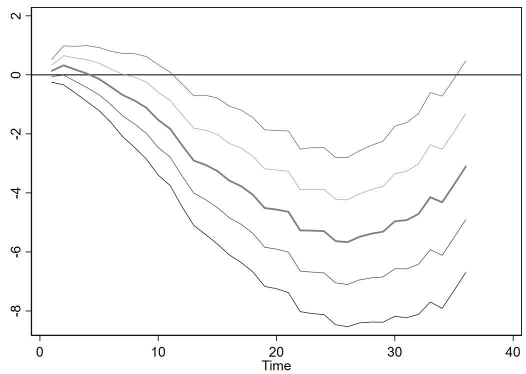

(a) Output Response Pre-1983

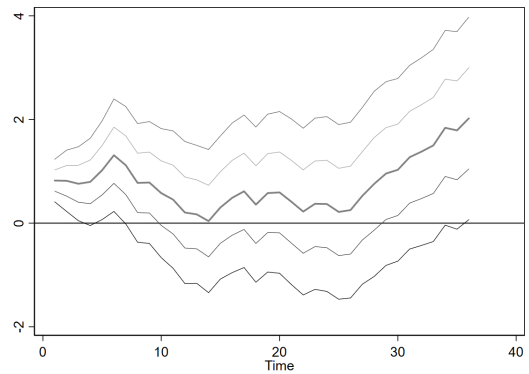

(b) Output Response Post-1983

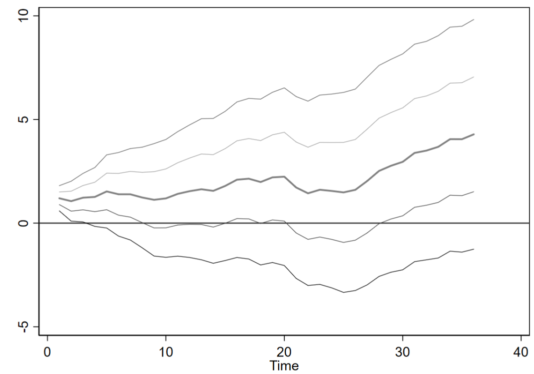

(c) Output Response Post-1983 (HF Shocks)

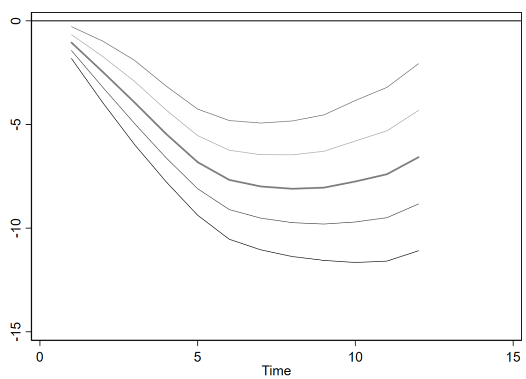

(d) Investment Response Pre-1983

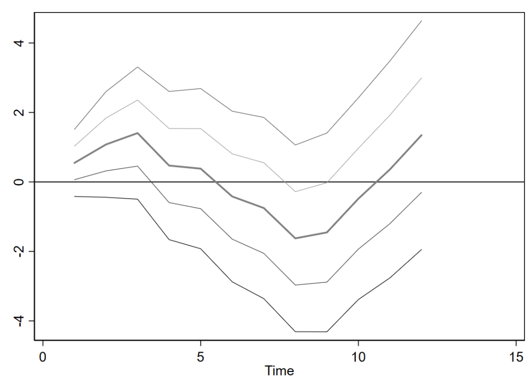

(e) Investment Response Post-1983

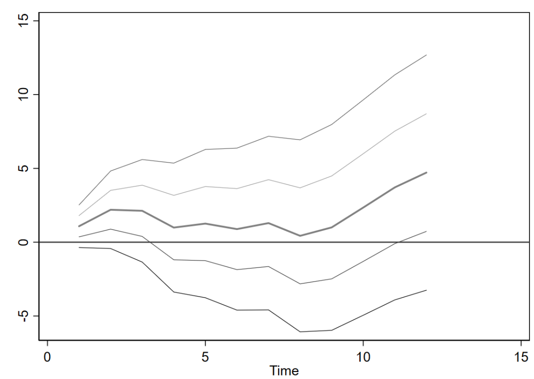

(f) Investment Response Post-1983 (HF Shocks)

Notes: The effect of a 1% hike in the federal funds rate on real industrial production (top row) and real gross private domestic fixed investment (bottom row) estimated over sample periods before and after 1983. "Time" is the number of periods after the shock, either in months (top row) or quarters (bottom row) depending on the frequency of the underlying data. Columns one and two use the Wieland and Yang extension of the Romer and Romer, and column three uses the high frequency shocks in Nakamura and Steinsson (2018) when they are available as a robustness check. One and two standard error bands are plotted, calculated as in Romer and Romer (2004).

3.7.2 Calculation of Domestic Labor Content

The computation of the domestic labor content begins with the make and use tables. Using the row-by-column convention, the use table is a commodity-by-industry table, where each column states the quantities of commodity intermediates used in a given industry's production. The make table is an industry-by-commodity table, where each column shows how the production of a commodity is distributed across industries. Dividing each column of the use table by industry gross output gives the NxN direct proportions matrix B, where N is both the number of industries and commodity types. Dividing each column in the make table the sum of each column gives the NxN make-shares matrix W. Defining $\mathbb{B} = BW$, which creates a commodity-by-commodity matrix, we compute the commodity-by-commodity Leontief inverse $(I - \mathbb{B})^{-1}$. Multiplying the Leontief inverse on a column vector of commodity demand for final use ε, the term $(I - \mathbb{B})^{-1}\varepsilon$ gives the vector of values the gross commodity output demanded to produce ε commodities for final delivery.

To account for imports, we compute the "total domestic requirements matrix" as defined by the BEA. This requires two steps in addition to the process outlined above to account for imports. First, the BEA provides total quantities of imports for each commodity type. This allows us to calculate an "import ratio" γ_j for each commodity type j, which is the fraction of imports over total domestic supply. Then, because the BEA does not have the data sources to identify how commodities are used across industries as intermediates and as final purchases, we assume imports represent a constant fraction of any use of commodity j. Then, before taking the Leontief inverse, we multiply each element of the row of the direct proportions matrix corresponding to commodity j by $1 - \gamma_j$. Then when computing the total commodity demands, each element of demand vector ε is also multiplied by the corresponding value of $1 - \gamma_j$. Taking the commodity "computers" as an example, imported computers are sometimes used as final purchases (e.g. consumption or investment) and sometimes used as an intermediate good in the production of other commodities.

