Three Essays on the Interaction of Microeconomic Frictions and Macroeconomic Outcomes

Inaugural-Dissertation

zur Erlangung des Grades eines Doktors

der Wirtschafts- und Gesellschaftswissenschaften

durch die

Rechts- und Staatswissenschaftliche Fakultät

der Rheinischen Friedrich-Wilhelms-Universität

Bonn

vorgelegt von

VOLKER TJADEN

aus Hannover

Bonn 2013

Dekan: Prof. Dr. Sandmann
Erstreferent: Prof. Dr. Bayer
Zweitreferent: Prof. Dr. Müller

Tag der mündlichen Prüfung: 12.08.2013

Acknowledgments

This dissertation benefited from the fruitful comments and support from many people. First and foremost, I want to thank my adviser Christian Bayer for his continued support and encouragement. I learned very much from our collaboration and countless discussions. I am also grateful that he encouraged my need for experimentation in terms of both research areas and methodology.

Felix Wellschmied, through years of common learning and research, has become both an intellectual companion and a friend to me. I want to thank him for our successful joint research and many hours of spirited and entertaining discussions during long hours of work.

I have been privileged to be a member of the *Bonn Graduate School of Economics*. I want to thank Urs Schweizer, Silke Kinzig, and Pamela Mertens for making this place the nourishing environment that it is. I have always had adequate financial resources to conduct my research and attend conferences and summer schools. In this context, I would also like to thank the German Research Foundation (DFG) for financial support.

I conducted an important part of the research that led to the third chapter of this dissertation during a guest stay at the economics department of the University of Pennsylvania. The *German-American Fulbright-Commission* supported this stay through a scholarship for which I am very grateful. I am also indebted to Dirk Krüger for inviting me to UPenn and want to thank everyone at the department for being such good hosts. Special thanks go to Cecilia Fieler, Jesús Fernández-Villaverde, and Iourii Manovskii for many helpful discussions of my work.

In Bonn, I conducted most of my work in the environment of the *Macroeconomics*

iv

Contents

Contents

Contents

List of Figures

List of Figures

x

List of Tables

Introduction

In the 1970's, what became known as the "'rational expectations revolution" transformed macroeconomic research. Up until then, macroeconomics had concentrated on estimating systems of ad hoc aggregate relations ("Cowles macroeconometrics") with little reference to individual decision making or to the underlying microeconomic heterogeneity. Following the contributions by, among others, Lucas, Sargent, and Wallace macroeconomic outcomes were cast as dynamic stochastic equilibria which were the result of rational optimal decision making by economic agents (Lucas, 1972; Sargent and Wallace, 1976). However, initial research on this new type of quantitative macroeconomic models which built on the influential work by Kydland and Prescott also abstracted from macroeconomic heterogeneity and instead made use of the assumption of a representative agent and firm (Kydland and Prescott, 1982). At that time, economists were lacking the conceptual and numerical tools for solving dynamic models that explicitly accounted for microeconomic heterogeneity and frictions. Also, it was not obvious that more detailed attention to these phenomena was important when trying to understand the business cycle dynamics of aggregate quantities and prices, or long-run growth.

In the last two decades, two simultaneous developments have changed this. First, microeconometric research in labor economics and industrial organization has revealed large and persistent cross-sectional dispersion and idiosyncratic volatility among individual market participants. The risk that a single household or firm faces is typically an order of magnitude larger than what is measured for aggregate variables. To give an example, even within four-digit SIC industries in the U.S. manufacturing sector, the average difference in logged total factor productivity (TFP)

between an industry's 90th and 10th percentile is 0.651. This means that a plant at the 90th percentile produces almost twice as much output with the same measured inputs as does the 10th percentile plant. The fact that these productivity differences are quite persistent begs the question what feature of the microeconomic structure in factor or product markets keeps factors from being reallocated to the more efficient production units and what this implies for aggregate productivity (Syverson, 2004; Foster et al., 2008).

The second important development was the dramatic reduction in the cost and availability of computing power[1] and the development of a conceptual framework for modeling the interaction of microeconomic heterogeneity and frictions with macroeconomic outcomes. Early contributions dealt with the steady state properties of an economy with entry, exit and within-industry productivity heterogeneity (Hopenhayn, 1992; Hopenhayn and Rogerson, 1993) and the effect of market incompleteness on aggregate savings and the real interest rate (Huggett, 1993; Aiyagari, 1994). An important milestone was the extension by Krusell and Smith which allowed to compute dynamic aggregate equilibria in the presence of aggregate business cycle risk and uninsurable risk while tracking the entire distribution of economic agents (Krusell and Smith, 1997, 1998). In recent times, ever more efficient algorithms for handling equilibrium dynamics under rational expectations and in the presence of microeconomic heterogeneity have become available (Algan et al., 2010; Den Haan and Rendahl, 2010; Kim et al., 2010; Malin et al., 2011; Reiter, 2010).

This dissertation contributes to the ongoing research agenda of trying to understand macroeconomic outcomes in their interdependence with the underlying microeconomic heterogeneity and frictions. The first chapter investigates the effects of modeling plant-level productivity heterogeneity and frictions to capital adjustment on aggregate investment dynamics in the context of a two-country general equilibrium framework. The third chapter extends this notion of plant-heterogeneity by introducing idiosyncratic differences in demand alongside productivity heterogeneity. The context there is a partial equilibrium framework where plants face sunk entry and fixed costs to be able to serve export markets. The economic application is a study of aggregate export elasticities to different kinds of aggregate shocks using

---

[1]This development includes the increase in the number of operations per second of individual CPUs, the much improved scalability and accesibility of CPU clusters and, more recently, shared memory parallelization in desktop computers using multi-core CPUs and specialized GPUs.

a version of the model that has been estimated on plant-level data. The second chapter differs from the others in that its focus is on worker heterogeneity together with a more general notion of job heterogeneity which may but does not have to result from differences in plant heterogeneity. Individual workers face a search friction in the labor market and we study its implications for resulting wage inequality. This application is set in a partial equilibrium setting as well. I give a more detailed summary of each chapter in the remainder of this introduction.

CHAPTER 1.[2] This chapter introduces fixed costs to capital adjustment at the plant-level into an otherwise standard two country real business cycle model. The presence of fixed adjustment costs implies increasing returns to scale in the investment technology. Plants therefore adjust in a "lumpy" fashion, i.e. in infrequent large bursts. A large literature that tries to micro-found aggregate investment has established this as an accurate description of plant investment behavior (Caballero et al., 1995; Doms and Dunne, 1998; Cooper et al., 1999; Bachmann and Bayer, 2011a,b). In contrast, international real business cycle models when fitting the volatility of investment series from national accounts have typically relied on convex adjustment costs to capital at the aggregate level (Baxter and Crucini, 1993; Schmitt-Grohe and Uribe, 2003). Previous studies in a closed economy general equilibrium setting found no effects of microeconomic fixed costs to capital adjustment on aggregate investment dynamics (Khan and Thomas, 2003, 2008).

The calibrated model yields two main results. First, unlike in the closed economy setting, in a two country model non-convex capital adjustment costs matter for the aggregate in that they dampen investment dynamics at the national level. The effect is the stronger, the more open an economy is to trade and it vanishes when letting the model converge towards two separate closed economies. Second, the aggregate dynamics of the model can be accurately replicated by the assumption of a homogeneous firm facing convex adjustment costs. This is a useful finding for applied work since the homogeneous firm model is much simpler to solve. Finally, while for any value of fixed adjustment costs there exists an accurate convex adjustment cost approximation, our results caution against attaching a structural interpretation to estimated convex adjustment costs from open economy models. While the mapping

---

[2]The chapter is based on the paper "Large Open Economies and Fixed Costs of Capital Adjustment", which I jointly wrote with Christian Bayer, (Bayer and Tjaden, 2013).

from fixed to convex adjustment costs is stable with respect to variations in demand side parameters (openness, investment tax credits), it is not invariant to variations in those parameters that directly enter the firm's trade-off between investment and non-adjustment, i.e. mark-up and idiosyncratic risk.

CHAPTER 2.[3] The second chapter leaves the realm of international economics and moves to labor economics instead. We try to understand how much of observed wage inequality is due the presence of a search friction in the labor market. Mincerian wage regressions explain only about a third of observed wage variation which means that much inequality is among observationally equivalent workers. If sampling job offers in unemployment takes time and is subject to the opportunity cost of foregone wages, identical workers rationally accept a range of heterogeneous job offers. Understanding how much of residual inequality results from search frictions as opposed to unobserved heterogeneity is of first order importance when evaluating the efficiency of labor markets and designing appropriate social insurance schemes. Previous research found more than 40 percent of wage inequality to be *frictional* (Postel-Vinay and Robin, 2002; Carrillo-Tudela, 2012).

A key mechanism for generating large frictional dispersion in search models is the ability to continue sampling job offers on the job (Hornstein et al., 2012). The more job offers workers receive on the job, the less of an option they are giving up when moving out of unemployment. This makes them more willing to accept relatively poor job offers and allows them to quickly move into very good matches which means many high wage workers.

We provide empirical evidence from the Survey of Income and Program Participation (SIPP) that an important share of job to job transitions is however not value improving. We build a structural search model that explicitly accounts for those losses. It includes a number of important channels that enlarge the set of acceptable job offers to the worker: skill accumulation on the job, skill loss in unemployment and search on the job. Nonetheless, and in contrast to previous findings, the model attributes only 14 percent of total wage inequality to the search friction. The crucial novelty that explains our different estimates is the introduction of reallocation shocks

---

[3]The chapter, "Quantifying the Contribution of Search to Wage Inequality" is based on a joint project with Felix Wellschmied, (Tjaden and Wellschmied, 2013). That paper previously circulated as "Exploring the Causes of Frictional Wage Inequality", (Tjaden and Wellschmied, 2012).

for employed workers. They leave the worker only the outside option of accepting a new job or moving into unemployment. When excluding that assumption, our estimated contribution of the search friction to wage inequality jumps to 38 percent, much closer to previous estimates.

CHAPTER 3.[4] The third chapter returns to the field of international economics. When manufacturing plants enter into the export market, they on average exhibit higher export revenue growth rates than incumbent exporters for a number of years and an exit hazard from exporting that is declining in tenure. These facts are at odds with standard fixed cost models typically used in empirical studies of export participation (Roberts and Tybout, 1997; Das et al., 2007; Willis and Ruhl, 2009). They hint at the presence of a market specific demand factor that entrants have to slowly accumulate. I explore the the macroeconomic implications of introducing this notion of *customer capital* into a dynamic model of plant exporting behavior. Other sources of heterogeneity are differences in revenue productivity and stochastic entry and fixed costs of exporting.

I structurally estimate the model on a large panel data set of German manufacturing plants between 1995 and 2008. The estimation method is a Simulated Method of Moments (SMM) procedure. The high dimensionality of the problem induces me to use a global particle swarm optimization algorithm to find the minimum of the objective function. The results provide a first estimate from plant level data of the costs of maintaining and expanding a costumer base in export markets. Implied costs are sizable and constitute by far the most important export associated costs. During the time of the sample, the average firm spends between 3 and 4 millions of 1995 euros on marketing activities. Average entry costs into exporting of around 33,467 euros are comparatively small.

In terms of predictive power, the model outperforms a standard fixed cost model of exporting by correctly predicting the sizable export expansion in the data after the year 2003. The model can also reconcile large predicted trade gains after a tariff reduction with a relatively low elasticity of substitution between exported and domestic goods in the export market. The discrepancy between high estimated substitution elasticities from trade reactions to tariff liberalizations and low elasticities

---

[4]The chapter is based on the paper "Foreign Customer Accumulation and Export Dynamics", (Tjaden, 2013).

needed to reproduce export dynamics at business cycle frequency had previously been called the *elasticity puzzle* of international economics.

Chapter 1

Large Open Economies and Fixed Costs of Capital Adjustment

1.1 Introduction

Since Backus et al. (1992) adapted the real business cycle model to the context of international economics, it has been well known that the model suffers in this context from excess volatility in capital reallocation across countries. The introduction of trade in intermediate goods as in Backus et al. (1994) tends to mitigate the effect, but this depends crucially on parameter values and model assumptions. An ample range of applications remains in which the assumption of unobstructed, frictionless capital flows across borders implies an investment volatility relative to output far in excess to what is consistent with national data.[1] As was first demonstrated by Baxter and Crucini (1993), the model's fit can be significantly improved by the introduction of convex adjustment costs to capital at the national level. Over time, this has become a standard practice.

However, this can only remain a kludge for removing excess investment volatility as macroeconomic research micro-founding aggregate investment behavior has found *fixed* and not convex adjustment costs to be the dominant friction to capital adjustment at the plant or firm level.[2] A priori, it is not clear whether a stand-in representative firm with quadratic adjustment costs is a good representation of many

---

[1] Common examples are the cases of perfect substitutability between consumption goods in multi-country models (e.g. Den Haan et al. (2011)), a small open economy setting (e.g. Schmitt-Grohe and Uribe (2003)) or the presence of nominal frictions (e.g. Chari et al. (2002)).

[2] Early studies using US data are Caballero et al. (1995), Doms and Dunne (1998), Caballero and Engel (1999) and Cooper et al. (1999). More recent examples are Gourio and Kashyap (2007) using US data and Bachmann and Bayer (2011a,b) using German firm-level data.

firms that exhibit lumpy investment. This leaves open what effect on the micro-level these representative firm open economy models actually capture; renders their estimated versions potentially subject to the Lucas critique; and finally challenges their policy predictions. We therefore introduce fixed capital adjustment costs into an otherwise standard two goods, two country real business cycle model in the spirit of Backus et al. (1994) and ask what the aggregate consequences are.

Specifically, we want to answer in a model of two large open economies the following three questions: First, do fixed adjustment costs at the firm level have aggregate consequences at all? Second – if so – can these effects be captured by a stand in aggregate model with convex adjustment costs and a representative firm? Third – if they can – how stable are the identified convex adjustment cost parameters of this stand-in model with respect to changes in the non-adjustment-cost parameters of the underlying lumpy-investment model?

We find that fixed adjustment costs do matter for the aggregate, that their effects can be captured by quadratic adjustment costs, but that the so identified quadratic costs change when non-adjustment cost parameters change, i.e. quadratic costs lack "fundamentalness".

Given that fixed costs and quadratic costs are often cast as rival specifications our research strategy, which is closely related to Chang et al. (2010),[3] may come as a surprise. Our reason for exploring the possibility of representing fixed adjustment costs by quadratic ones is the following: All papers studying the role of fixed adjustment costs in general equilibrium have found approximate aggregation in the sense of Krusell and Smith (1998). In particular, they found that a log-linear law of motion describes the dynamics of the aggregate stock of capital, the only endogenous aggregate state variable in these models, just as it does in a quadratic adjustment-cost model. This suggests that such a model can capture the aggregate dynamics of the heterogenous firm, lumpy investment model.