Let θ_k^L be the industry k's labor share of gross output and let $\bar{\theta}^L$ be a row vector of the industry labor shares of gross output. Let ε^i be the column vector of commodity demand shares for final use i, where an element ε_j^i expresses the domestic expenditure demanded of commodity j by a dollar of demand for final use i, $\sum_j \varepsilon_j^i = 1$. Let γ be the vector of import ratios of commodities, with $D(\gamma)$ have the elements of γ on the diagonal with off-diagonal elements equaling 0. Then the domestic labor content is measured as $dlc^i = W(I - \mathbb{B})^{-1}D(\gamma)\varepsilon^i\bar{\theta}^L$, which is a scalar. From section 2.3.4, $W(I - \mathbb{B})^{-1}D(\gamma)\varepsilon^i$ is a vector where the elements are the quantities of industry gross output demanded ω_{ik}.

We can also compute the share of each dollar spent on final use i paid domestically, which is the domestic share of expenditure. Let θ_k^v be the value added per gross output of industry k. Then let $\bar{\theta}^v$ be a row vector, of which θ_k^v is the k^{th} element. Then, the share of domestic expenditure $1 - m^i$ is equal to total domestic value added per dollar of final expenditure, computed as $(1 - m^i) = W(I - \mathbb{B})^{-1}D(\gamma)\varepsilon^i\bar{\theta}^v$. Lastly, we can back out the labor share of domestic production $(1 - \alpha^i) = dlc^i/(1 - m^i) =$ (domestic labor content)/(domestic share of expenditure).

3.7.3 Other Considerations

Fixed Costs and Mark-ups When estimating the domestic labor content, we also hypothesized that rising mark-ups may decrease the *marginal* domestic labor content - the amount of labor income generated by a marginal puchase of final demand component i. We used industry level estimates of fixed costs and mark-ups from De Ridder (2019), allowing marginal sales to not all be allocated to marginal factor payments and instead to non-labor profit. While the inclusion of fixed costs and markups lowers the measured of the domestic labor content, it minimally affects the change over time.

Re-imports One concern in using the input-output tables is the potential for domestic import demand to generate demand for US exports through global input-output linkages. For example, if demand for autos increases and the cars are primarily manufactured in Mexico but parts are supplied from the US, this may increase demand for labor in parts-supplying industries. However, investigations using world I-O tables from the World I-O Database reveal trivial effects of reimports on the estimates of the domestic labor income generated.

3.7.4 Decomposition of Hand-to-Mouth Consumption

First, consider log linearizing the definition of the real wage bill,

$$W_tN_t = W_tN_t^x + W_tN_t^i + W_tN_t^c,$$

around a steady state in which φ is chosen to normalize steady state $N = 1$,

$$\widehat{W_tN_t} = N^x\widehat{W_tN_t^x} + N^i\widehat{W_tN_t^i} + N^c\widehat{W_tN_t^c}$$

and using the fact that the wage bill is proportional to output in each sector:

$$\widehat{W_tN_t} = N^x\widehat{\left(\frac{P_t^xX_t}{P_t^c}\right)} + N^i\widehat{\left(\frac{P_t^kI_t}{P_t^c}\right)} + (1 - N^x - N^i)\widehat{C}_t$$

Finally, note that $\widehat{C}_t$ can be decomposed into Ricardian and hand-to-mouth consumption: working from the definition in equation (2.1), log linearize

$$\widehat{C}_t = \frac{\chi WN}{C}\widehat{W_tN_t} + \left(1 - \frac{\chi WN}{C}\right)\widehat{C}_{r,t}$$

Note that in the model, we have $\chi WN/C = \chi(1-\alpha_c)(1-m_c)/N^c$. Rewriting and eliminating the aggregate wage bill with $\widehat{W_tN_t} = \widehat{C_{k,t}}$:

$$\underbrace{\widehat{C_{k,t}}}_{\substack{\text{Hand-to-Mouth}\\ \text{Consumption}}} = \frac{N^x}{\Omega}\underbrace{\widehat{\left(\frac{P_t^xX_t}{P_t^c}\right)}}_{\text{Real Exports}} + \frac{N^i}{\Omega}\underbrace{\widehat{\left(\frac{P_t^kI_t}{P_t^c}\right)}}_{\text{Real Investment}} + \left(1 - \frac{N^x}{\Omega} - \frac{N^i}{\Omega}\right)\underbrace{\widehat{C_{r,t}}}_{\substack{\text{Ricardian}\\ \text{Consumption}}}$$

where $\Omega \equiv 1 - \chi(1-\alpha_c)(1-m_c)$.

3.8 Appendix to Chapter 3

3.8.1 Robustness Checks

This appendix reports three sets of robustness checks to the paper's main IV results: adding controls, using the lagged level of the exchange rate as an instrument instead of the signed, squared, lagged distance from target, and measuring intervention Q_t using dollar sales (instead of converting to pounds and scaling by UK M0).

3.8.1.1 Adding Controls to IV Regressions

A simple portfolio balance channel model, outlined in Section 3.8.4, suggests the following controls: let r_t (r_t^*) be the interest rate on riskless pound (dollar) bonds and let h be their maturity, then the exchange rate is given by

$$\ln e_t - \ln e_{t-1} = \beta_0 + \beta_1 Q_t + \beta_2 \Delta E_t \Big[\ln e_{t+h} \Big] + \beta_3 \Delta r_t + \beta_4 \Delta r_t^* + X_t + \mu_t, \tag{3.3}$$

where β_3 is negative, β_4 is positive, and the coefficient on the h-period forecast revision β_2 is unity in theory. In practice, we use futures markets to compute forecast revisions and changes in policy rates in both countries for r_t and r_t^*; the vector X_t includes other interest rate controls, in addition to two lags of the dependent variable, as well as day-of-week, month and year dummies.[28] Including controls will change our estimate for β_1 if the controls are correlated with intervention Q_t (i.e. if they belong in the true "policy rule" for intervention) and the included controls are correlated with exchange rate growth, as suggested by the model.

Table 3.7 demonstrates that regardless of the sample period, our IV approach flips the sign of the OLS regression and yields a precise estimate of the effect of sterilized intervention within the two standard error bands of the Table 3.1 results: with the full sample, a sale of dollars/purchase of pounds equivalent to 1% of UK M0 would cause a four basis point appreciation of the pound. To give a sense of magnitude, a 1% intervention would be a large but far from abnormal daily intervention (see Figure 3.2 and Table 3.3), and the median daily change in the exchange rate (in absolute value) is 2.2 basis points in our sample. Thus, our results are consistent with the view that sterilized foreign exchange intervention was a useful tool for managing daily fluctuations.

[28] The full list of additional controls, not shown in Table 3.7, is the change in US 3-month treasury rates (available at a daily frequency) and the change in Treasury Bill Rates (available monthly). The additional UK controls include changes in consol yields, commercial paper rates, and UK M0, all available at a monthly frequency.

Table 3.7: Effect of Intervention on the Change in the Exchange Rate by Subsample [1952-1971]

	OLS	IV: Distance				
	(1) Full Sample	(2) Full Sample	(3) Pre-Devaluation	(4) Post-Devaluation	(5) Drop Nov. '67	(6) After 1958
Intervention	0.01^{*}	-0.04^{***}	-0.02^{**}	-0.07^{**}	-0.05^{***}	-0.04^{***}
	(0.00)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)
L.FX Growth	-0.12^{**}	-0.07	-0.02	-0.09	-0.06	-0.09^{*}
	(0.04)	(0.04)	(0.02)	(0.07)	(0.04)	(0.04)
L2.FX Growth	-0.07^{***}	-0.05^{*}	-0.00	-0.06	-0.04	-0.05^{*}
	(0.02)	(0.02)	(0.01)	(0.04)	(0.02)	(0.03)
Change in US Policy Rate	0.01^{**}	0.01^{**}	0.00	0.01	0.01^{**}	0.01^{**}
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Change in UK Policy Rate	-0.05^{***}	-0.05^{***}	-0.05^{***}	-0.07^{**}	-0.05^{***}	-0.06^{***}
	(0.00)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)
1-mo. Exp. Revision	0.59^{***}	0.64^{***}	0.82^{***}	0.56^{***}	0.65^{***}	0.62^{***}
	(0.05)	(0.05)	(0.03)	(0.08)	(0.05)	(0.06)
Observations	4227	4227	3322	905	4209	3105

Notes: Heteroskedasticity-robust standard errors in parentheses. The dependent variable is the log growth in the value of the dollar relative to the pound, and "Intervention" is the daily quantity of dollar sales by the Bank of England (divided by UK M0) undertaken to appreciate the pound. We also present point estimates for two lags of the dependent variable, changes in both the UK "Bank Rate" and the effective federal funds rate, and one month ahead expectation revisions read from futures markets. The IV results are estimated over several subsamples, including one which drops the entire month of the devaluation (November of 1967), and another which keeps only the period after an important liberalization in UK capital markets in 1958 when current account convertibility as a result of the European Monetary Agreement. While the capital account was still not completely liberated, the policy meant much larger capital flows in and out of the UK. Therefore, we separate this period in a sub-sample. All regressions include day-of-week, month and year dummies, as well as additional interest rate controls described in the text, and drop the first trading day after the November 1967 devaluation.
Stars indicate: * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$

3.8.1.2 Using the Lagged Level of the Exchange Rate As an Instrument

The distance instrument used in the paper is $Z_t \equiv (\ln e_{t-1} - \ln e_{t-1}^{target})^2 \times \mathrm{sign}(\ln e_{t-1} - \ln e_{t-1}^{target})$, where the target is time-varying only because of the devaluation in sample. We motivated this instrument by arguing that if yesterday's market closed far from the target, then regardless of today's developments the dealers may be more aggressive in intervening. In practice, the variation in this instrument comes from changes in the lagged, log level of the exchange rate, $\ln e_{t-1}$, which is why we frame discussion of e.g. the exclusion restriction in these terms.

This section of the Appendix demonstrates that using the lagged, log level of the exchange rate alone as an instrument yields very similar results to those obtained with our "Distance" instrument. Table 3.8 reports estimates of our benchmark regression without controls, estimated over the pre-devaluation and post-devaluation periods as well as the entire sample. When using the whole sample, it is important to allow the effect of the level to vary across the devaluation and concomittant change in the exchange rate target in 1967 (since it is really the distance from target that influences the dealers' decisionmaking). Table 3.9 reports the first stage, which accords with intuition, and finally Table 3.10 reports results adding controls motivated by the model in Appendix 3.8.4. The results are quite similar. The lagged, signed, squared distance from target instrument remains our headline result because it has a stronger first-stage, consistent with the idea that the dealers respond more aggressively when the exchange rate is further from target.

Table 3.8: Effect of Intervention on the Change in the Exchange Rate [1952-1971]

	(1) OLS	(2) OLS (Pre)	(3) OLS (Post)	(4) IV: Level	(5) IV: Level (Pre)	(6) IV: Level (Post)
Intervention	0.02***	0.02*	0.03***	-0.07***	-0.05***	-0.12**
	(0.01)	(0.01)	(0.00)	(0.02)	(0.01)	(0.04)
Observations	5244	4278	966	5244	4278	966

Notes: Heteroskedasticity-robust standard errors in parentheses. The dependent variable is the log growth in the value of the dollar relative to the pound, and "Intervention" is the daily quantity of dollar sales by the Bank of England (divided by UK M0) undertaken to appreciate the pound. Columns (1)-(3) present OLS estimates for the whole sample (1) and two sub-periods: pre-devaluation (2) and post-devaluation (3). The IV results using the lagged, log level of the exchange rate are similarly divided and imply that an intervention equivalent to 1% of UK M0 appreciates the pound by between 5-12 basis points, depending on the sample period. All regressions include day-of-week, month and year dummies and drop the first trading day after the November 1967 devaluation.
Stars indicate: * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$

Table 3.9: First Stage Regressions: Effect on Intervention (Dollar Sales)

	(1) Level	(2) Level (Pre)	(3) Level (Post)
Logged, Lagged FX Lvl.	0.18***	0.14***	0.32***
	(0.02)	(0.02)	(0.05)
Logged, Lagged FX Lvl (after Devaluation).	0.03***		
	(0.00)		
Constant	19.02***	14.38***	27.75***
	(1.78)	(1.71)	(3.97)
Observations	5244	4278	966
R^2	0.071	0.029	0.077
F	74.66	71.00	49.37

Notes: Heteroskedasticity-robust standard errors in parentheses. The dependent variable is the daily quantity of intervention (measured as dollar sales divided by UK M0). The signs confirm the economic intuition underlying the relevance assumption of each approach: when the lagged level is high, the pound is weak relative to target and the Bank of England sells dollars to strengthen it. All regressions drop the first trading day after the November 1967 devaluation.
Stars indicate: * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$

Table 3.10: Effect of Intervention on the Change in the Exchange Rate [1952-1971]

	OLS	IV: Level		
	(1) Full Sample	(2) Full Sample	(3) Pre Dep.	(4) Post Dep.
Intervention	0.01*	-0.03***	-0.02*	-0.05**
	(0.00)	(0.01)	(0.01)	(0.02)
L.FX Growth	-0.12**	-0.08*	-0.03	-0.12
	(0.04)	(0.04)	(0.02)	(0.07)
L2.FX Growth	-0.07***	-0.05*	-0.01	-0.08*
	(0.02)	(0.02)	(0.01)	(0.04)
Change in US Policy Rate	0.01**	0.01**	0.00	0.01*
	(0.00)	(0.00)	(0.00)	(0.00)
Change in UK Policy Rate	-0.05***	-0.05***	-0.05***	-0.06**
	(0.00)	(0.01)	(0.01)	(0.02)
1-mo. Exp. Revision	0.59***	0.63***	0.81***	0.54***
	(0.05)	(0.05)	(0.03)	(0.08)
Observations	4227	4227	3322	905

Notes: Heteroskedasticity-robust standard errors in parentheses. The dependent variable is the log growth in the value of the dollar relative to the pound, and "Intervention" is the daily quantity of dollar sales by the Bank of England (divided by UK M0) undertaken to appreciate the pound. We also present point estimates for two lags of the dependent variable, changes in both the UK "Bank Rate" and the effective federal funds rate, and one month ahead expectation revisions read from futures markets. The IV results are estimated over the full sample (2) and two others: pre-devaluation (3) and post-devaluation (4). All regressions include day-of-week, month and year dummies, as well as additional interest rate controls described in the text, and drop the first trading day after the November 1967 devaluation.