In a one-sector closed economy model, Khan and Thomas (2003, 2008) show this to be true in a very specific sense. They find fixed adjustment costs to be entirely irrelevant for aggregate dynamics. Hence, firms can be represented by a single firm not facing any adjustment costs. In their closed economy general equilibrium model, this

---

[3]They have recently argued for more caution when aggregating over ex post heterogeneous micro units in the presence of frictions. They look at an incomplete-markets, heterogeneous household setup, and show that the estimated parameters of homogeneous agents models in these settings can often lack "fundamentalness" in the sense of the Lucas critique.

irrelevance result arises because the household's desire to smooth consumption does not allow for much variation in savings behavior. This yields that small additional changes in the interest rate undo all potential aggregate effects of microeconomic lumpiness in a closed economy because indivdiual investment timing is very sensitive to interest rate movements notwithstanding the fixed adjustment costs, see House (2008), while savings are not. In an open economy setting, domestic savings are not the only means to finance investment and consumption smoothing can also be achieved via movements in the current account. This should, in theory, dampen interest rate responses which leaves room for fixed adjustment costs to matter.

This intuition turns out to be right. In a two country model, non-convex capital adjustment costs matter for the aggregate in that they dampen investment dynamics at the national level. The effect is the stronger, the more open an economy is to trade (i.e., the smaller its home bias in consumption), such that the Khan and Thomas (2003, 2008) result obtains when letting the model converge towards a model of two separate closed economies.

Finding a dampening effect of fixed costs and approximate aggregation lets us investigate then our conjecture of approximate representation. Indeed, a homogeneous firm facing convex adjustment costs can act as a handy stand-in to replicate the aggregate dynamics. Finally, we asses the "fundamentalness" of these convex adjustment cost approximations. To do so, we construct matches between convex and non-convex adjustment cost parameters while varying other model parameters, in particular openness to trade, the introduction of an investment tax credit, variations in the idiosyncratic profitability risk, and the curvature of the production function (characterizing the mark-up firms can charge). It turns out that the link between the two cost specifications is stable with respect to variations in those model parameters that only characterize the aggregate trade-off between investment and consumption, i.e. openness and the tax credit. Yet, it is not stable to variations in those parameters that directly enter the (firm's) trade-off between investment and non-adjustment, i.e. the mark-up and the idiosyncratic risk.

The intuition for these seemingly contradictory findings is rather straightforward if one thinks of the equilibrium as the solution to a social planner's problem. A social planner chooses sequences of distributions of capital across production units in order to maximize utility of the representative household from consumption and leisure. In choosing these distributions, the planner needs to take into account both the direct costs from capital adjustment as well as the indirect, efficiency costs from

having otherwise equal plants employing different levels of capital. Between the two costs there is a trade-off. The more frequent firms adjust, the more adjustment costs are paid but the more efficiently is the aggregate stock of capital distributed. This means non-adjustment at the firm level has an efficiency cost, which is a function of the curvature of the production function and the dispersion of bliss-points in capital stocks. More specifically, non-adjustment is the more costly, the higher the production function's curvature and the faster the distribution of bliss points flattens out due to idiosyncratic shocks, i.e. the more these shocks are dispersed; and if non-adjustment at the firm level is more costly, the more often will a social planner change idiosyncratic capital stocks, which – as a byproduct – allows to be more reactive to aggregate shocks, too.[4] Similarly, if the curvature is low, the planner can use the intensive margin of those firms adjusting to react to aggregate shocks.

This in mind, our first result of approximate representation (beyond approximate aggregation) implies that the cost of adjustment in the aggregate are approximately quadratic. Second, since changes in parameters outside the firm's problem (openness, investment tax-credit) do not affect the trade-off between efficiency and adjustment frequency, they do not change the approximate representative firms' problem. Third, since changes in curvature or risk effectively change this trade-off, changes in the production function (or in productivity heterogeneity) change the adjustment costs of the approximately representative firm.

The remainder of the chapter is organized as follows: Section 2 reviews a number related recent contributions to the literature. Section 3 presents the model. Section 4 briefly introduces the numerical solution method. Section 5 explains parameter choices. Section 6 presents our main results – fixed adjustment costs matter but aggregate dynamics are indistinguishable from a representative firm model with quadratic adjustment costs. Section 7 discusses how stand-in quadratic adjustment costs co-depend on other model parameters. Finally, Section 8 concludes. An appendix provides more detailed information concerning the calibration of fixed adjustment costs and the numerical solution procedure.

1.2 Related Literature

---

[4]Berger and Vavra (2010) show a similar result for variations in risk in a sticky price model. When idiosyncratic risk increases, then price setting becomes more flexible in their model.

A number of other recent papers have shown applications in which the non-convexity of plant-level decisions does matter in shaping aggregate dynamics. Bachmann et al. (2010) show that lumpiness in the capital adjustment decision helps explain the pro-cyclicality of the aggregate investment response to TFP shocks in U.S. data. Fiori (2012) introduces a two-sector RBC model in which non-convex capital adjustment costs in the investment goods producing sector allow the model to replicate a hump-shaped response of aggregate investment to productivity shocks. Given that a two-country model can also be interpreted as a model with two large sectors, our result reinforces the importance of fixed capital adjustment costs in inter-sectoral reallocation. Bachmann and Ma (2012) solve a closed economy model in which aggregate savings can also take the form of inventory accumulation. They show that a fixed cost to restocking inventories and fixed capital adjustment costs have mutually reinforcing effects on aggregate dynamics. In Sustek (2011) plants face non-convex costs to using different forms of shift work. In consequence, output volatility is reduced and becomes countercyclical. Most similar in spirit to our work is a paper by Miao and Wang (2011) developed parallel which derives conditions under which the aggregate dynamics in a model where firms face both fix and convex capital adjustment costs can be represented by a model where only convex costs are present and Tobin's Q is a sufficient statistic for describing investment dynamics. Their paper demonstrates that the form of the convex cost function in the isomorphic representation depends on the size of non-convex adjustment costs. In contrast to our approach, Miao and Wang (2011) assume constant returns to scale in production at the micro level and then use convex capital adjustment costs alongside fixed ones to avoid a degenerated firm problem. By dropping the constant returns to scale assumption, we can show that the curvature of the production function and the distribution of idiosyncratic profitabilities enters in the convex adjustment cost representation.

1.3 The Model

We model a world economy composed of two countries **Home** and **Foreign** (where necessary, country specific variables will be distinguished by the superscripts H and F respectively). Each country is populated by a representative household and a continuum of firms producing an intermediate good which differs between the two

countries. Competitive final goods producers use these inputs to produce a local composite good used for investment and consumption. There exists a complete set of contingent claims which ensures international consumption risk sharing. The challenge in solving the model lies in the solution of the intermediate goods producers' problem in both countries. Here, we closely follow Khan and Thomas (2008) and can therefore be brief in referring the reader to these papers for further explanations. Our focus will instead lie on the necessary adaptations to the solution method for it to be applicable to our model.

1.3.1 Households

There is a continuum of identical households in both economies who work and consume and who have access to complete international asset markets. Their felicity function is defined on the consumption of their local consumption good and in (indivisible) labor, which they supply on the local labor market:

$$U(C_j, N_j) = \log(C_j) - AN_j, \tag{1.1}$$

where C_j denotes consumption in country j and N_j the households labor supply in country j.

Households hold wealth as one-period shares in plants denoted by the measure λ_j. Given the prices they receive for their current shares $\rho_0^j(\epsilon, l; \hat{z}, m)$ and the real wage rate W_j / P_j^C, households choose current consumption C_j, labor effort N_j and the number of new shares $\lambda_j'(\epsilon, k)$ to buy at prices $\rho_1^j(\epsilon', k'; \hat{z}, m)$. $(\hat{z}, m)$ summarizes the aggregate state and is defined further below. Households maximize the expected discounted present value of intertemporal utility:

$$W(\lambda_j; \hat{z}, m) = \max_{C_j, N_j, \lambda'} \left[U(C_j, N_j) + \beta \mathbb{E}\left[W(\lambda'; \hat{z}', m') \right] \right]$$

subject to

$$C_j + \int \rho_1^j(\epsilon', k'; \hat{z}, m) \, d(\epsilon' \times k') \leq W_j / P_j^C + \int \rho_0^j(\epsilon, k; \hat{z}, m) \, d(\epsilon \times k)$$

Let λ be the Lagrangian multiplier on the household's intertemporal budget con-

straint. We obtain the first-order conditions with respect to consumption

$$\lambda P_j^C = U_C(C_j, N_j) = \frac{1}{C_j},\qquad(1.2)$$

where P_j^C is the current price of the final consumption good in country j. With respect to labor we obtain

$$\lambda W_j = -U_N(C_j, N_j)\qquad(1.3)$$

where W_j is the nominal wage in country j. Combining this with the first order condition on consumption and plugging in the assumed functional forms we obtain

$$W_j/P_j^C = AC_j.$$

Note that with complete international financial markets the resulting allocation must be efficient. This, together with the assumption of symmetric initial endowments, implies equal Pareto weights and hence the risk sharing condition $\frac{U_C(C_F, N_F)}{P_F^C} = \frac{U_C(C_H, N_H)}{P_H^C}$.

1.3.2 Final Goods Producers

In both countries, consumption and investment use a composite good produced by a competitive final goods producer. The final goods producer in country j combines intermediate goods $X_j^{H,F}$, where X_j^H (X_j^F) are intermediate goods produced in the Home (Foreign) country and used in country j. Final consumption goods in country j are produced using the constant returns to scale production function:

$$G_j(X_j^H, X_j^F) = \left[\omega^{\frac{1}{\sigma}} X_j^{j\,\frac{\sigma-1}{\sigma}} + (1-\omega)^{\frac{1}{\sigma}} X_j^{\neg j\,\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \ j = H, F,$$

where ω measures the home-bias or importance of local intermediate goods for the final goods production, $\neg j$ denotes the respective other country.

Final goods markets are competitive. Let P_j^X be the price of the intermediate good produced in country j. Then final goods producers solve the cost minimization problem:

$$\min_{X_j^H, X_j^F} P_H^X X_j^H + P_F^X X_j^F\qquad(1.4)$$

s.t.

$$G_j(X_j^H, X_j^F) = 1$$

This cost minimization and perfect competition imply that the price of the consumption good P_j^C in country j is given by

$$P_j^C = \left[\omega \left(P_j^X \right)^{1-\sigma} + (1 - \omega) \left(P_{\neg j}^X \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}.$$

Using the Home country intermediate good as a numeraire and normalizing P_H^X to one we obtain as prices for the final consumption good:

$$P_H^C(\tau) = [\omega + (1 - \omega)\tau^{1-\sigma}]^{\frac{1}{1-\sigma}}$$
$$P_F^C(\tau) = [\omega\tau^{1-\sigma} + (1 - \omega)]^{\frac{1}{1-\sigma}} = \tau P_H^C \left(\tau^{-1} \right)$$

where $\tau = \frac{P_F^X}{P_H^X}$ denotes the terms of trade.

1.3.3 Intermediate Goods Producers

The more complicated planning problem is the one of the intermediate goods producer. In both countries, intermediate goods producers employ predetermined capital and labor and produce according to a Cobb-Douglas decreasing-returns-to-scale production function

$$y = z\epsilon(k^\chi n^{1-\chi})^{\frac{1}{\eta}}$$

where z is stochastic total factor productivity common to all firms in the country and ϵ is firm-specific productivity. A way of reading the decreasing returns-to-scale assumption is as constant returns-to-scale in production with capital share χ *cum* monopolistic competition in intermediate goods, where firms earn a mark-up of η on their sales. This implies revenue elasticities of capital $\theta = \frac{1}{\eta}\chi$ and $\nu = \frac{1}{\eta}(1 - \chi)$ of labor.

We assume that **Home** and **Foreign** technology follow the joint process

$$\begin{bmatrix} \log\left(z_{H,t}\right) \\ \log\left(z_{F,t}\right) \end{bmatrix} = \begin{bmatrix} \rho_1 & \rho_2 \\ \rho_2 & \rho_2 \end{bmatrix} \begin{bmatrix} \log\left(z_{H,t}\right) \\ \log\left(z_{F,t}\right) \end{bmatrix} + \begin{bmatrix} \hat{\nu}_1 \\ \hat{\nu}_2 \end{bmatrix}, \text{with} \begin{bmatrix} \nu_1 \\ \nu_2 \end{bmatrix} \sim \text{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_{\epsilon_1}^2 & \sigma_{\epsilon_1\epsilon_2} \\ \sigma_{\epsilon_1\epsilon_2} & \sigma_{\epsilon_2}^2 \end{bmatrix} \right)$$

As in Backus et al. (1994) and Heathcote and Perri (2002), technology spillovers are assumed symmetric with $\rho_1, \rho_2 > 0$. This allows for an important simplification which results in the elimination of one state variable in our numerical solution algorithm. The focus of our study is on the excess volatility of the investment series in national economies caused by international capital reallocation in response to productivity differentials between the two countries. From now on, we therefore focus on relative technology $\log\left(\hat{z}_t\right) = \log\left(z_{H,t}\right) - \log\left(z_{F,t}\right)$ only. $\hat{z}_t$ follows an AR(1) process

$$\log\hat{z}_t = \rho\log\hat{z}_{t-1} + \hat{v}$$

with

$$\hat{v} \sim N(0, \underbrace{\sigma_{\epsilon_1}^2 + \sigma_{\epsilon_2}^2 - 2\sigma_{\epsilon_1\epsilon_2}}_{\equiv \sigma^2}).$$

and $\rho = \rho_1 - \rho_2$. We set $z_t^H = \hat{z}_t$ and $z_t^F = \hat{z}_t^{-1}$ and discretize $\hat{z}_t$ into a 13-state Markov process using Tauchen's (1986) method. The idiosyncratic profitability process follows a 15-state Markov process which is an approximation to a continuous AR(1) process for log profitability with Gaussian innovations.

Each firm produces an intermediate good but needs to raise capital in terms of the national composite good. At the beginning of a period a firm receives an idiosyncratic i.i.d. fixed adjustment cost draw $\xi \geq 0$, which is denominated in units of labor. It is drawn from a distribution $G : [0, \bar{\xi}] \to [0, 1]$. This distribution is common to all firms:

$$G \sim U(0, \bar{\xi}).$$

We initially denote the firm's planning problem in units of the local capital-consumption good. The intra-period timing is a follows: After having observed innovations to aggregate and idiosyncratic productivity and its adjustment cost draws, the firm optimally adjusts labor, produces output and harvests flow profits. Afterwards, the firm decides whether to pay the adjustment cost and adjust its capital stock to the current target level or whether to exercise its option to wait and see and

let its capital depreciate. Upon investment, the firm incurs a fixed cost of $w\xi$, where w is the current real wage rate defined in local intermediate goods $w_j := W_j/P_j^X$. Capital depreciates at rate δ. Table 1.1 summarizes the evolution of the firm's capital stock (in efficiency units) between two consecutive periods, from k to k'.