Stars indicate: * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$

3.8.2 Results Measuring Intervention Using Dollar Sales

In the paper we express interventions as a percentage of the money supply. Here we re-estimate our main IV using dollar amounts of intervention instead. This allows us to compare our results to the literature which estimates the effects of intervention measured this way. The main results are in Table 3.11 and Table 3.12 below. Specifically, we estimate β_1 in

$$(\ln e_t - \ln e_{t-1}) \times 100 = \beta_0 + \beta_1(\text{Dollar Sales in Millions of USD}_t) + \epsilon_t$$

using the same instrument and sample periods, but without converting dollar sales into pounds using the exchange rate target and deflating by UK M0. The results demonstrate that the transformation does not change the results beyond scaling the size of the estimated coefficient.[29] Given the IV point estimate in column 2, a sale of \$10mn USD would appreciate the pound by 1.2 basis points, and a sale of \$1bn USD would appreciate the pound by 1.2 percentage points. This is quite close to other estimates found in the literature: Arango-Lozano, Menkhoff, Rodríguez-Novoa, and Villamizar-Villegas (2020) report an average effect of one percentage point in a meta study of 74 papers.

[29] To see why, note that since M0 is measured monthly and the exchange rate target is constant except for the devaluation, in practice most of the variation in $Q_t = \frac{(\text{Dollar Sales in Millions of USD}_t) \times e_t^{target}}{\text{UK M0 in the previous month}}$ comes from dollar sales, and $corr\left(Q_t, \text{Dollar Sales in Millions of USD}_t\right) = 0.995$

Table 3.11: Effect of Intervention on the Change in the Exchange Rate by Subsample [1952-1971]

	OLS	IV				
	(1) Full Sample	(2) Full Sample	(3) Pre-Devaluation	(4) Post-Devaluation	(5) Drop Nov. '67	(6) After 1958
Dollar Sales (in Millions USD)	0.0002***	-0.0012**	-0.0008***	-0.0020*	-0.0016***	-0.0012**
	(0.0001)	(0.0004)	(0.0002)	(0.0009)	(0.0005)	(0.0004)
Observations	5244	5244	4278	966	5224	3277

Notes: Heteroskedasticity-robust standard errors in parentheses. The dependent variable is the percentage increase in the value of the dollar relative to the pound, and "Intervention" is the daily quantity of dollar sales by the Bank of England (in millions of USD) undertaken to appreciate the pound. Columns (1) and (2) present OLS and IV estimates of the effects of intervention, demonstrating the bias in OLS and suggesting that a sale of $1mn USD appreciates the pound by 0.12 basis points. Columns (3)-(6) present IV estimates for subsamples, where (5) drops the entire month of the devaluation (November of 1967) and (6) keeps only the period after an important liberalization in UK capital markets in 1958 when current account convertibility was restored as a result of the European Monetary Agreement. All regressions include day-of-week, month and year dummies and drop the first trading day after the November 1967 devaluation.
Stars indicate: * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$

Table 3.12: First Stage Regressions: Effect on Intervention (Dollar Sales) by Subsample

	(1) Full Sample	(2) Pre-Devaluation	(3) Post-Devaluation	(4) Drop Nov. '67	(5) After 1958
Lagged, Squared Distance from Target	28.25***	19.66***	46.25***	23.99***	43.62***
	(3.26)	(3.17)	(6.79)	(2.63)	(5.49)
Constant	-3.02***	-0.19	-16.45***	-3.15***	-5.31***
	(0.46)	(0.44)	(1.62)	(0.38)	(0.71)
Observations	5244	4278	966	5224	3277
R^2	0.036	0.019	0.090	0.036	0.050
F	75.12	38.42	46.43	83.03	63.02

Notes: Heteroskedasticity-robust standard errors in parentheses. The dependent variable is the daily quantity of intervention (measured as dollar sales in millions of USD). The signs confirm the economic intuition underlying the relevance assumption: when the lagged distance from target instrument is positive, the pound is "too weak" relative to target and the Bank of England acts to strengthen it. All regressions drop the first trading day after the November 1967 devaluation.
Stars indicate: * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$

3.8.3 Treating Holiday-Driven Deviations from a "Policy Rule" as Shocks to Intervention

This section estimates a policy rule for central bank intervention based on non-holiday trading days, and uses this rule to calculate a counterfactual for the quantity of intervention that would have occurred on each holiday if the bank had been open as normal. Similarly, we estimate an exchange rate forecasting equation on non-holiday trading days and use this to create a counterfactual for exchange rate growth. Then we can regress the deviation of the exchange rate from its counterfactual on the deviation of intervention from its counterfactual and obtain an estimate of the effect of intervention.

We restrict our attention to holidays because we know on these days that the deviation from the policy rule was due to the bank being closed, and thus unrelated to developments in world currency markets. Specifically, we restrict attention to the deviations that occur on UK-specific holidays, during which the Bank of England was closed (and rarely intervened) while the pound continued to trade in New York, Zurich, and other world currency markets.[30] For example, throughout our sample the last Monday of August is a secular holiday called the "Late Summer Bank Holiday" where all British markets and banks close while the rest of world traded as usual. We also include a second secular bank holiday in the winter and Good Friday in addition to Easter Monday. The results are consistent with our earlier IV results (of the same sign and magnitude) despite relying on a very different identification assumption.

Formally, we report point estimates from the following regression: denote $\hat{Q}_t^i$ as predicted intervention using some policy rule i (which we will discuss below). Similarly, let $\hat{e}_t^i$ be a forecast for the exchange rate. Then we estimate the following β^i via OLS:

$$\underbrace{(\ln e_t - \ln e_{t-1}) - \left(\ln \hat{e}_t^i - \ln e_{t-1}\right)}_{\substack{\text{Forecast Error for} \\ \text{the Growth in } e}} = \beta_0^i + \beta^i \times \underbrace{\left(Q_t - \hat{Q}_t^i\right)}_{\substack{\text{Deviation from} \\ \text{the Policy Rule } i}} + \gamma_t^i \tag{3.4}$$

where β_0^i is a constant, γ_t^i is an error term, and we can interpret β^i as the effect of intervention, which has the same economic interpretation and units as the coefficient estimated in our IV regressions above.[31]

In practice, our forecasting procedure is as follows: we use adaptive lasso to choose the policy rule from up to ten lags of the dependent variables and ten lags of all included and excluded instruments used above, in addition to various dummies and time trends. We use adaptive lasso as recent evidence suggests it performs well in time series contexts (Medeiros and Mendes, 2016).