Table 1.1: Evolution of the Firm's Capital Stock

	Fixed cost paid	$\gamma k'$
$i \neq 0$:	$w\xi$	$(1-\delta)k + i$
$i = 0$	0	$(1-\delta)k$

Notes: This table summarizes the evolution of the firm's capital stock conditional on the binary choice between investing and not investing.

The distributions of firms over capital and idiosyncratic productivity states (ϵ, k) in the two countries are summarized using the probability measures μ^H and μ^F. They are sufficient to describe differences between firms and their evolution over time given the i.i.d. nature of the adjustment costs. Define $m \equiv [\mu^H(k, \epsilon), \mu^F(k, \epsilon)]$ so that the aggregate state of the economy is described by $(\hat{z}, m)$. The distributions evolve over time according to a mapping Γ from the current aggregate state $m' = \Gamma(\hat{z}, m)$ which will be defined below.

Let $v_j(\epsilon, k, \xi; \hat{z}_j, m)$ denote the expected discounted value - measured in local consumption goods - of a firm in country j that is in idiosyncratic state (ϵ, k, ξ), given the aggregate state $(\hat{z}, m)$. Its expected value prior to drawing its adjustment cost draw is then given by:

$$\bar{v}_j(\epsilon, k; \hat{z}^j, m) = \int_0^{\overline{\xi}} v_j(\epsilon, k, \xi; \hat{z}^j, m) G(d\xi) \tag{1.5}$$

The dynamic programming problem of a firm in country j is described by:

$$v(\epsilon, k, \xi; \hat{z}^j, m)_j = cf_j + \max \left\{ v_j^{dep}, \max_{k'} \left(-ac_j + v_j^{adj} \right) \right\}, \tag{1.6}$$

where cf are flow profits, v_j^{dep} is the firm's continuation value if it chooses inaction and lets its capital depreciate, and v_j^{adj} the continuation value, net of adjustment costs, if the firm chooses to invest and adjust its capital stock to the current target

level. These functions are given by:

$$cf_j = \max_n \big[\hat{z}^j \epsilon (k^\chi n^{1-\chi})^{\frac{1}{\eta}} - w_j(\hat{z}, m) n \big] \frac{P_j^X}{P_j^C} \tag{1.7a}$$

$$v_j^{dep} = \mathbb{E}\big[d_j\left(\hat{z}', m'\right) \bar{v}(\epsilon', \frac{(1-\delta)}{\gamma} k; \hat{z}', m') \big] \tag{1.7b}$$

$$ac_j = \xi w_j(\hat{z}, m) \frac{P_j^X}{P_j^C} \tag{1.7c}$$

$$v_j^{adj} = -i + \mathbb{E}\big[d_j\left(\hat{z}', m'\right) \bar{v}(\epsilon', k'; \hat{z}', m') \big] \tag{1.7d}$$

where both expectation operators average over next period's realization of the average and idiosyncratic productivity states, conditional on this period's values, and we recall that $i = \gamma k' - (1-\delta)k$. The stochastic discount factor of the local representative household is $d_j\left(\hat{z}', m'\right) = \beta \frac{U_C(C_j(\hat{z}', m'), N_j(\hat{z}', m'))}{U_C(C_j(\hat{z}, m), N_j(\hat{z}, m))}$.

We can eliminate the stochastic discount factor by rephrasing the firm's value function in terms of utils (more details can be found in Khan and Thomas, 2008). This allows us compute equilibrium by solving a single Bellman equation that combines the plant-level optimization problem in equations (1.5)-(1.7) with the household first order conditions (1.2)-(1.3). Given that investment uses the composite consumption good, we define its price relative to the firm's output as (P_H^X has been normalized to one):

$$q_j(\tau) = \begin{cases} \frac{P_H^C}{P_H^X} = P_H^C(\tau) & for\ j = \mathbf{H} \\ \frac{P_F^C}{P_F^X} = \frac{P_F^C}{\tau} = P_H^C(\tau^{-1}) & for\ j = \mathbf{F} \end{cases}. \tag{1.8}$$

Denoting the marginal utility of consumption by $\varrho_j \equiv U_C^j(C_j, N_i)$, we obtain due to efficient risk sharing between the economies:

$$\varrho_j(\tau, C_H) = \begin{cases} U_C^H & for\ j = \mathbf{H} \\ U_C^F \frac{P_F^X}{P_F^C} q_F(\tau) = \varrho_H(\tau, C_H) \frac{P_F^X}{P_H^C} q_F(\tau) = \varrho_H(\tau, C_H) \tau \frac{q_F(\tau)}{q_H(\tau)} & for\ j = \mathbf{F} \end{cases} \tag{1.9}$$

U_C^H is the marginal utility of consumption in the **Home**-economy. Importantly, we can express the marginal utility of the foreign household as a function of home marginal utility and terms of trade.

Let $V_j(\epsilon, k, \xi; \hat{z}_i, m) = v_j U_C^j(C_i, N_i)$ now denote the expected discounted value in

utils of the respective representative household of a firm. This is:

$$V_j(\epsilon, k, \xi; \hat{z}^j, m) = CF_j + \max\left\{V_j^{dep}, \max_{k'}\left(-AC_j + V_j^{adj}\right)\right\}, j = H, F \qquad (1.10)$$

with the components defined analogously to before. These are given by:

$$CF_j = \max_n\left[z_j(\hat{z})\,\epsilon(k^\chi n^{1-\chi})^{\frac{1}{\eta}} - w_j(\hat{z}, m)n\right]\frac{\varrho_j(\tau, C_H)}{q_j(\tau)} \qquad (1.11a)$$

$$V_j^{dep} = \beta\mathbb{E}\left[\bar{V}_j(\epsilon', \frac{1-\delta}{\gamma}k; \hat{z}', m')\right] \qquad (1.11b)$$

$$AC_j = \xi w_j(\hat{z}, m)\frac{\varrho_j(\tau, C_H)}{q_j(\tau)} \qquad (1.11c)$$

$$V_j^{adj} = -i\varrho_j(\tau, C_H) + \beta\mathbb{E}\left[\bar{V}_j(\epsilon', k'; \hat{z}', m')\right] \qquad (1.11d)$$

$$\bar{V}_j(\epsilon, k; \hat{z}_j, m) = \int_0^{\bar{\xi}} V_j(\epsilon, k, \xi; \hat{z}_j, m)G(d\xi). \qquad (1.11e)$$

Given (ϵ, k, ξ) and equilibrium prices $w_j(\hat{z}, m)$, $\varrho_j\left[\tau(\hat{z}, m), C_H(\hat{z}, m)\right]$ and $q\left[\tau(\hat{z}, m)\right]$ the plant chooses employment and whether to invest or let its capital depreciate.[5] Denote as $N_j = N(\epsilon, k; \hat{z}, m)$, $K_j = K(\epsilon, k, \xi; \hat{z}, m)$ the intermediate firm policy functions. Since capital is predetermined, the optimal employment decision is independent of the current adjustment cost draws. We denote the total intermediate goods output in country j by Y_j.

1.3.4 Recursive Equilibrium

A *recursive competitive equilibrium* for this economy is completely described by the set:

$$\{w_j, \varrho_H, \tau, V_j, N_j^D, N_j^S, K_j, C_j, X_j^H, X_j^F, \Gamma\}_{j=H,F}$$

that satisfy

1. *Firm optimality*: Taking w, τ, λ and Γ as given, V_j satisfy (1.10)-(1.11e) and N_j^D, K_j are the associated policy functions.

---

[5]Note that the problem is symmetric for both countries, which can be exploited to save computation time.

2. *Household optimality*: Taking w, τ and λ as given, the households' consumptions C_j and labor supplies N_j^S satisfy (1.2) and (1.3).

3. X_j^H, X_j^F solve (1.4).

4. *Labor markets clearing*:

$$N_j^S(\hat{z}, m) = \int \left\{ N_j^D(\epsilon, k; \hat{z}, m) + \int \xi \mathbb{I} \left[\frac{1-\delta}{\gamma} k - K_j(\epsilon, k, \xi; \hat{z}, m) \right] dG \right\} d\mu_j$$

where $\mathbb{I}(x) = 0$ if $x = 0$; $\mathbb{I}(x) = 1$ otherwise.

5. *Final goods markets clearing*:

$$C_j + \int \int_0^{\overline{\xi}} \left[\gamma K_j(\epsilon, k, \xi; \hat{z}, m) - (1-\delta)k \right] dG d\mu^H = G^j(X_j^H, X_j^F)$$

6. *Intermediate goods markets clearing*:

$$\sum_{l=H,F} X_l^j = Y^j$$

7. *Model consistent dynamics*: The evolution of the cross-sectional distributions that characterize the economy in both countries, $m' = \Gamma(\hat{z}, m)$, is induced by $\{K_j(\epsilon, k, \xi; \hat{z}, m)\}_{j=H,F}$ and the exogenous processes for $\hat{z}$ and ϵ.

1.4 Numerical Solution

The aggregate state contains two infinite dimensional objects: The distributions of intermediate producers in both countries over capital and idiosyncratic productivity states. Following Krusell and Smith (1998, 1997) we approximate those distributions by a finite number of distributional moments. Let $\hat{m} = [\overline{k}_H, \overline{k}_F]$ denote our approximate aggregate state and $\hat{\Gamma}(\hat{m}, \hat{z}))$ denote its law of motion, such that $\hat{m}' = \hat{\Gamma}(\hat{m}, \hat{z})$. In our applications, first moments over capital, $\overline{k}_H$ and $\overline{k}_F$ turn out to contain sufficient information to accurately forecast prices. A number of accuracy tests including R^2s are reported in the numerical appendix. We specify simple log-linear rules to describe price forecasts for ϱ and τ and the evolution of capital stocks $\hat{\Gamma}$. Instead of

laws of motion for $\log(\overline{k}_H)$ and $\log(\overline{k}_F)$, we work with a rotated version of the system in the sum of of log-world capital and its difference[6].

We impose some economic structure to minimize the effect of simulation and estimation uncertainty inherent in a Monte-Carlo method such as Krusell and Smith's (1998) algorithm (see DenHaan, 1997). For this reason, exploiting symmetry, we assume that world capital depends only on previous world capital (and not on its distribution over countries) and the difference in capital stocks between countries depends only on previous differences and not on world capital stocks. Moreover, we impose log linear effects of aggregate productivity on the dynamics of capital stocks as well as prices:[7]

$$\left[\log(\overline{k}_H) + \log(\overline{k}_F)\right]' = \alpha_0^{world} + \alpha_1^{world}\left[\log(\overline{k}_H) + \log(\overline{k}_F)\right] \tag{1.12a}$$

$$\left[\log(\overline{k}_H) - \log(\overline{k}_F)\right]' = \alpha_1^{\Delta}\left[\log(\overline{k}_H) - \log(\overline{k}_F)\right] + \alpha_2^{\Delta}\log(\hat{z}) \tag{1.12b}$$

$$\log(\varrho_H) = \alpha_0^{\varrho} + \alpha_1^{\varrho}\log(\overline{k}_H) + \alpha_2^{\varrho}\log(\overline{k}_F) + \alpha_3^{\varrho}\log(\hat{z}) \tag{1.12c}$$

$$\log(\tau) = \alpha_1^{\tau}\left[\log(\overline{k}_H) - \log(\overline{k}_F)\right] + \alpha_3^{\tau}\log(\hat{z}). \tag{1.12d}$$

The solution algorithm consists of two steps which are repeated successively until

---

[6]One may obtain the specification below by pre-multiplying the system

$$\left[\begin{array}{c} \log\left(\overline{k}_H\right)' \\ \log\left(\overline{k}_F\right)' \end{array}\right] = \left[\begin{array}{c} \alpha_0^H \\ \alpha_0^F \end{array}\right] + \left[\begin{array}{cc} \alpha_1^H & \alpha_2^H \\ \alpha_1^F & \alpha_2^F \end{array}\right]\left[\begin{array}{c} \log\left(\overline{k}_H\right) \\ \log\left(\overline{k}_F\right) \end{array}\right] + \left[\begin{array}{c} \alpha_1^H \\ \alpha_2^F \end{array}\right]\log\left(\hat{z}\right)$$

with the matrix

$$\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}.$$

Symmetry of Home and Foreign country implies that $\alpha_0^H = \alpha_0^F$; $\alpha_1^H = \alpha_2^F$; $\alpha_1^F = \alpha_2^H$ and $\alpha_3^H = -\alpha_3^F$. It follows that

$$\alpha_0^{world} = 2\alpha_0^H$$
$$\alpha_0^{\Delta} = 0$$
$$\alpha_1^{world} = \alpha_1^H + \alpha_1^F = \alpha_2^H + \alpha_2^F$$
$$\alpha_1^{\Delta} = \alpha_1^H - \alpha_1^F = \alpha_2^H - \alpha_2^F$$
$$\alpha_2^{world} = 0$$
$$\alpha_2^{\Delta} = 2\alpha_2^H$$

[7]We checked whether these restriction we imposed actually restrict the dynamics, by estimating versions without the imposed restrictions and check whether the restrictions would be rejected by a Wald test. We found that the imposed restrictions would not be rejected in equilibrium.

parameters of the aggregate laws of motion converge. Using an initial guess for the parameters of the aggregate laws, we solve the dynamic programming problem posed by equations (1.10) - (1.11e) which becomes computationally feasible once the approximate aggregate state is used. A number of the problem's features facilitate the solution considerably. First, the firms' employment decision is static and independent of its investment adjustment cost draw so that it can be maximized out using the respective first order condition:

$$N(\epsilon, k; \hat{z}, \hat{m}) = \left(\frac{w}{\hat{z}\epsilon\nu k^\theta} \right)^{\frac{1}{\nu-1}}$$

Second, the optimal capital stock chosen conditional on adjustment is independent of the firm's current individual capital stock. This optimization problem therefore needs to be solved only once for each point on the aggregate state grid. Given that adjustment is costly and that it always holds that $V_j^{adj}(\epsilon; \hat{z}, \hat{m}) \geq V_j^{dep}(k, \epsilon; \hat{z}, \hat{m})$, the value of the adjustment cost draw, $\hat{\xi}(\epsilon, k; \hat{z}, \hat{m})$, at which the firm is just indifferent between adjusting and exercising its option to wait and see (i.e. letting its capital depreciate) is given by:

$$\hat{\xi}_j(\epsilon, k; \hat{z}, \hat{m}) = \frac{q_j\left[\tau(\hat{z}, \hat{m})\right]\left[V_j^{adj}(\epsilon; \hat{z}, \hat{m}) - V_j^{dep}(k, \epsilon; \hat{z}, \hat{m})\right]}{\varrho_j(\hat{z}, \hat{m})w\left[\tau(\hat{z}, \hat{m})\right]} \tag{1.13}$$

Denoting the target capital stock to which a firm with idiosyncratic productivity ϵ in country j adjusts in the absence of frictions by $k_j^*(\epsilon; \hat{z}, \hat{m})$ allows us to compute the firms' second policy function determining investment:

$$k' = K_j(\epsilon, k, \xi; \hat{z}, \hat{m}) = \begin{cases} k_j^*(\epsilon; \hat{z}, \hat{m}) & \text{if } \xi \leq \hat{\xi}_j(\epsilon, k; \hat{z}, \hat{m}), \\ (1-\delta)k/\gamma & \text{otherwise.} \end{cases} \tag{1.14}$$

Given firm policy functions, we simulate the economy in the second step. In order to more efficiently exploit parallel computing resources, instead of using one long draw of relative productivities, we generate observations for aggregate variables using several shorter draws of $\hat{z}_t$. During the simulation, market clearing values of ϱ and τ are computed exactly. This procedure generates a total of T=4800 observations of $\{\hat{m}_t, \varrho_t, \tau_t\}$ which we use then to update the α-coefficients in the aggregate laws of motion by simple OLS regression. We iterate these steps until an F-Test

finds all parameter estimates from two successive steps statistically indistinguishable. Upon convergence, we have obtained the Krusell-Smith recursive equilibrium of our economy for a given set of parameters.[8]

1.5 Parameter Choices

The model parameters to calibrate are relatively standard. They involve the discount factor, β, the disutility of labor, A, the parameters of the production function, χ and η, the law of motion for aggregate productivity, the substitution elasticity in final goods production, σ, as well as the home bias parameter, ω. The parameters somewhat less standard are of the idiosyncratic productivity process, ρ and σ_ϵ and the adjustment cost parameter $\bar{\xi}$.

1.5.1 Open Economy Parameters

The substitution elasticity, σ, between intermediate input goods in the production function for the final consumption goods is set to 1.5. A common range in the open economy literature is $[1, 2]$. A recent estimate for the bilateral productivity process come from Heathcote and Perri (2002) who use data for the US and the rest of the world as other economy. Their estimates imply values for our process of log-relative TFP of $\rho = .945$ and $\sigma_\nu^2 = 0.0087$. For our baseline, we set $\omega = 0.7$ which is about average for OECD countries.

1.5.2 Parameters for the National Economies

The model period is a quarter. Our choice for $\beta = 0.99$ implies an annual interest rate of 4 percent. We set $A = 1.852$ to match a steady state labor supply of 1/3.

---

[8]Our abstraction from variations in world-TFP implies a near constant world capital stock. In the stochastic steady state of the system, $\bar{k}_H$ and $\bar{k}_F$ are therefore almost perfectly negatively correlated. This poses a serious problem for the estimation of our Krusell-Smith rules if using only observations from the stochastic steady state.

We solve this by first letting the system settle into its stochastic steady state during fifty initial periods, then we lower the capital stock of every firm by 20% and observe the adjustment path of the economies back to the steady state. This gives enough additional information to identify also the law of motion for world capital. More details on the solution procedure can be found in the appendix.

We set the coefficients of the revenue function $\chi = 1/3$ and $\eta = 4/3$ which implies a mark-up of 33 % and the implied output elasticities of labor and capital are $\nu = 1/2$ and $\theta = 1/4$ respectively, as in Bloom (2009). This is also close to the empirical estimate in Bachmann and Bayer (2011a) for manufacturing. We calibrate γ to imply a technological growth rate of 1.4% p.a. and depreciation to 9.4 % p.a and assume for idiosyncratic productivity

$$\log \epsilon' = \rho_\epsilon \log \epsilon + \sigma_\epsilon u \tag{1.15}$$

where $u \sim N(0,1)$ and set $\rho_\epsilon = 0.98$ and $\sigma_\epsilon = 0.0459$ in line with Bachmann and Bayer (2011a), who report annual cross-sectional firm level data.

Our baseline specification of $\bar{\xi} = 1.7$ matches a cross-sectional skewness in annual plant investment rates of 2.2 reported in Bachmann and Bayer (2011a) and also matches roughly the fraction of spike-adjusters (13%) they report for manufacturing.[9] This estimate is well in line with estimates by Bachmann et al. (2010), Caballero and Engel (1999), Cooper and Haltiwanger (2006), or Bloom (2009). In comparing the point estimates, one should keep in mind, however, that we calibrate to the quarterly frequency. As adjustment costs are i.i.d. random variables, this implies that actual adjustment costs paid are on average much lower. The estimate of $\bar{\xi}$ itself therefore cannot be directly compared to papers calibrated to an annual frequency, but the implied average resources spent on adjustment can be. Our adjustment cost estimate implies that conditional on adjustment roughly 29% of the adjusting firm's annual labor force is used for installing capital (worth 14% of its annual output) which is in line with the estimates cited above, see Table 3 in Bachmann et al. (2010). Table 1 below summarizes our parameter choices[10] .

---

[9]If we had matched the fraction of spike adjusters in LRD (19%), a somewhat smaller estimate of $\bar{\xi} \approx 0.4$ would have been obtained. Note, however, that even at this smaller level, fixed adjustment costs significantly dampen aggregate fluctuations in investment. On the other hand, taking additionally unit aggregation into account would increase the estimate of $\bar{\xi}$ instead. Moreover, note in comparing our $\bar{\xi} = 1.7$ to other calibrations the quarterly frequency to which we calibrate, which means that firms have four draws of ξ in a year and can choose the smallest one.

[10]At first sight, our calibrated maximum adjustment costs of 1.7 look much larger than the 0.0083 in Khan and Thomas (2008). Indeed, one might wonder whether the effects we report below are simply the effect of higher adjustment costs and not of introducing the open economy dimension. There are two main arguments why this is not likely to be the case. For one, as we report in Table 1.4, the dampening effect of fixed adjustment costs on aggregate investment is a decreasing

Table 1.2: Parameters of the Baseline Calibration

Preferences		
	HH discount rate	$\beta = 0.99$
	Disutility of labor	$A = 1.852$
Firm production		
	Output elasticity of capital	$\chi = 1/3$
	Mark-up in intermediate goods markets	$\eta = 4/3$
	Implied labor revenue elasticity	$\nu = 1/2$
	Implied capital revenue elasticity	$\theta = 1/4$
	Rate of technological progress	$\gamma = 1.0035$
	Depreciation rate	$\delta = 0.0235$
	Persistence in id. prod.	$\rho_\epsilon = 0.98$
	Std. of innovations to id. prod.	$\sigma_\epsilon = 0.0459$
Open Economy		
	EOS in composite good	$\sigma = 1.5$
	Persistence in relative TFP	$\rho = 0.945$
	Variance of innovations to TFP	$\sigma_v^2 = 0.0087$
	Import share in consumption	$1 - \omega = 0.3$

1.6 Results

We obtain three sets of results from the simulation of our model. First, under our baseline specification, fixed adjustment costs dampen capital reallocation between economies. This effect is the stronger, the more open the economies are to trade and vanishes as $\omega \to 1$, i.e. the economies become more similar to a closed economy, so that the main Khan and Thomas (2003, 2008) result obtains. Second, the aggregate behavior of the economy with fixed adjustment costs can be represented by an economy with quadratic adjustment costs. Third, the size of the quadratic costs that yield almost identical aggregate behavior co-depends on the curvature of the revenue function w.r.t. capital and on the size of idiosyncratic risk σ_ϵ, but does not depend

---

function of the openness to trade of the two economies. When approaching the limiting case of two closed economies the effect disappears. Second, as a robustness exercise, Khan and Thomas also solve a version of their model with much larger fixed costs of about .21 and still find only very small effects on aggregate investment. At an annual frequency, this model specification is actually relatively close to ours. For example, the probability of receiving a fixed cost draw of at most .05 in any year in their specification is 24.1 percent. In our case, given the quarterly frequency of the model, the likelihood of that event is 11.26 percent so only about half as big.

Table 1.3: Cyclical Properties

	Frictionless Model	Non-convex cost Model	Quadratic cost Model
Standard deviations in %			
Output	1.826	1.775	1.777
Standard Deviation Relative to Output			
Investment	6.327	4.426	4.417
Consumption	0.188	0.208	0.207
Employment	0.581	0.495	0.498
Exports	0.869	0.506	0.505
Imports	0.870	0.507	0.505
NX	0.693	0.360	0.378
ToTs	0.939	1.043	1.037
Correlation with Output			
Investment	0.860	0.915	0.915
Consumption	0.940	0.985	0.985
Employment	0.961	0.983	0.984
Exports	-0.226	0.354	0.343
Imports	0.225	-0.354	-0.343
NX	-0.679	-0.656	-0.658
ToTs	0.940	0.985	0.985
Persistence			
Output	0.694	0.696	0.696
Investment	0.581	0.627	0.627
Consumption	0.816	0.746	0.746
Employment	0.617	0.655	0.655
Exports	0.747	0.944	0.943
Imports	0.747	0.944	0.943
NX	0.601	0.668	0.667
ToTs	0.816	0.746	0.746

Notes: All statistics are averages from 100 simulations of the economy over 200 quarters. All data are in logs except for net-exports (NX) and HP(1600)-filtered. Net-exports are relative to Output.

on openness or the presence of an investment tax-credit. This third set of results are discussed in detail in Section 1.7.

1.6.1 Fixed Adjustment Costs Matter in the Open Economy

We first asses whether non-convex adjustment costs at the firm level matter at all in our setting in shaping aggregate investment dynamics. In other words, we check whether our initial intuition that the potential to achieve consumption smoothing via movements in the current account brings back a role for fixed adjustment costs to capital in shaping the business cycle.

We compare the aggregate behavior of our economy to a frictionless reference and a version with partial adjustment that introduces quadratic adjustment costs of the form $\frac{\phi}{2} w_t \left(\frac{k_{t+1} - k_t}{k_t} \right)^2$. We set $\phi = 3.08$, details for this parameter are deferred to the next section[11]. Table 1.3 summarizes volatilities, correlation and persistence of the key variables for the three specifications. Two important results can be read from this table. First, unlike in the closed economy, see Khan and Thomas (2008), non-convex adjustment costs to capital at the plant level have an effect on the dynamic behavior of macro aggregates in our open economy setup. Their main effect is to dampen the volatility of aggregate investment relative to output by reducing impact responses of investment to productivity innovations and to mildly increase its persistence. Moreover, they change the cyclicality of exports and imports. When comparing the reported moments to empirical moments of international business cycles, one should keep in mind that we have basically abstracted from world tfp movements by focusing on movements in relative productivity. That, for example, explains the relative consumption volatility of only .208 in our baseline specification. When evaluating the reported results, the focus should therefore lie on the quantitative differences between the columns.

Given the strong difference in results for the open and closed economy, we explore whether our specification nests the closed economy result as a limiting case of an economy with perfect home bias in consumption. We therefore solve our model for various values of the openness parameter, ω, ranging from 0.55 to 0.9. Table 1.4

---

[11] We chose capital adjustment costs to be labor denominated to keep as close a possible to the baseline fixed cost specification. We also experimented with output denominated costs. This did not qualitatively alter results. Quantitatively, calibrated quadratic adjustment costs did not change much either.

Table 1.4: Investment Volatility relative to output for different openness parameter values

	$\omega = 0.55$	$\omega = 0.6$	$\omega = 0.7$	$\omega = 0.85$	$\omega = 0.9$
Frictionless	8.121	7.395	6.327	5.103	4.799
$\bar{\xi} = 1.7$	2.981	3.721	4.426	4.413	4.302
Percentage difference	63.3%	49.7%	30.0%	13.5%	10.4%

Notes: All other parameters are kept at their baseline value.

displays volatility of investment relative to output for different parameter values for a frictionless economy and compares it to one featuring non-convex adjustment costs of $\bar{\xi} = 1.7$. Clearly, the effect of fixed adjustment costs on investment volatility is stronger, the more open an economy is to trade. Vice versa, at $\omega = 0.9$ the effect has weakened substantially. The closed economy setting of Khan and Thomas (2008) then obtains as a limiting case. When our economies are completely closed to trade, fixed adjustment costs once again wash out when looking at aggregate statistics.

Coming back to Table 1.3, there is a second observation. The effects of microeconomic fixed adjustment costs, when looking at aggregate dynamics, are almost indistinguishable from our quadratic adjustment cost specification. Figure 1.1 puts this in a graphical version and displays the demeaned time series of investment for the same draw of aggregate productivity shocks for both the quadratic cost model and the fixed adjustment cost model calibrated to investment at the micro level.[12] In fact, the investment series of the quadratic cost model (partial adjustment) and the fixed cost model perfectly align. More technically one could say that we find "approximate representation" beyond "approximate aggregation" in the sense of Krusell and Smith (1998). In the following section we investigate this similarity of fixed adjustment costs and quadratic ones in more detail.

---

[12]Both series slightly differ in their respective means, because the stochastic costs of capital reallocation introduce a precautionary motive into the aggregate investment decision in the fixed cost model. Meanwhile, we account for firm heterogeneity only implicitly in the quadratic cost model by simulating one firm whose TFP is adjusted upwards to account for higher overall productivity due to log-normally distributed idiosyncratic productivities.

Figure 1.1: Time-Series of Investment in Fixed Cost and Partial Adjustment Model
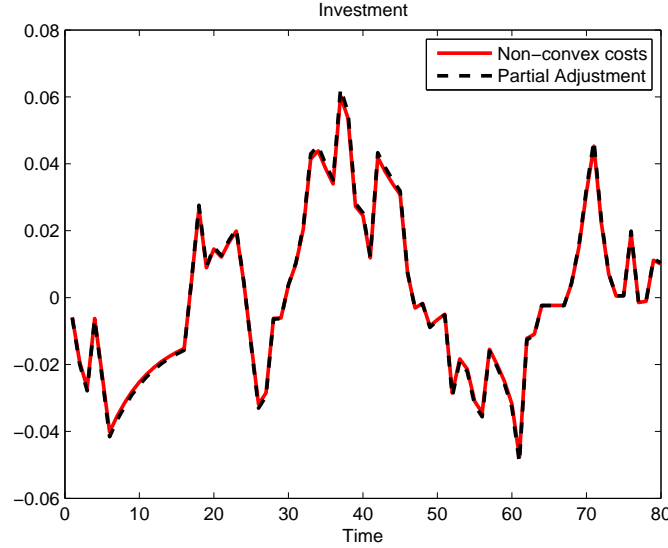
Notes: The figure shows series of aggregate investment generated by the fixed adjustment cost model and a matched quadratic adjustment cost specification simulated over the same draw of aggregate productivities. Both series have been demeaned.