We use lasso because simple rules do not predict either exchange rates or intervention well, and complicated rules require discipline, as it is not *a priori* obvious what belongs in the

[30]As discussed in Section 3.2, while the Bank of England's dealers could call and request that other central banks intervene in offshore markets during holidays, in practice they rarely did so, as they were not in the office: intervention Q_t is zero on approximately 90% of our holidays, and the 10% of holidays with non-zero values are not driving our results (dropping them leaves point estimates unchanged).

[31]Note that using a forecast error on the left hand side is formally almost identical to simply including the variables chosen by the adaptive lasso as controls on the right hand side.

policy rule.[32] Unlike with conventional monetary policy, where the arguments of the central bank's policy rule are well understood, contemporaries were vague on the determinants of day-to-day operations even in their secret, internal communications. As Harry Siepmann unhelpfully wrote in 1936, in a section labeled "tactics":

> The tactics and management of the EEA naturally attract a good deal of attention and comment, but the fact is that the technique of day-to-day operations is not susceptible of much development or variation. Once the objectives have been set by policy, the question of method is a matter for practical judgment and opportunism, which necessarily depends upon the state of the market. In the press and elsewhere an attempt is occasionally made to propound some kind of theory of management… All such hypothetical arguments have the advantage that, by their very nature, they cannot be disproved.

While it is possible that operations could have become more systematic in the period studied here, we conclude that an atheoretic approach has some appeal. Finally, using adaptive lasso to estimate the policy rule (instead of OLS) mitigates concerns of overfitting. If we overfit, this will effectively add "noise" to our policy shock measure, biasing estimates of β^i toward zero.

Note that even having chosen adaptive lasso as our estimator, we still have some freedom over its implementation. Accordingly, we present results from three different forecasting procedures that yield distinct models, labelled according to the information criterion involved in the implementation: either AIC (Akaike Information Criterion), BIC (Bayesian Information Criterion) or Extended BIC (EBIC).[33] In practice it is not clear which model we should favor in our context, which is why we present multiple; see Section 3.8.3.1 below for details. The results of estimating β^i in equation (3.4) via OLS given the different forecasts from the two models are largely consistent, and presented in Table 3.13. Note that the coefficients on intervention have the same interpretation as previous tables, and we thus compare them directly.

The first three columns of Table 3.13 present results for all non-holiday trading days, suggesting that a sale of dollars equivalent to 1% of UK M0 actually depreciates the pound by 3 basis points. We include these columns to show that use of a policy rule alone is not helpful in achieving identification and simply recovers the positive OLS estimates in the first columns of Table 3.1 and Table 3.7. As the Bank of England official Harry Siepmann stated above, day-to-day operations respond in real time to changes in market conditions, meaning that deviations from the rule on the right-hand-side in our setting are generally endogenous to any shocks driving exchange rate movements.

---

[32] To see why simple policy rules have little power, note the first stage of our IV regression in Table 3.2 has an R-squared of .04; fitting parsimonious such simple rules would leave us with little power. Use of Lasso allows us to make nontrivial predictions and raise the R-squared values without concerns of overfitting (the Lasso models described below are capable of explaining 20-25% of the variance in-sample for intervention). Note that this approach benefits heavily from the fact that in our setting the central bank is intervening almost every day, so that a policy rule can be meaningfully estimated.

[33] A fourth option, the corrected AIC or AICc, yields similar results as for the AIC and is thus omitted.

Columns (4) to (6) restrict our sample to holidays, when we know deviation from the policy rule is driven by the closure of the Bank of England and not current developments in financial markets. Here we obtain results of the same sign and magnitude as the IV results in Table 3.7, suggesting that a sale of dollars equivalent to 1% of UK M0 appreciates the pound by 6-7 basis points, depending on the specification. Finally, as a robustness check, Table 3.14 demonstrates the robustness of point estimates for intervention using holidays to the choice of subsample, obtaining similar results.

Table 3.13: Effects of Sterilized Intervention on the Log Growth in the Value of the Dollar using Policy Rules and Forecasts Estimated via Adaptive Lasso: Results by Choice of Information Criterion

	All Dates			Holidays		
	(1) BIC	(2) EBIC	(3) AIC	(4) BIC	(5) EBIC	(6) AIC
Intervention: BIC	0.03^{***} (0.01)			-0.07^{*} (0.03)		
Intervention: EBIC		0.03^{***} (0.01)			-0.06^{+} (0.03)	
Intervention: AIC			0.03^{**} (0.01)			-0.07^{*} (0.03)
Observations	3453	3849	3170	41	52	38
R^2	0.047	0.051	0.045	0.088	0.052	0.117

Notes: Heteroskedasticity-robust standard errors in parentheses. The dependent variable is the log growth in the value of the dollar relative to the pound, and "Intervention" is the daily quantity of dollar sales by the Bank of England (divided by UK M0) undertaken to appreciate the pound, each given as a deviation from a forecast computed using the adaptive lasso and either the BIC, EBIC or AIC as an information criterion (see Appendix 3.8.3.1; results with the AICc are identical to the AIC in our context and omitted). The table compares the results using just holiday dates to the results using all dates in columns (1)-(3). With all dates we replicate the OLS results, while with holidays we recover results consistent with the IV columns in Table 3.7. Note we are missing some holidays because forecasts could not be computed (due to missing data) and that this explains the smaller number of observations when using less parsimonious models (e.g. the AIC selected model).

Superscripts indicate: $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$

Table 3.14: Effects of Sterilized Intervention on the Log Growth in the Value of the Dollar using Policy Rules and Forecasts Estimated via Adaptive Lasso: Results by Choice of Information Criterion and Subsample

	Results Using Holidays:								
	Pre-Devaluation			After 1958			Pre-Devaluation, After 1958		
	(1) BIC	(2) EBIC	(3) AIC	(4) BIC	(5) EBIC	(6) AIC	(7) BIC	(8) EBIC	(9) AIC
Intervention: BIC	-0.08^{+}			-0.05***			-0.04***		
	(0.05)			(0.01)			(0.01)		
Intervention: EBIC		-0.09^{+}			-0.03^{+}			-0.06**	
		(0.05)			(0.02)			(0.02)	
Intervention: AIC			-0.08^{+}			-0.04**			-0.04*
			(0.05)			(0.01)			(0.01)
Observations	30	38	27	33	39	32	22	25	21
R^2	0.074	0.096	0.111	0.189	0.058	0.156	0.206	0.196	0.166

Notes: Heteroskedasticity-robust standard errors in parentheses. The dependent variable is the log growth in the value of the dollar relative to the pound, and "Intervention" is the daily quantity of dollar sales by the Bank of England (divided by UK M0) undertaken to appreciate the pound, each given as a deviation from a forecast computed using the adaptive lasso and either the BIC, EBIC or AIC as an information criterion (see Appendix 3.8.3.1 for details on computation; results with the AICc are identical to the AIC in our context and are omitted). The table demonstrates the robustness of the results in Table 3.13 to the choice of subsample. Note we are missing some holidays because forecasts could not be computed (due to missing data) and that this explains the smaller number of observations when using less parsimonious models (e.g. the AIC selected model). The small number of holidays prevents meaningful estimation over the post-devaluation sample, so these results are omitted.