1.7 "Fundamentalness" of Quadratic Capital Adjustment Costs

As discussed in the introduction, the problem of excess volatility in investment is widespread in the literature on open economy business cycle models. The problem is most severe when countries trade perfectly substitutable goods as in Backus et al. (1992), but it occurs in varying degrees of severity also in settings more comparable to ours. The introduction of convex or more specifically quadratic adjustment costs at the aggregate level has become a standard kludge to bring the model closer to the data.

We noted in our first take on the results that the aggregate dynamics in a setting with microeconomic lumpiness in investment look almost indistinguishable from one where a representative stand-in firm faces quadratic costs to capital adjustment. In the aggregate, quadratic costs therefore provide a reasonable approximation to the dynamics of the more complicated underlying heterogeneous firm fixed adjustment cost model ("approximate representation"). Versions of the representative agent model estimated on aggregate data only will thus never be rejected even if our

lumpy investment model was the data generating process. More importantly, to the researcher who is interested in matching aggregate dynamics only, quadratic costs provide an easily manageable tool and an additional degree of freedom. He may scale investment responses to productivity innovations without missing aggregate dynamics implied by more detailed microeconomic foundations.

1.7.1 Do Estimates of Quadratic Costs have Structural Meaning?

We have shown so far that in an open economy setting, quadratic costs to capital adjustment provide a good aggregate representation of a model which realistically models the firm's investment problem. This does not yet prove, however, that the estimated quadratic adjustment costs have a structural interpretation. We know how model predictions may be invalidated when a macro outcome is ad hoc described as the result of a microeconomic decision problem one knows to differ from the actual one. There is widespread evidence for the main adjustment costs at the plant or firm level being non-convex and not quadratic (see Cooper and Haltiwanger (2006)). Consequently, estimated versions of quadratic cost models can be subject to the Lucas critique in the sense that their cost parameter lacks fundamentalness with respect to changes in policy and other deep parameters. Meanwhile, we also know from the example of Hansen (1985) that it is sometimes innocuous to represent the behavior of many agents by a decision problem that no agent actually faces if this is justified by aggregation itself.

The crucial question is, whether the quantitative link between estimates of convex aggregate and non-convex idiosyncratic adjustment costs is invariant to policy interventions. Calibrating the fixed costs from micro data – as we did – or alternatively estimating quadratic adjustment costs from aggregate time-series would thus be equivalent as it always yielded the same aggregate behavior. If, however, the macro-equivalent quadratic adjustment costs are not invariant to changes in non-adjustment-cost parameters, then using an estimated quadratic adjustment cost model is subject to the Lucas critique.[13]

To investigate this issue, we ask whether the matched quadratic adjustment costs

---

[13]Obviously, a prerequisite for this issue is whether approximate representation holds for a wide range of model parameters or is particular to our calibration. Throughout this section we find approximate representation to hold.

change if we vary key non-adjustment cost parameters of the model (that have a strong influence on the dynamics of aggregate investment in the model). We do so in four experiments. Each of these experiments has a policy interpretation, some maybe not perfect, but this imperfect interpretation comes at the advantage of a minimal intervention: not changing the model structure but only the model parameters.

First we look at variations in openness, ω. These variations can be thought of as introduction (removal) of tariffs or other trade barriers. As we have seen before, variations in openness have a strong impact on both the aggregate investment volatility and the ability of fixed adjustment costs to dampen this volatility. Second, we look at the introduction of a stochastic investment tax credit, modeled as a subsidy that changes the relative price of investment goods (which was so far fixed to one). Third, we investigate changes in the mark-up, η, which can be thought of as the result of competition policy. Fourth, we vary idiosyncratic risk, σ_ϵ. While this last experiment does not have a very clear policy analogue, it relates closely to the recent literature on time-varying uncertainty (see e.g. Bloom (2009) or Bachmann and Bayer (2011b)). It also helps to better understand the mixed results that we obtain with respect to the "fundamentalness" of quadratic adjustment costs.

1.7.2 Methodology: Matching Quadratic to Fixed Costs of Capital Adjustment

For the various parameter combinations we experiment with, we search for a size ϕ of quadratic adjustment costs such that the law of motion for aggregate capital coincides between fixed and quadratic adjustment cost specification. To be precise, we minimize the distance between the resulting parameters of the log-linear law of motion for relative capital stocks $\left(\alpha_1^\Delta \ \alpha_2^\Delta\right)$ in the two model specifications:

$$\min_\phi \Psi\left(\phi, \bar{\xi}\right) = \min_\phi \sum_{i=1,2} \left[\alpha_i^{\Delta, fixed}\left(\bar{\xi}\right) - \alpha_i^{\Delta, quad.}\left(\phi\right)\right]^2.$$

The justification for this matching strategy is that once accurate rules for predicting the two capital stocks have been found, at least in the quadratic adjustment cost case, the sequence of allocations can be pinned down from these rules and labor market clearing.[14] The obtained matches are very close in terms of Ψ, with square-root

---

[14]In order to obtain a solution to the firm problem in the quadratic cost setting, we again employ

differences ranging around 2%-6% of $\left(\alpha_1^\Delta + \alpha_2^\Delta\right)$. We report estimates of matched quadratic costs and minimized distances for all experiments in Section 1.B of the appendix.

1.7.3 Variations in Openness

Figure 1.2: Quadratic adjustment cost estimates for different values of ω
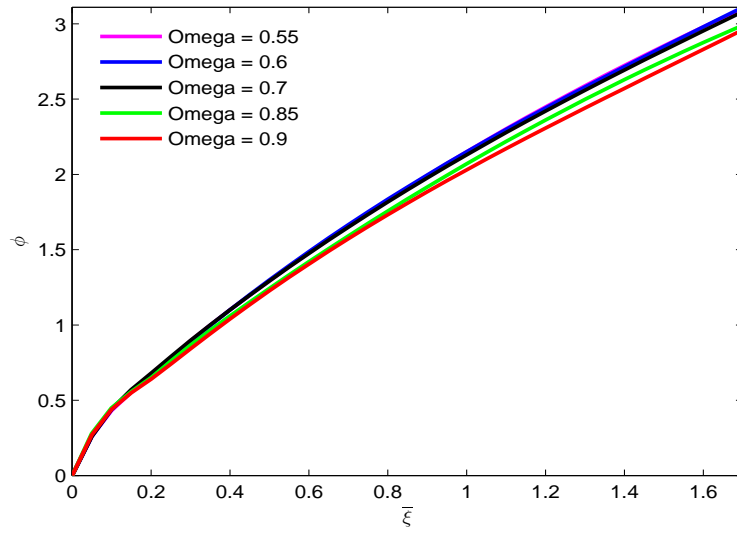
Notes: The figure displays estimates of quadratic capital adjustment costs for different values of the home bias parameter ω and for a broad range of fixed adjustment cost specifications $\overline{\xi}$. We obtained estimates of the quadratic adjustment cost parameter ϕ by minimizing the distance between the laws of motion of the relative capital stocks $\left(\alpha_1^\Delta \ \alpha_2^\Delta\right)$ for the fixed and quadratic cost model.

We start by varying the home-bias parameter from a value of 0.55, an economy very open to trade (say the Netherlands) to a value of 0.9, an economy almost closed to the world (say, the United States up to the mid 1990's). For each value for openness, we estimate matching quadratic adjustment costs across a broad range of fixed costs values. Figure 1.2 summarizes the results of this experiment. The figure displays the matched $\left(\phi, \overline{\xi}\right)$ pairs for five levels of openness, ω. In order to for the quadratic

---

a global solution technique using a variant of the solution method for the fixed cost case. In matching the two models, we take into account that firm productivities in the heterogeneous firm model are log-linearly distributed which implies an upwards adjustment of mean productivity by $\frac{\sigma_\epsilon^2}{2\sqrt{1-\rho_\epsilon^2}}$ and thus raises steady state capital stocks in both countries compared to a homogenous firm version of the model.

adjustment cost estimate to be ascribed structural or "fundamental" meaning, its value should be invariant across the different home bias specifications.

We find first that "approximate representation" holds for all values of adjustment costs and openness considered. Secondly, importantly, and perhaps more surprisingly, the lines in Figure 1.2, which represent the implied mapping $\phi(\bar{\xi})$ for various levels of openness, basically fall on top of each other.[15] Moreover, the obtained quadratic adjustment costs estimate is also broadly in line with the numbers typically considered when calibrating the quadratic adjustment cost model to aggregate data (With 3.08 being the estimate for our baseline specification). Accordingly, the link between fixed and quadratic adjustment costs passes this first test.

1.7.4 An Investment Tax Credit

Another more subtle trade policy is an investment tax credit (ITC). We assume that in country **H** there is an ad valorem subsidy τ_{inv} on buying investment goods, so that the effective price for the investment good decreases and (1.11d) in the firm's problem becomes:

$$V_H^{adj} = -(1 - \tau_{inv})i\varrho_H(\tau, C_H) + \beta\mathbb{E}\big[\bar{V}_H(\epsilon', k'; \hat{z}', m')\big]$$

However, to render the problem more interesting, we assume that the policy is stochastic. We model the tax policy state as a Markov chain with state vector $\begin{bmatrix} 0 & \tau_{inv} \end{bmatrix}$ and transition matrix Π_{tax}:

$$\Pi_{tax} = \begin{bmatrix} 0.97 & 0.03 \\ 0.25 & 0.75 \end{bmatrix}$$

The country applies the subsidy 10 percent of the time and once the policy is in place its expected duration is 4 quarters.

Under the policy, firms in country **H** pay less for investment goods. The subsidy is financed by a lump-sum tax levied on all households (through risk-sharing also effectively on the country **F** households). Technically this means that we can simply distort the firm's decision problem and leave all other equilibrium constraints, in par-

---

[15]The variation lies in the order of magnitude of the precision of our numerical procedure.

ticular the goods market clearing condition, unchanged. Of course this introduces an additional aggregate state to the problem and makes the firm problems asymmetric across countries. Consequently, the KS-rules need to be adjusted and as the policy only applies to one country need to accommodate the asymmetric setting now where we can no longer expect small variations in the world stock of capital.

Figure 1.3: Levels of ϕ corresponding $\bar{\xi} = 1.7$ for various sizes of ITC, τ_{inv}
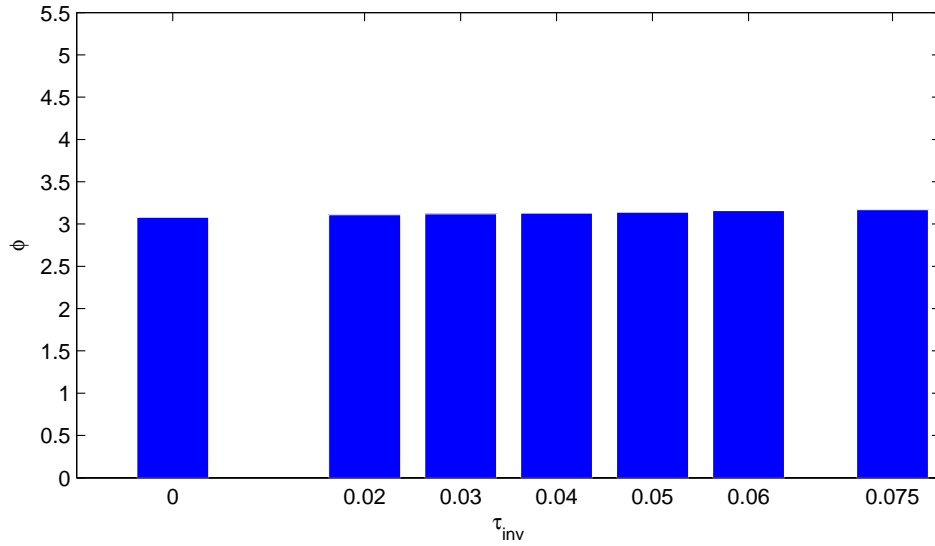
Notes: The figure displays estimates of matched quadratic adjustment costs after introducing a stochastic investment tax credit (ITC) in country **H** into the model environment. One bar corresponds to one value of the investment tax credit τ_{inv}. Estimates of the quadratic adjustment cost parameter ϕ were obtained by minimizing the distance between the laws of motion of the relative capital stocks $\left(\alpha_1^\Delta \ \alpha_2^\Delta\right)$ for the fixed and quadratic cost model.

Figure 1.3 displays the results for the quadratic-cost match for this experiment. The unconditional business cycle statistics are available in the appendix. While the introduction of an investment tax credit has substantial impact on the business cycle statistics (output and more so investment volatility goes up, consumption becomes less pro-cyclical), the impact on the size of the matched quadratic adjustment costs is minor. By and large also for this experiment we can view the representative-firm, quadratic-adjustment-cost model as a quasi "fundamental" stand in for the heterogenous-firm, fixed-adjustment-cost model.

Figure 1.4: Levels of ϕ corresponding $\bar{\xi} = 1.7$ for various η and σ_ε

| (a) Varying η | (b) Varying σ_ϵ |
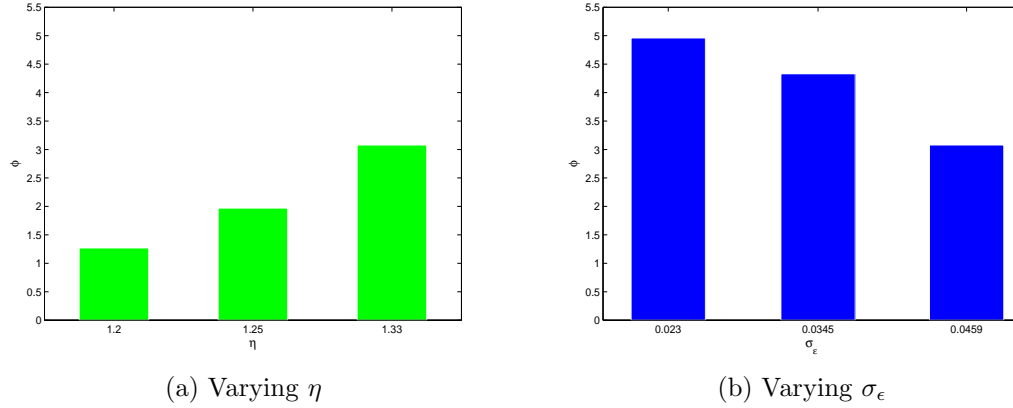
Notes: The figure displays estimates of matched quadratic adjustment costs for variations in the curvature of the production function (η) and the standard deviation of idiosyncratic productivity (σ_ϵ). For all experiments, we keep $\bar{\xi}$ fixed at its baseline value of 1.7. Estimates of the quadratic adjustment cost parameter ϕ were obtained by minimizing the distance between the laws of motion of the relative capital stocks $\left(\alpha_1^\Delta \; \alpha_2^\Delta\right)$. for the fixed and quadratic cost model.