Superscripts indicate: $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$

3.8.3.1 Adaptive Lasso Implementation

We use the "lassopack" package in Stata (version 1.4.1) provided by Ahrens, Hansen, and Schaffer (2020) to implement adaptive lasso, using default values for implementation, and note the importance of the choice of information criterion.

We allow for up to ten lags of the following: intervention as a fraction of M0, gold reserves as a fraction of M0, lagged squared distance from target (i.e. the instrument used in our earlier IV regressions), growth in the exchange rate, forecast error revisions, changes in the Bank of England policy rate and changes in the Fed Funds rate. We also allow for a linear time trend, a time trend with a break after the devaluation, a dummy for being post-devaluation, day-of-week, month and year dummies, and all of our previous interest rate controls (which were available at a monthly frequency).

Adaptive lasso is a shrinkage estimator; formally, we pick parameters λ and ω_j and solve the resulting optimization problem:

$$\hat{\boldsymbol{\beta}} = \arg\min_{\beta} \frac{1}{2n} \sum_{i=1}^{n} (y_i - \boldsymbol{x}_i \boldsymbol{\beta}')^2 + \lambda \sum_{j=1}^{p} \omega_j |\beta_j|$$

where y_i is either the growth in the exchange rate or intervention as a fraction of UK M0, and x is a vector of p potential predictors. Adaptive lasso uses a set of initial OLS estimates to pick the ω_j, and λ is chosen so that the resulting model minimizes a particular information criterion: either the Akaike information criterion (AIC), Bayesian information criterion (BIC), or the Extended BIC (EBIC). The BIC and EBIC yield more parsimonious models, while the AIC has a greater in-sample fit and produces more substantial forecasts. In Table 3.13 and Table 3.14 of the main text, we presented results for all three approaches.[34]

To obtain the coefficients shown in Table 3.13, we regress the forecast error for exchange rate growth on the forecast error for intervention (i.e. the deviation from the policy rule). Denote $\hat{Q}_t^i$ as predicted intervention using either the $i = BIC, AIC$ or $EBIC$ model. Similarly, let $\hat{e}_t^i$ be the forecast for the exchange rate. Then we estimate the following β^i for each model via OLS:

$$\underbrace{(\ln e_t - \ln e_{t-1}) - \left(\ln \hat{e}_t^i - \ln e_{t-1}\right)}_{\substack{\text{Forecast Error for} \\ \text{the Growth in } e}} = \beta_0^i + \beta^i \times \underbrace{\left(Q_t - \hat{Q}_t^i\right)}_{\substack{\text{Deviation from} \\ \text{the Policy Rule } i}} + \gamma_t^i$$

where γ_t^i is an error term. As Table 3.13 shows, in this step it is critical to restrict t to be in the set of dates which are holidays: there, we know the deviations from the policy rule on the right-hand-side are plausibly exogenous. Note the economic interpretation of β^i is similar to our IV regressions with controls and we should expect similar point estimates if both approaches are truly identifying as-good-as-random variation in Q_t.

[34] The software package we use also produces estimates based on the corrected AIC (AICc) but in our context this delivers identical results to using the AIC.

3.8.4 A Reduced-Form Portfolio Balance Channel Model of Sterilized Foreign Exchange Intervention

We use a simple portfolio balance channel model of sterilized foreign exchange intervention to discipline our regression specifications and choice of controls.[35] We also use this model to illustrate a case where the distance instrument's exclusion restriction only holds "approximately" due to the presence of mean-reverting fundamental shocks to the level of the exchange rate.

Define e_t as the exchange rate in terms of the home currency (pounds) per unit of foreign currency (dollars), so that an increase in e is a depreciation of home's currency. World demand for home (pound) bonds is determined by some function D which is increasing in their excess return over foreign (dollar) bonds. Letting R_t and R_t^* denote gross interest rates on h-period home and foreign currency bonds,

$$\text{World Portfolio Share of Home Bonds} \equiv \chi_t D\left(R_t - R_t^*\left(\frac{E_t[e_{t+h}]}{e_t}\right)\right).$$

This is a standard reduced-form model of UIP deviations.[36] We assume that the log of the demand shifter χ_t follows an AR(1) process:

$$\ln \chi_t = \rho \ln \chi_{t-1} + \delta_t$$

where $\rho \in [0,1]$ and δ_t is a white noise process. We assume the supply of bonds available for the private sector to hold is given by total home (UK) government debt, denoted B_t, less central bank holdings, denoted A_t. If we let W_t denote global wealth, the world portfolio share of home bonds must be:

$$\text{World Portfolio Share of Home Bonds} \equiv \frac{B_t - A_t}{W_t}$$

Given the supply of home currency bonds, equilibrium in the market for home debt is achieved through the exchange rate adjusting today, given gross interest rates and expectations of the exchange rate. The intuition is simple: if the expected returns on the two bonds are unchanged, their relative price must adjust when the central bank alters the supply. This is achieved through exchange rate adjustment.

Setting supply equal to demand, normalizing the financial stock variables A_t, B_t and W_t by UK M0 at time t (denoting them, a_t, b_t, and w_t, respectively) yields

$$\frac{b_t - a_t}{w_t} = \chi_t D\left(R_t - R_t^*\left(\frac{E_t[e_{t+h}]}{e_t}\right)\right). \tag{3.5}$$

[35]This is sometimes also called the "imperfect asset substitutability" or "risk premium" channel.

[36]This model follows e.g. Frankel (1984), but assuming a representative investor and a more general relationship between the portfolio share and excess return. This standard model is used in undergraduate texts (e.g. Krugman, Obstfeld, and Melitz, 2015) and applied work, often with extensions (e.g. Cushman, 2007). A similar relationship also holds in microfounded models such as Gabaix and Maggiori (2015), who obtain a similar demand function assuming incentive-compatibility constraints prevent risk-neutral investors from arbitraging away UIP deviations.

Log-linearizing around a steady state in which all variables in (3.5) are constant yields a framework for our regressions. Then, since we do not observe the size of the bank's balance sheet at a daily frequency, but only the changes, we first difference. In what follows, for any variable z_t, we use z to denote the steady state of z_t; writing $\ln R = \ln(1+r) \approx r$, and letting Δ be one period differences over t, we obtain

$$\Delta \ln e_t = E_t[\ln e_{t+h}] - E_{t-1}[\ln e_{t+h-1}] - \frac{R}{R^*}\Delta r_t + \Delta r_t^* - \frac{\varphi}{b-a}\Delta a_t + \epsilon_t$$

where ϵ_t collects unobserved structural errors, and φ is a positive, uninteresting collection of steady-state values.[37] The sign accords with intuition regarding sterilized foreign exchange interventions, as Δa_t in this model corresponds with the sterilized intervention data collected and divided by M0, or Q_t in the text.[38]

For the instrument to work, we need the *level* of the exchange rate to be uncorrelated with the *changes* in the structural shocks ϵ_t. To see how mean reversion in fundamental shocks violates this, consider the case where b_t and w_t are constant, so that the error simplifies to

$$\epsilon_t \equiv \varphi \Delta \ln \chi_t$$

Since the distance instrument contains a lag of the exchange rate, it contains the shock χ_{t-1}. It will generally only be orthogonal to the error term if $\Delta \ln \chi_t = (\rho - 1)\ln \chi_t + \delta_t$ is i.i.d. which is only the case if $\rho = 1$ and the demand shifter is a random walk. Thus, we note that if ρ is close to one then the exclusion restriction will be almost satisfied.