1.7.5 Variations in Mark-up and Idiosyncratic Uncertainty

We so far found a stable quantitative link between estimates of quadratic and fixed capital adjustment costs across variations of what could broadly be described as "demand side" parameters.

Next we show that the quadratic adjustment costs are not fundamental with respect to all parameter variations. The quantitative link breaks down, once we start to consider variations on the "production side". Besides the size of adjustment costs $\bar{\xi}$ also the curvature of the production function (which in our case is pinned down by η) determines the aggregate importance of fixed adjustment costs. This is the key insight from Gourio and Kashyap (2007) who show that the capital-elasticity of revenues, i.e. the curvature of the reduced form revenue function, is central in determining the effect of non-convex adjustment costs.

The intuition for their finding can be grasped from thinking about the problem from a social planner's perspective. The planner chooses sequences of capital distributions (and labor) in order to maximize utility from consumption and leisure. Fixed adjustment costs imply a trade-off for the social planner between paying high adjustment costs to have firms adjusting frequently or accepting inefficiency from having equally productive firms employing different levels of capital. The social costs of

capital adjustments now result as the minimum loss the social planner can achieve by trading off higher adjustment frequencies with inefficient distributions of capital over production units.

This trade-off is absent in the quadratic cost model. Hence, anything that affects its severity has to be captured by the size of ϕ. The efficiency costs of unequal capital stocks depend on the capital-elasticity of revenues and are largest around an elasticity of 1/2. Consequently, the social planner accepts larger capital dispersions for lower η and lets firms adjust less frequently. For the same reason, the social planner uses more the intensive margin of adjusters to accommodate aggregate shocks. By a similar line of argument, the same should hold true for variations in productivity uncertainty σ_ε. Higher σ_ε mean more diverse capital stocks for equally productive firms if the frequency of adjustment is kept constant.

In fact, our "quadratic-fixed costs equivalence scale" breaks down once we vary the mark-up η or uncertainty σ_ε, see Figure 1.4. We consider $\eta \in \{4/3, 5/4, 6/5\}$ and $\sigma_\varepsilon \in \{0.23, 0.345, 0.46\}$. The figure displays the various equivalent levels of ϕ corresponding to fixed adjustment costs $\bar{\xi} = 1.7$. As one can see, the equivalent ϕ depends on the mark-up and is – as expected – decreasing in the revenue elasticity of capital (i.e. increasing in η) and decreasing in idiosyncratic uncertainty σ_ε. Hence the equivalent ϕ increases when uncertainty decreases. A similar result is found by Berger and Vavra (2010) for price adjustment.[16]

These results show that it is problematic to interpret the parameters of the representative firm model in a deep sense. Suppose for example a researcher estimates the aggregate representative firm model before and after a liberalization of intermediate goods markets, such that market power of single firms decreases and effectively η decreases. Then this researcher would conclude that the liberalization was accompanied by a decrease in capital adjustment costs. Similarly, a researcher that estimates the representative firm model from aggregate data once in turbulent and another time in more tranquil times, when σ_ε is small, would be led to the (false) conclusion that capital adjustment costs have increased.

---

[16]We also experimented with recalibrating $\bar{\xi}$ as the distribution of micro-investment rates changes with η and σ_ε. Qualitatively we obtained the same results.

1.7.6 Quantitative importance of deviations from fundamentalness

Yet, beyond the changes in estimated parameters, how severe is the non-fundamentalness of quadratic adjustment costs in terms of the ability of the model to produce forecasts of aggregate business cycle statistics? As discussed, one can also consider the mark-up experiment as an example of a policy, as changes in η can result from competition policy. Suppose, now, a researcher would study the effect of such policy on the business cycle behavior of the economy, but takes the stand-in representative firm model calibrated to aggregate data for the status quo ante of the reform. Assume the data is generated from our model with fixed adjustment costs and $\eta = 4/3$. Then assume the researcher uses the quadratic cost model to predict the macroeconomic implications of the policy change. According to our previous findings, he should have decreased the stand in quadratic adjustment costs, ϕ, as a result of the change in η. If the researcher fails to do so, this representative firm model under-predicts investment fluctuations.

Figure 1.5 shows that the size of the error can be sizable and close to 18% for the examples studied. The figure displays the relative difference of the investment volatility of the fixed adjustment cost model to the quadratic adjustment cost model, keeping ϕ at its value obtained as the stand-in quadratic adjustment costs for our baseline specification, $\phi = 3.08$. Changes in uncertainty have a very similar effect, again see Figure 1.5.

1.8 Conclusions

Chang et al. (2010) have pointed out that in many macroeconomic models where aggregate dynamics are represented as the decisions of a representative agent, insufficient care is applied when explicitly aggregating across potentially heterogeneous microeconomic units. We take this criticism seriously and apply it to the context of two country open economy models. In that literature, the problem of excess volatility in national investment dynamics is widespread and has usually been addressed by the introduction of quadratic costs to capital adjustment at the level of a representative firm.

However, this specification lacks empirical microeconomic foundation. Instead, much of the literature on plant level investment dynamics in the last two decades has

Figure 1.5: Relative deviations in investment volatility when keeping ϕ constant

(a) Varying η (b) Varying σ_ϵ
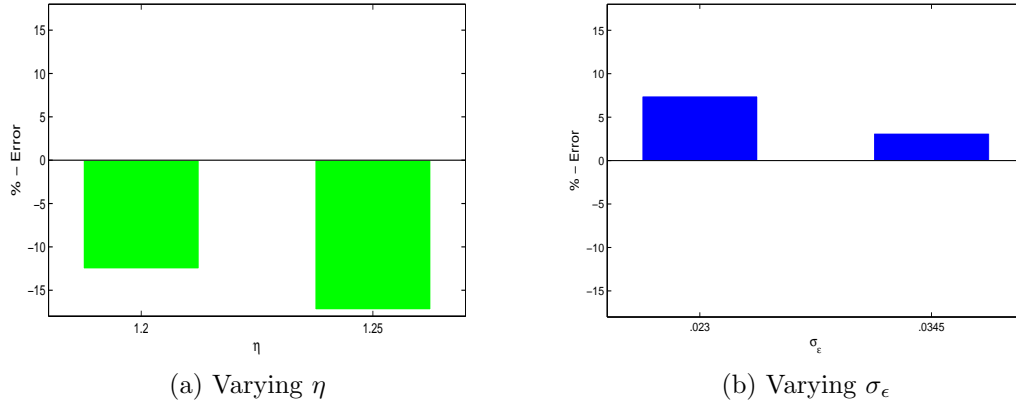
Notes: The figure shows prediction errors in investment volatility from the quadratic adjustment cost model relative to the underlying fixed adjustment cost model when varying "supply side" parameters. In panel (a) the mark-up parameter η changes and panel (b) varies idiosyncratic uncertainty. Quadratic adjustment costs are always kept at their baseline estimate of 3.08.

focused on so called lumpy investment models where micro units are faced with fixed costs to capital adjustment. These models have proven able to reproduce salient features of plant level investment behavior: long time-spans of virtual inactivity interspersed by occasional outbreaks of large and concentrated adjustments of the capital stock. How this behavior affects aggregation is a priori not clear.

In this chapter we have solved a relatively standard two-country real business cycle model of differentiated goods where firms face idiosyncratic productivity risk and stochastic fixed costs to capital adjustment. We demonstrate that this cost specification matters for shaping aggregate dynamics. The smoothing effect of fixed adjustment costs is the stronger, the more open an economy is to trade and the irrelevance result of Khan and Thomas (2003, 2008) obtains only as a limiting case of no openness to trade. Secondly, our model serves to rationalize the assumption of convex costs to capital adjustment in a representative firm setting as the aggregate effects of non-convex adjustment costs turn out to be indistinguishable from a quadratic cost setting. We argue that for aggregate purposes one may view the representative agent quadratic adjustment costs model as a suitable approximation. This approximation is not only a statistically sensible representation of the fixed adjustment cost model but can even be viewed as an economically sensible approximation in as far as the "deep" adjustment cost parameter of the quadratic adjustment cost model

does only co-depend on the parameters of the deeper fixed-costs model that regard the production side (revenue elasticity of capital and productivity heterogeneity). With this caveat in mind, we view estimates of the quadratic cost specification as a reasonable macro representation of the underlying investment technology in a wide range of applications.

Appendix to Chapter 1

1.A Calibrating Convex Adjustment Costs

The investment adjustment cost parameter $\overline{\xi}$ of our baseline specification was chosen to match a cross-sectional investment rate skewness of 2.2 which Bachmann and Bayer (2011a,b) find in firm level data. Table 1.5 displays a number of statistics for our baseline calibration which summarize annual plant level investment behavior as a function of investment adjustment costs. We always calculate gross annual investment for a plant, simulating a cross-section of 6000 plants. As spike adjuster, we qualify a plant whose annual gross investment rate exceeds 20 percent ($\frac{i_t}{.5(k_{t+1}-k_t)} >$ .2). Mean average spike adjustment refers to mean capital stock growth conditional on the adjustment being counted as spike. Finally, we calculate the average cost that firms have to bear from realized adjustments as a fraction of their annual output and wage bill.

Table 1.5: Annual Plant Level Statistics

η	1.33								1.25	1.2
$\overline{\xi}$	8E-3	2E-2	0.05	0.2	0.25	0.4	1.0	1.7	1.7	1.7
Mean fraction of adjusters	0.95	0.87	0.74	0.65	0.64	0.62	0.58	0.56	0.57	0.57
Std. fraction of adjusters	0.01	0.01	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02
Investment rate skewness	0.05	0.16	0.48	0.94	1.04	1.27	1.84	2.20	1.91	1.68
Mean fraction of spike adjusters	0.38	0.36	0.29	0.22	0.21	0.19	0.15	0.13	0.12	0.19
Std. fraction of spike adjusters	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Mean average spike adjustments	0.49	0.50	0.56	0.64	0.65	0.69	0.77	0.82	0.93	1.03
Std. of average spike adjustments	0.21	0.22	0.24	0.26	0.27	0.28	0.31	0.33	0.36	0.39
Adj. cost paid / % annual output	0.11	0.27	0.85	1.98	2.27	3.03	5.34	7.46	9.70	12.58
Adj. cost paid / % annual wages	0.23	0.55	1.70	3.97	4.55	6.06	10.67	14.92	18.18	22.64

Notes: Results are means from 100 simulations of 6000 plants over 200 periods (quarters). All statistics refer to annualized plant level investment rates. An adjusting plant is a plant that at least once within a given year chooses not to let its capital stock depreciate but to invest. A plant is counted as spike adjuster if its annual investment rate $\frac{i_t}{.5(k_{t+1}-k_t)}$ exceeds .2.

1.B Matching Convex to Non-convex Adjustment Costs

We minimize the Euclidian distance between the log-linear law of motion of relative capital stocks in a non-convex and a convex adjustment cost setting as described in the main chapter. Table 1.6 provides the resulting parameter matches and the minimized distance measure.

Table 1.6: Parameter Matches and Minimal Distances Ψ

	Variations in Openness									
	$\omega = 0.55$		$\omega = 0.6$		$\omega = 0.7$		$\omega = 0.85$		$\omega = 0.9$	
	ϕ	Ψ	ϕ	Ψ	ϕ	Ψ	ϕ	Ψ	ϕ	Ψ
$\bar{\xi} = 0.002$	0.0300	0.0053	0.0200	0.0051	0.0100	0.0027	0.0000	0.0009	0.0000	0.0015
$\bar{\xi} = 0.008$	0.0700	0.0071	0.0500	0.0032	0.0300	0.0007	0.0050	0.0015	0.0300	0.0012
$\bar{\xi} = 0.05$	0.2600	0.0025	0.2600	0.0010	0.2600	0.0003	0.2800	0.0012	0.2700	0.0010
$\bar{\xi} = 0.2$	0.6600	0.0022	0.6700	0.0021	0.6800	0.0014	0.6500	0.0008	0.6400	0.0006
$\bar{\xi} = 0.25$	0.7800	0.0014	0.7800	0.0018	0.7900	0.0012	0.7600	0.0006	0.7400	0.0005
$\bar{\xi} = 0.4$	1.1000	0.0003	1.1000	0.0008	1.1000	0.0007	1.0600	0.0004	1.0400	0.0003
$\bar{\xi} = 1.0$	2.1500	0.0005	2.1500	0.0002	2.1300	0.0001	2.0700	0.0002	2.0300	0.0002
$\bar{\xi} = 1.7$	3.1000	0.0008	3.1100	0.0006	3.0800	0.0003	2.9900	0.0001	2.9600	0.0001

Investment Tax Credit

	ϕ	Ψ		
$\tau_{inv} = 0.02$	3.11	0.0035		
$\tau_{inv} = 0.03$	3.12	0.0051		
$\tau_{inv} = 0.04$	3.13	0.0067		
$\tau_{inv} = 0.05$	3.14	0.0083		
$\tau_{inv} = 0.06$	3.16	0.0098		
$\tau_{inv} = 0.075$	3.17	0.0121		

Markup-variations

	ϕ	Ψ		
$\eta = 1.25$	1.97	0.0005		
$\eta = 1.2$	1.27	0.0013		

Variations in plant-uncertainty

	ϕ	Ψ		
$\sigma_\epsilon = 0.0345$	4.33	0.0013		
$\sigma_\epsilon = 0.023$	4.96	0.0031		

The table reports quadratic adjustment cost estimates ϕ and corresponding minimized distances ψ for all model specifications presented in the main chapter. In the experiments varying the investment tax credit, mark-ups and plant-level idiosyncratic uncertainty, openness (ω) is fixed at its baseline value of 0.7

1.C Numerical Solution Method

In order to numerically solve our model, we need to accurately approximate the laws of motion of the two cross-sectional distributions over capital and idiosyncratic productivity states in both countries. Given that our value function is of dimension $3 + 2n$ where $n \in \mathbb{R}^n$ is our number of moments by which to approximate μ^H and μ^F, we are limited in the number of moments for approximation. Another challenge is the simultaneous solution for the two prices τ and λ during the simulation step.

1.C.1 Accuracy of the Approximating Laws of Motion

Furtunately, we find that, as in related papers, it is sufficient to approximate the cross-sectional distributions by their first moments over capital in order to forecast prices judging by all standard measures for accuracy. We use the following general log-linear forms for the aggregate laws of motion:

$$\left[\log(\overline{k}_H) + \log(\overline{k}_F)\right]' = \alpha_0^{world} + \alpha_1^{world}\left[\log(\overline{k}_H) + \log(\overline{k}_F)\right] + \alpha_2^{world}\log(\hat{z}) + \alpha_3^{world}\tau_{inv} \tag{1.16a}$$

$$\left[\log(\overline{k}_H) - \log(\overline{k}_F)\right]' = \alpha_0^{\Delta} + \alpha_1^{\Delta}\left[\log(\overline{k}_H) - \log(\overline{k}_F)\right] + \alpha_2^{\Delta}\log(\hat{z}) + \alpha_3^{\Delta}\tau_{inv} \tag{1.16b}$$

$$\log(\varrho_H) = \alpha_0^{\varrho} + \alpha_1^{\varrho}\log(\overline{k}_H) + \alpha_2^{\varrho}\log(\overline{k}_F) + \alpha_3^{\varrho}\log(\hat{z}) + +\alpha_4^{\varrho}\tau_{inv} \tag{1.16c}$$

$$\log(\tau) = \alpha_0^{\tau} + \alpha_1^{\tau}\left[\log(\overline{k}_H) - \log(\overline{k}_F)\right] + \alpha_2^{\tau}\log(\hat{z}) + +\alpha_3^{\tau}\tau_{inv} \tag{1.16d}$$

For all model versions (including our baseline) where there is no investment tax credit α_3^{world}, α_3^{Δ}, α_3^{ϱ} and α_3^{τ} are obviously restricted to zero. One can view the laws of motion for world capital and capital differences as rotated versions of laws of motion for capital in each national economy. Also, for those variants theory predicts that such forecasting rules for each national economy should be symmetric, as the economies we model are symmetric. For our rotated rules, this symmetry with respect to the individual stocks of capital implies the following additional restrictions:

1. $\alpha_2^{world} = 0$

2. $\alpha_0^{\Delta} = 0$

3. $\alpha_0^{\tau} = 0$

Tables 1.7 to 1.11 report the parameter estimates which solve the Krusell-Smith algorithm for the model parameterizations that we consider in our main chapter. As first measure of accuracy we report the R^2 (Table 1.12- 1.13) along with the regression results. Note that forecast error variance is generally very low and that in our baseline calibration the minimum R^2 for all regression results is 99,99%.

As argued by Den Haan (2010) very low one-step-ahead prediction errors by themselves are no guarantee for having found an accurate description of the aggregate behavior of the economy. We therefore compute an additional measure which allows us to asses accuracy over longer forecast horizons. Specifically, we simulate the economy over a fresh, long productivity draw of 5000 periods, calculating market clearing prices and the behavior of the aggregate capital stock stocks exactly. We record these capital stocks as reference and then simulate the time path of the same variables using our estimated laws of motion as the data generating process. Our measure of accuracy is 99th percentile of percentage deviations between the two series.[17] Results are reported in Table 1.14. The fact that deviations between the two series are almost always below 1 % and for our baseline calibration below 0.5% gives us confidence in having found accurate solutions. Also note that there is no systematic relation between either forecast length and percentage error or cyclical position and percentage error, i.e. market clearing and forecasting rules do not diverge.

1.C.2 Simulating the economy

For simulating the economy and obtaining updates for the estimates of our aggregate rules we use the method proposed in Young (2010). This means approximating the distribution of firms over capital and productivity by a histogram over a fine but fixed grid and assigning all probability mass in transitions to the grid-points by linearly splitting it up between adjacent ones. Every period, the two equilibrium prices are computed exactly using the trust-region-dogleg algorithm of MATLAB's fsolve as described in Coleman and Li (1996) with a termination tolerance on prices

---

[17]This is a variant of the accuracy measure proposed in Den Haan (2010) who takes the maximum percentage deviation as accuracy measure. We therefore require the two series to be very close 99 % of the time. Our numerical solution for market clearing extremely rarely creates very short-lived spike-deviations. We view these blips rather as numerical outliers of the root finding algorithm. In fact the maximum deviations change depending on solving on a 64-bit or 32-bit machine. Therefore, taking the maximum error is likely to understate the quality of our results.

of 1E-7. During simulations, we use on-grid maximization methods to solve the firm problems when finding the root of our market clearing conditions. In order to do that, we interpolate value functions to our histogram capital of grid points using the current forecasts from the aggregate laws of motion.

Table 1.7: Krusell-Smith-Rules for $\left[\log(\overline{k}_H) + \log(\overline{k}_F)\right]'$

Variations in Opennes	$\omega = 0.55$		$\omega = 0.6$		$\omega = 0.7$		$\omega = 0.85$		$\omega = 0.9$	
	α_1^{world}	α_2^{world}	α_1^{world}	α_2^{world}	α_1^{world}	α_2^{world}	α_1^{world}	α_2^{world}	α_1^{world}	α_2^{world}
Frictionless	0.3002	0.9229	0.3013	0.9226	0.2998	0.9230	0.2999	0.9229	0.2999	0.9230
$\overline{\xi} = 0.002$	0.3010	0.9219	0.3008	0.9220	0.2999	0.9222	0.2997	0.9223	0.2998	0.9223
$\overline{\xi} = 0.008$	0.2998	0.9223	0.2993	0.9224	0.2990	0.9225	0.3000	0.9223	0.2998	0.9223
$\overline{\xi} = 0.05$	0.2957	0.9234	0.2955	0.9234	0.2958	0.9233	0.2964	0.9232	0.2965	0.9232
$\overline{\xi} = 0.2$	0.2916	0.9243	0.2916	0.9243	0.2916	0.9243	0.2918	0.9242	0.2918	0.9242
$\overline{\xi} = 0.25$	0.2905	0.9245	0.2905	0.9245	0.2905	0.9245	0.2906	0.9245	0.2908	0.9245
$\overline{\xi} = 0.4$	0.2876	0.9252	0.2876	0.9252	0.2876	0.9252	0.2878	0.9251	0.2879	0.9251
$\overline{\xi} = 1.0$	0.2798	0.9270	0.2797	0.9270	0.2797	0.9270	0.2798	0.9270	0.2798	0.9270
$\overline{\xi} = 1.7$	0.2732	0.9285	0.2732	0.9285	0.2732	0.9285	0.2733	0.9285	0.2734	0.9285

Investment Tax Credit	α_0^{world}	α_1^{world}	α_2^{world}	α_3^{world}
$\tau_{inv} = 0.02$	0.2768	0.9275	0.0005	0.0040
$\tau_{inv} = 0.02$	0.2763	0.9276	0.0008	0.0060
$\tau_{inv} = 0.02$	0.2756	0.9277	0.0011	0.0080
$\tau_{inv} = 0.02$	0.2748	0.9280	0.0014	0.0101
$\tau_{inv} = 0.02$	0.2737	0.9282	0.0016	0.0122
$\tau_{inv} = 0.02$	0.2719	0.9287	0.0017	0.0154

Markup-variations	α_1^{world}	α_2^{world}
$\eta = 1.25$	0.2894	0.9284
$\eta = 1.2$	0.3024	0.9280

Variations in plant-uncertainty	α_1^{world}	α_2^{world}
$\sigma_\epsilon = 0.0345$	0.2611	0.9297
$\sigma_\epsilon = 0.023$	0.2549	0.9298

Notes: This table reports the Krusell-Smith equilibrium laws of motion for the sum of world log-capital stocks for all model specifications presented in the main chapter. In the experiments varying the investment tax credit, mark-ups and plant-level idiosyncratic uncertainty, openness (ω) is fixed at its baseline value of 0.7.

Table 1.8: Krusell-Smith-Rules for $\left[\log(\overline{k}_H) - \log(\overline{k}_F)\right]'$

Variations in Opennes	$\omega = 0.55$		$\omega = 0.6$		$\omega = 0.7$		$\omega = 0.85$		$\omega = 0.9$	
	α_1^Δ	α_2^Δ	α_1^Δ	α_2^Δ	α_1^Δ	α_2^Δ	α_1^Δ	α_2^Δ	α_1^Δ	α_2^Δ
Frictionless	0.4542	0.7118	0.6376	0.6306	0.7907	0.5265	0.8821	0.4236	0.8995	0.3984
$\overline{\xi} = 0.002$	0.4657	0.6840	0.6372	0.6197	0.7884	0.5238	0.8816	0.4244	0.8989	0.3997
$\overline{\xi} = 0.008$	0.4950	0.6604	0.6474	0.6154	0.7924	0.5229	0.8808	0.4222	0.8985	0.3967
$\overline{\xi} = 0.05$	0.5683	0.5455	0.6758	0.5514	0.7999	0.4982	0.8833	0.4134	0.9003	0.3903
$\overline{\xi} = 0.2$	0.6518	0.4189	0.7151	0.4656	0.8103	0.4607	0.8863	0.4018	0.9023	0.3819
$\overline{\xi} = 0.25$	0.6689	0.3958	0.7247	0.4477	0.8134	0.4520	0.8871	0.3986	0.9029	0.3795
$\overline{\xi} = 0.4$	0.7053	0.3464	0.7468	0.4062	0.8211	0.4298	0.8892	0.3899	0.9043	0.3731
$\overline{\xi} = 1.0$	0.7716	0.2559	0.7931	0.3181	0.8407	0.3720	0.8949	0.3646	0.9082	0.3538
$\overline{\xi} = 1.7$	0.8044	0.2122	0.8190	0.2702	0.8542	0.3330	0.8995	0.3443	0.9115	0.3378

Investment Tax Credit	α_0^Δ	α_1^Δ	α_2^Δ	α_3^Δ
$\tau_{inv} = 0.02$	-0.0009	0.8543	0.3314	0.0122
$\tau_{inv} = 0.02$	-0.0013	0.8548	0.3306	0.0184
$\tau_{inv} = 0.02$	-0.0018	0.8553	0.3297	0.0248
$\tau_{inv} = 0.02$	-0.0022	0.8557	0.3289	0.0313
$\tau_{inv} = 0.02$	-0.0027	0.8561	0.3283	0.0379
$\tau_{inv} = 0.02$	-0.0033	0.8564	0.3275	0.0481

Markup-variations	α_1^Δ	α_2^Δ
$\eta = 1.25$	0.8421	0.3791
$\eta = 1.2$	0.8322	0.4177

Variations in plant-uncertainty	α_1^Δ	α_2^Δ
$\sigma_\epsilon = 0.0345$	0.8659	0.3020
$\sigma_\epsilon = 0.023$	0.8713	0.2922

Notes: This table reports the Krusell-Smith equilibrium laws of motion for differences in log-capital stocks between **Home** and **Foreign** for all model specifications presented in the main chapter. In the experiments varying the investment tax credit, mark-ups and plant-level idiosyncratic uncertainty, openness (ω) is fixed at its baseline value of 0.7.

Table 1.9: Krusell-Smith-Rules for $\log(\varrho_H)$

Variations in Openness	$\omega = 0.55$				$\omega = 0.6$			
	α_0^ϱ	α_1^ϱ	α_2^ϱ	α_3^ϱ	α_0^ϱ	α_1^ϱ	α_2^ϱ	α_3^ϱ
Frictionless	1.0123	-0.2357	-0.2014	-0.0704	1.0144	-0.2576	-0.1805	-0.1271
$\bar{\xi} = 0.002$	1.0190	-0.2356	-0.2022	-0.0713	1.0191	-0.2577	-0.1801	-0.1272
$\bar{\xi} = 0.008$	1.0179	-0.2348	-0.2019	-0.0713	1.0175	-0.2561	-0.1804	-0.1283
$\bar{\xi} = 0.05$	1.0135	-0.2323	-0.2011	-0.0727	1.0137	-0.2532	-0.1803	-0.1317
$\bar{\xi} = 0.2$	1.0091	-0.2295	-0.2005	-0.0745	1.0093	-0.2494	-0.1808	-0.1365
$\bar{\xi} = 0.25$	1.0078	-0.2288	-0.2003	-0.0748	1.0079	-0.2484	-0.1808	-0.1375
$\bar{\xi} = 0.4$	1.0043	-0.2272	-0.1996	-0.0755	1.0045	-0.2461	-0.1808	-0.1398
$\bar{\xi} = 1.0$	0.9944	-0.2232	-0.1974	-0.0768	0.9945	-0.2404	-0.1802	-0.1446
$\bar{\xi} = 1.7$	0.9859	-0.2201	-0.1953	-0.0774	0.9860	-0.2364	-0.1790	-0.1474
	$\omega = 0.7$				$\omega = 0.85$			
	α_0^ϱ	α_1^ϱ	α_2^ϱ	α_3^ϱ	α_0^ϱ	α_1^ϱ	α_2^ϱ	α_3^ϱ
Frictionless	1.0114	-0.3011	-0.1355	-0.2164	1.0102	-0.3679	-0.0680	-0.3237
$\bar{\xi} = 0.002$	1.0181	-0.3016	-0.1357	-0.2167	1.0179	-0.3682	-0.0690	-0.3258
$\bar{\xi} = 0.008$	1.0174	-0.3005	-0.1360	-0.2172	1.0189	-0.3686	-0.0687	-0.3280
$\bar{\xi} = 0.05$	1.0143	-0.2976	-0.1362	-0.2228	1.0152	-0.3656	-0.0687	-0.3337
$\bar{\xi} = 0.2$	1.0095	-0.2934	-0.1368	-0.2316	1.0098	-0.3613	-0.0691	-0.3433
$\bar{\xi} = 0.25$	1.0082	-0.2923	-0.1371	-0.2336	1.0085	-0.3602	-0.0693	-0.3459
$\bar{\xi} = 0.4$	1.0046	-0.2894	-0.1376	-0.2387	1.0049	-0.3575	-0.0696	-0.3526
$\bar{\xi} = 1.0$	0.9947	-0.2819	-0.1388	-0.2519	0.9947	-0.3499	-0.0708	-0.3725
$\bar{\xi} = 1.7$	0.9863	-0.2763	-0.1393	-0.2608	0.9864	-0.3438	-0.0718	-0.3884
	$\omega = 0.9$							
	α_0^ϱ	α_1^ϱ	α_2^ϱ	α_3^ϱ				
Frictionless	1.0100	-0.3911	-0.0447	-0.3596				
$\bar{\xi} = 0.002$	1.0181	-0.3918	-0.0455	-0.3631				
$\bar{\xi} = 0.008$	1.0190	-0.3921	-0.0452	-0.3653				
$\bar{\xi} = 0.05$	1.0152	-0.3891	-0.0452	-0.3711				
$\bar{\xi} = 0.2$	1.0099	-0.3850	-0.0455	-0.3808				
$\bar{\xi} = 0.25$	1.0085	-0.3839	-0.0456	-0.3834				
$\bar{\xi} = 0.4$	1.0050	-0.3813	-0.0459	-0.3903				
$\bar{\xi} = 1.0$	0.9947	-0.3739	-0.0468	-0.4112				
$\bar{\xi} = 1.7$	0.9864	-0.3680	-0.0477	-0.4283				

Notes: This table reports the Krusell-Smith equilibrium forecasting rules for ϱ_H for the variations in openness and fixed adjustment costs.