This exclusion assumption was motivated by the observation that most models of the exchange rate do not out-perform a simple random walk model in out-of-sample forecasting exercises, particularly over short horizons (Meese and Rogoff, 1983; Rossi, 2013); the hope is that any mean reversion (and corresponding downward bias in our IV estimates of the effects of intervention) is small enough to be ignored. Appendix 3.8.5 explores this formally, showing through simulation that empirically-plausible levels of mean reversion are unlikely to explain the IV results obtained in the main paper alone (under the null hypothesis that intervention is not effective).

Regarding our specification with controls, we estimate the following version of the above equation:

$$\Delta \ln e_t = \beta_0 + \beta_1 Q_t + \beta_2 \Delta E_t[\ln e_{t+h}] + \beta_3 \Delta r_t + \beta_4 \Delta r_t^* + X_t + \mu_t \tag{3.6}$$

where $\Delta \ln A_t$ has been replaced by Q_t and X_t includes day-of-week, month and year fixed effects, two lags of the dependent variable and various interest rate controls. We use changes in one month forward rates to proxy for $\Delta E_t[\ln e_{t+h}]$, and policy rates in each country as

[37] Formally, $\varphi \equiv \frac{D'(R-R^*)}{D(R-R^*)}R^*$ and $\epsilon_t \equiv \varphi\left(\frac{b}{b-a}\Delta \ln b_t - \Delta \ln w_t - \Delta \ln \chi_t\right)$.

[38] To see the correspondence between the "dollar sales" variable Q_t and the change in the central bank's holding of pound bonds $\Delta \ln a_t$, consider an example where the Bank of England (the home country) sells dollars and buys pounds. When the Bank of England buys home bonds to sterilize the intervention, a_t increases.

our measures of r_t and r_t^*. Note that from the model's perspective, ideally we would use changes in risk-free h-month rates, but these are not always available at daily frequencies (an exception is the US 3-month treasury rate which is available at a daily frequency from FRED). We thus also include various changes in available interest rates at monthly frequencies as described in Appendix Section 3.8.6.

3.8.5 Long-Run Mean Reversion and Identification Concerns

The exclusion assumption in our IV regressions is that the log-level of the exchange rate is not useful for forecasting its own growth rate, conditional on the (secret) interventions by the central bank. This assumption is motivated by the fact that fundamentals-based models of the exchange rate usually underperform a simple random walk without drift in out-of-sample forecasting exercises (Meese and Rogoff, 1983; Rossi, 2013). However, a large literature documents that shocks to the exchange rate may be mean-reverting to some fundamental value in the very long run.[39] Could plausibly-small deviations from the random walk model, consistent with long-run mean reversion, explain our results?

To show that this is unlikely, this appendix simulates data from a model where intervention is actually ineffective (true $\beta_1 = 0$) but where our IV estimates will nevertheless tend to be negative because of mean reversion in the exchange rate. After running multiple simulations, we can examine the distribution of results and conclude whether or not our actual IV estimates are likely to come from such a model. We choose parameters to match key features of the data, including our observed strong first stage and positive OLS coefficient (note that since we assume $\beta_1 = 0$, replicating weakly-positive but near zero OLS coefficients necessarily implies little bias in OLS). We conclude by noting that our actual IV estimates lie well outside the 90% range of estimates from these simulations except for cases where mean reversion is unreasonably high, suggesting that we can reject that this model of the world is behind our IV estimates. We conclude that our IV estimates suggest that intervention is effective, but caution that they may be biased downward.

First consider the simplest estimating equation from our paper, where $\ln e_t \equiv s_t$:

$$s_t - s_{t-1} = \hat{\beta}_0 + \hat{\beta}_1 Q_t + \hat{\epsilon}_t$$

and where we instrumented for Q_t with s_{t-1}. The results from this exercise are in Table 3.8 and Table 3.9. Assume the true model is

$$s_t = \rho s_{t-1} + \beta_1 Q_t + \epsilon_t \tag{3.7}$$

and the policy rule for intervention is

$$Q_t = \phi s_{t-1} + \epsilon_t + \mu_t \tag{3.8}$$

[39] See e.g. Sweeney (2006); Bessec (2003); Taylor and Peel (2000). Murray and Papell (2002) surveyed the literature on convergence to PPP and found that deviations have a half life of 3-5 years, but argued that the appropriate range is closer to 2-20 years.

where ϕ indicates the responsiveness to the lagged level of the exchange rate from a target; we assume the target is zero here for simplicity. The presence of ϵ_t reflects the fact that the central bank responds to current shocks (offsetting them if $\beta_1 < 0$). μ_t is a white noise process reflecting idiosyncratic decisions by the central bank's dealers and perhaps an imperfect knowledge of the shocks ϵ_t.

The exclusion restriction requires that $\rho = 1$, but long-run convergence to some mean suggests ρ may be slightly less than one.[40] This would bias the IV estimates of β_1 downward. However, the low levels of ρ needed to quantitatively explain our results under the null hypothesis that $\beta_1 = 0$ are inconsistent with modest long-run mean reversion.

To show this, we simulate data from the data generating process given by (3.7) and (3.8) where $\beta_1 = 0$ (by assumption) and $\phi = .14$ (as estimated from the pre-devaluation data).[41] We choose $\epsilon_t \sim N(0, .05)$ to match a standard deviation of changes in the log exchange rate of .05, and we choose $\mu_t \sim N(0, \sigma)$, where $\sigma = .75$ is chosen to approximate the F-statistic and R^2 of the first stage regression.[42]

Given these choices, and a choice for ρ, we simulate 500 IV estimates on samples of 5,000 trading days each, reporting the mean and 90% distribution. Figure 3.8.1 plots the average IV estimate and 90% distribution for specific choices of $\rho \in [.9975, .9999]$ corresponding to decay rates of 0-5% at the monthly level (assuming 21 trading days in a month) or, equivalently, to shock half-lives of as little as a year for comparision to the literature on long-run convergence to PPP. Our point estimate generally lies well below the simulated results, even for large amounts of mean reversion, suggesting that mean reversion alone is unlikely to explain our IV results.

3.8.6 Data Sources

For exchange rates, we rely on data collected by Accominotti et al. (2019) from the *Financial Times*, and patch in missing data from Global Financial Data (GFD).[43] We also take one-month forward premiums from this same source, which we use to construct the forecast revision $\Delta E_t[\ln e_{t+h}]$ in equation (3.6). All other interest rate controls and UK M0 are downloaded from the Federal Reserve Economic Database (FRED), except for the UK policy rate which is taken from the Bank of England.

[40] Note even if $\rho = 1$ and $\beta_1 < 0$, then the central bank's policy rule implies the exchange rate will not follow a random walk, but instead follow a stable AR(1); to see this, substitute (3.8) into (3.7).

[41] Calibrating the simulation using the post-devaluation sample and investigating the ρ necessary to replicate those results only makes the conclusion more stark, both because our point estimate $\hat{\beta}_1$ is much more negative (-0.17 instead of -0.05) and the coefficient on the first stage is larger ($\phi = .32$), which generally improves the behavior of the IV estimator in our simulations. We thus report only the more precisely estimated pre-devaluation sample.

[42] A low R^2 requires σ to be large, but results are not sensitive to this choice (we tried $\sigma \in \{.5, .75, 1.0\}$).

[43] We prefer to use the Accominotti et al. (2019) data from *Financial Times* since it is better documented, and patch in for dates when data is missing due to e.g. bad scans of the *Financial Times*. In particular we also use the GFD data to obtain prices for holiday dates when the world market was still trading. In practice, on the days when they overlap, the correlation between the *Financial Times* data and the GFD data is effectively one.

Figure 3.8.1: Bias in IV under the Null Hypothesis that $\beta_1 = 0$ with Mean Reversion

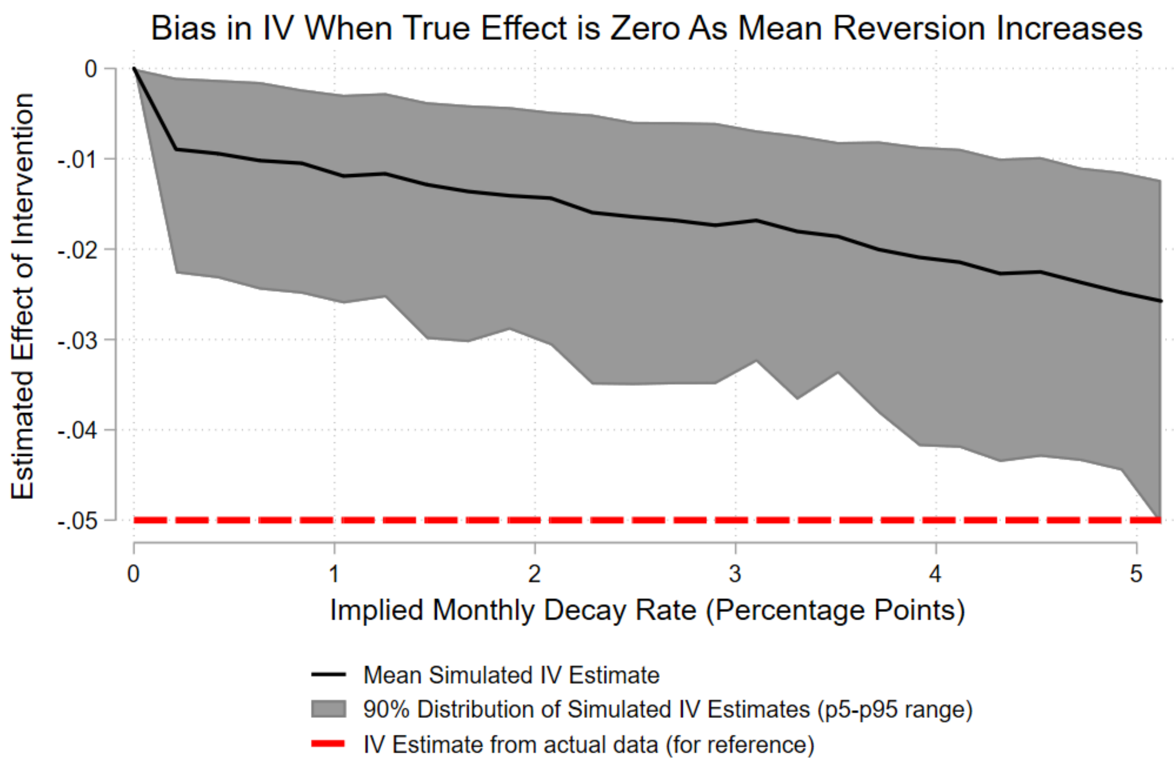

Notes: Bias in IV under the null hypothesis that $\beta_1 = 0$ for $\rho \in [.9975, .9999]$, which corresponds to the monthly decay rates on the horizontal axis. Note a monthly decay rate of e.g. 3% is quite large (as this implies a half life of shocks of as little as two years).

For our intervention variable, we deflate by the previous months UK M0. We also include the following monthly changes as controls in some specifications (with FRED series names): changes in UK M0 (MBM0UKM); changes in US Treasury Bill rates (INTGSTUSM193N); changes in 3-month Treasuries (TB3MS); changes in UK console yields (YCLTUK); and changes in UK commercial paper rates (DRSTPUKM). For daily interest rates, changes in the Bank Rate were downloaded from the Bank of England's web site; we also included changes in the US policy rate as captured by the effective Fed Funds rate (DFF) and changes in 3-month treasury rates (DTB3), though these series only begin in 1954.

One non-trivial data cleaning issue bears mentioning: foreign exchange markets in Europe were open on Saturdays from April 15th, 1955 to October 17th, 1964, and the Bank of England intervened over the weekend as a result. However, the Bank of England recorded its intervention for both Friday and Saturday jointly, so that we do not observe how much intervention occurred on each day. Rather than impute, we instead treat Friday and Saturday as one trading day for the purposes of estimation, and construct changes from end-of-day Thursday to end-of-day Saturday when creating our controls and non-intervention variables. All references to the number of observations made in the text account for this, counting Friday and Saturday together as one trading day instead of two.

3.8.6.1 Details and Treatment of the 1967 Devaluation

The announcement of the devaluation occurred on a Saturday (November 18th, 1967) when the market was closed. We drop the first trading day after the devaluation (Tuesday, November 21st 1967). Monday was declared a "bank holiday" which is also dropped from our analysis (Forrest, 2010).[44] We "drop" these days by not including them as observations in any regressions or tabulations and also set the appropriate values to missing for our regressors (so that they are not influencing our results when we take e.g. lags of the growth rate of the exchange rate). This procedure is what is meant by the line in each table caption, "All regressions include day-of-week, month and year dummies, as well as additional interest rate controls described in the text, and drop the first trading day after the November 1967 devaluation."

Our regression specifications that drop the entirety of the month of November 1967 further ensure that the devaluation episode is not driving our results. The motivation behind this exercise is that discussion of devaluation began internally prior to November 18th (on November 16th, the Chancellor of the Exchequer recommended a devaluation to the Cabinet) and one might be skeptical that this was completely private information. However, as shown in the original results that estimated regressions over the full sample, it does not appear that including the month of November, 1967 drives our results.

[44]Note we do not include such bank holidays in our "holidays" approach, as these closures are far from randomly assigned.