Table 1.10: Krusell-Smith-Rules for $\log(\varrho_H)$ (cont.)

Investment Tax Credit	α_0^ϱ	α_1^ϱ	α_2^ϱ	α_3^ϱ	α_4^ϱ
$\tau_{inv} = 0.02$	0.9923	-0.2778	-0.1413	-0.2613	0.0101
$\tau_{inv} = 0.03$	0.9926	-0.2776	-0.1419	-0.2615	0.0154
$\tau_{inv} = 0.04$	0.0031	-0.2773	-0.1426	-0.2618	0.0208
$\tau_{inv} = 0.05$	0.9943	-0.2772	-0.1434	-0.2620	0.0265
$\tau_{inv} = 0.06$	0.9960	-0.2775	-0.1442	-0.2621	-0.0324
$\tau_{inv} = 0.075$	0.9996	-0.2783	-0.1455	-0.2621	-0.0416
Markup-variations	α_0^ϱ	α_1^ϱ	α_2^ϱ	α_3^ϱ	
$\eta = 1.25$	1.0381	-0.2910	-0.1415	-0.2540	
$\eta = 1.2$	1.0685	-0.3011	-0.1427	-0.2481	
Variations in plant-uncertainty	α_0^ϱ	α_1^ϱ	α_2^ϱ	α_3^ϱ	
$\sigma_\epsilon = 0.0345$	1.0119	-0.2727	-0.1409	-0.2675	
$\sigma_\epsilon = 0.023$	1.0367	-0.2727	-0.1433	-0.2692	

Notes: This table reports the Krusell-Smith equilibrium forecasting rules for ϱ_H for the experiments varying the investment tax credit, mark-ups and plant-level idiosyncratic uncertainty.

Table 1.11: Krusell-Smith-Rules for $\log(\tau)$

Variations in Opennes	$\omega = 0.55$		$\omega = 0.6$		$\omega = 0.7$		$\omega = 0.85$		$\omega = 0.9$	
	α_1^τ	α_2^τ	α_1^τ	α_2^τ	α_1^τ	α_2^τ	α_1^τ	α_2^τ	α_1^τ	α_2^τ
Frictionless	0.3433	1.4083	0.3855	1.2716	0.4140	1.0825	0.4284	0.9249	0.4330	0.8992
$\bar{\xi} = 0.002$	0.3352	1.4266	0.3877	1.2726	0.4150	1.0841	0.4274	0.9309	0.4328	0.9078
$\bar{\xi} = 0.008$	0.3291	1.4270	0.3786	1.2844	0.4114	1.0865	0.4285	0.9373	0.4336	0.9131
$\bar{\xi} = 0.05$	0.3104	1.4575	0.3640	1.3185	0.4033	1.1147	0.4241	0.9537	0.4299	0.9279
$\bar{\xi} = 0.2$	0.2879	1.4927	0.3425	1.3666	0.3914	1.1588	0.4174	0.9813	0.4244	0.9521
$\bar{\xi} = 0.25$	0.2834	1.4990	0.3374	1.3765	0.3879	1.1688	0.4157	0.9886	0.4229	0.9586
$\bar{\xi} = 0.4$	0.2738	1.5126	0.3255	1.3994	0.3793	1.1943	0.4112	1.0079	0.4193	0.9761
$\bar{\xi} = 1.0$	0.2565	1.5374	0.3007	1.4478	0.3574	1.2604	0.3987	1.0648	0.4089	1.0282
$\bar{\xi} = 1.7$	0.2481	1.5493	0.2870	1.4741	0.3424	1.3048	0.3887	1.1102	0.4004	1.0712

Investment Tax Credit	α_0^τ	α_1^τ	α_2^τ	α_3^τ
$\tau_{inv} = 0.02$	0.0010	0.3426	1.3060	-0.0135
$\tau_{inv} = 0.02$	0.0015	0.3425	1.3061	-0.0205
$\tau_{inv} = 0.02$	0.0020	0.3424	1.3064	-0.0276
$\tau_{inv} = 0.02$	0.0025	0.3423	1.3065	-0.0349
$\tau_{inv} = 0.02$	0.0030	0.3423	1.3066	-0.0423
$\tau_{inv} = 0.02$	0.0038	0.3423	1.3065	-0.0538

Markup-variations	α_1^τ	α_2^τ
$\eta = 1.25$	0.3734	1.2714
$\eta = 1.2$	0.3958	1.2422

Variations in plant-uncertainty	α_1^τ	α_2^τ
$\sigma_\epsilon = 0.0345$	0.3297	1.3371
$\sigma_\epsilon = 0.023$	0.3247	1.3437

Notes: This table reports the Krusell-Smith equilibrium forecasting rules for the terms of trade τ for all model specifications presented in the main chapter. In the experiments varying the investment tax credit, mark-ups and plant-level idiosyncratic uncertainty, openness (ω) is fixed at its baseline value of 0.7.

Table 1.12: R^2s

| | $\omega = 0.55$ | | | | $\omega = 0.6$ | | | |
	α_x^{world}	α_x^{Δ}	α_x^{ϱ}	α_x^{τ}	α_x^{world}	α_x^{Δ}	α_x^{ϱ}	α_x^{τ}
$\bar{\xi} = 0.002$	1.0000	0.9950	1.0000	0.9999	1.0000	0.9978	1.0000	0.9998
$\bar{\xi} = 0.008$	1.0000	0.9982	1.0000	0.9999	1.0000	0.9990	1.0000	0.9998
$\bar{\xi} = 0.05$	1.0000	0.9994	1.0000	1.0000	1.0000	0.9993	1.0000	0.9999
$\bar{\xi} = 0.2$	1.0000	0.9998	1.0000	1.0000	1.0000	0.9996	1.0000	0.9999
$\bar{\xi} = 0.25$	1.0000	0.9999	1.0000	1.0000	1.0000	0.9997	1.0000	0.9999
$\bar{\xi} = 0.4$	1.0000	0.9999	1.0000	1.0000	1.0000	0.9998	1.0000	1.0000
$\bar{\xi} = 1.0$	1.0000	1.0000	1.0000	1.0000	1.0000	0.9999	1.0000	1.0000
$\bar{\xi} = 1.7$	1.0000	1.0000	1.0000	1.0000	1.0000	0.9999	1.0000	1.0000

| | $\omega = 0.7$ | | | | $\omega = 0.85$ | | | |
	α_x^{world}	α_x^{Δ}	α_x^{ϱ}	α_x^{τ}	α_x^{world}	α_x^{Δ}	α_x^{ϱ}	α_x^{τ}
$\bar{\xi} = 0.002$	1.0000	0.9995	0.9999	0.9997	0.9999	0.9999	0.9999	0.9998
$\bar{\xi} = 0.008$	1.0000	0.9998	0.9999	0.9997	1.0000	0.9999	0.9998	0.9998
$\bar{\xi} = 0.05$	1.0000	0.9998	0.9999	0.9997	1.0000	1.0000	0.9998	0.9997
$\bar{\xi} = 0.2$	1.0000	0.9998	0.9999	0.9997	1.0000	1.0000	0.9998	0.9997
$\bar{\xi} = 0.25$	1.0000	0.9998	0.9999	0.9997	1.0000	1.0000	0.9998	0.9997
$\bar{\xi} = 0.4$	1.0000	0.9998	0.9999	0.9998	1.0000	1.0000	0.9998	0.9997
$\bar{\xi} = 1.0$	1.0000	0.9999	0.9999	0.9999	1.0000	1.0000	0.9998	0.9997
$\bar{\xi} = 1.7$	1.0000	0.9999	1.0000	0.9999	1.0000	1.0000	0.9998	0.9998

| | $\omega = 0.9$ | | | |
	α_x^{world}	α_x^{Δ}	α_x^{ϱ}	α_x^{τ}
$\bar{\xi} = 0.002$	1.0000	0.9999	0.9999	0.9998
$\bar{\xi} = 0.008$	1.0000	1.0000	0.9998	0.9998
$\bar{\xi} = 0.05$	1.0000	1.0000	0.9998	0.9997
$\bar{\xi} = 0.2$	1.0000	1.0000	0.9998	0.9997
$\bar{\xi} = 0.25$	1.0000	1.0000	0.9998	0.9997
$\bar{\xi} = 0.4$	1.0000	1.0000	0.9998	0.9997
$\bar{\xi} = 1.0$	1.0000	1.0000	0.9998	0.9997
$\bar{\xi} = 1.7$	1.0000	1.0000	0.9998	0.9997

Notes: This table reports the R^2 of the regressions estimating the Krusell-Smith laws of motion for capital and the price forecasting rules for the variations in openness and fixed adjustment costs.

Table 1.13: R^2s (cont.)

Investment Tax Credit	α_0^ϱ	α_1^ϱ	α_2^ϱ	α_3^ϱ	α_4^ϱ
$\tau_{inv} = 0.02$	1.0000	0.9999	0.9999	0.9999	
$\tau_{inv} = 0.02$	1.0000	0.9999	0.9999	0.9999	
$\tau_{inv} = 0.02$	1.0000	0.9999	0.9999	0.9999	
$\tau_{inv} = 0.02$	1.0000	0.9999	0.9999	0.9999	
$\tau_{inv} = 0.02$	0.9999	0.9999	0.9999	0.9999	
$\tau_{inv} = 0.02$	0.9999	0.9999	0.9999	0.9999	
Markup-variations	α_0^ϱ	α_1^ϱ	α_2^ϱ	α_3^ϱ	
$\eta = 1.25$	1.0000	0.9999	0.9999	0.9999	
$\eta = 1.2$	1.0000	0.9999	0.9999	0.9999	
Variations in plant-uncertainty	α_0^ϱ	α_1^ϱ	α_2^ϱ	α_3^ϱ	
$\sigma_\epsilon = 0.0345$	1.0000	0.9999	1.0000	0.9999	
$\sigma_\epsilon = 0.023$	1.0000	0.9999	0.9998	0.9999	

Notes: This table reports the R^2 of the regressions estimating the Krusell-Smith laws of motion for capital and the price forecasting rules for the experiments varying the investment tax credit, mark-ups and plant-level idiosyncratic uncertainty.

1.C.3 Business Cycle Statistics for ITC Model

Table 1.15 displays the business cycle statistics for the model with investment tax credits of various sizes.

Table 1.14: Long-run accuracy

Variations in Openness

	$\omega = 0.55$	$\omega = 0.6$	$\omega = 0.7$	$\omega = 0.85$	$\omega = 0.9$
$\bar{\xi} = 0.002$	0.50	0.54	0.53	0.42	0.41
$\bar{\xi} = 0.008$	0.39	0.44	0.39	0.33	0.32
$\bar{\xi} = 0.05$	0.30	0.41	0.40	0.36	0.33
$\bar{\xi} = 0.2$	0.27	0.39	0.44	0.37	0.36
$\bar{\xi} = 0.25$	0.25	0.37	0.43	0.38	0.36
$\bar{\xi} = 0.4$	0.21	0.32	0.40	0.37	0.36
$\bar{\xi} = 1.0$	0.20	0.26	0.35	0.35	0.35
$\bar{\xi} = 1.7$	0.20	0.22	0.31	0.33	0.34

Investment Tax Credit

$\tau_{inv} = 0.02$	0.3486
$\tau_{inv} = 0.03$	0.3796
$\tau_{inv} = 0.04$	0.4276
$\tau_{inv} = 0.05$	0.5136
$\tau_{inv} = 0.06$	0.5980
$\tau_{inv} = 0.075$	0.7473

Markup-variations

$\eta = 1.25$	0.3810
$\eta = 1.2$	0.4414

Variations in plant-uncertainty

$\sigma_\epsilon = 0.0345$	0.2596
$\sigma_\epsilon = 0.023$	0.2731

This table reports the 99th percentile of percentage deviations between market clearing capital stocks and the ones implied by the aggregate laws of motion over a simulation of 5000 periods using a draw for aggregate productivity different from the one that was used in the Krusell-Smith solution algorithm.

Table 1.15: Cyclical Properties with Investment Tax Credit (Matched ϕ)

τ_{inv}	0.02	0.03	0.04	0.05	0.06	0.075

Standard Deviations in %

	0.02	0.03	0.04	0.05	0.06	0.075
Output	1.842	1.934	2.056	2.207	2.374	2.268

Standard Deviations relative to Output

	0.02	0.03	0.04	0.05	0.06	0.075
Investment	6.00	7.326	8.586	9.840	11.269	14.297
Consumption	0.243	0.285	0.327	0.364	0.396	0.434
Employment	0.699	0.841	0.978	1.100	1.203	1.326
Exports	0.922	1.228	1.507	1.746	1.944	2.183
Imports	1.078	1.461	1.798	2.082	2.312	2.580
NX	0.708	0.928	1.122	1.284	1.412	1.555
ToTs	0.9666	0.963	0.958	0.953	0.948	0.941

Correlation with Output

	0.02	0.03	0.04	0.05	0.06	0.075
Investment	0.778	0.752	0.745	0.741	0.727	0.674
Consumption	0.580	0.318	0.102	-0.069	-0.204	-0.356
Employment	0.878	0.842	0.832	0.836	0.845	0.862
Exports	-0.018	-0.172	-0.292	-0.386	-0.460	-0.546
Imports	0.062	0.214	0.331	0.422	0.494	0.576
NX	-0.538	-0.551	-0.584	-0.621	-0.655	-0.699
ToTs	0.858	0.725	0.577	0.429	0.293	0.116

Persistence

	0.02	0.03	0.04	0.05	0.06	0.075
Output	0.693	0.687	0.679	0.671	0.664	0.653
Investment	0.577	0.559	0.545	0.529	0.507	0.455
Consumption	0.666	0.627	0.600	0.572	0.569	0.555
Employment	0.612	0.586	0.569	0.557	0.549	0.542
Exports	0.632	0.576	0.548	0.532	0.523	0.513
Imports	0.604	0.560	0.540	0.530	0.525	0.521
NX	0.562	0.538	0.527	0.520	0.517	0.515
ToTs	0.732	0.717	0.701	0.684	0.670	0.650

Chapter 2

Quantifying the Contribution of Search to Wage Inequality

2.1 Introduction

Mincerian wage regressions explain only about a third of the observed inequality in wage data. Search theoretic models of the labor market offer a compelling explanation for this phenomenon. Their central assumption is that sampling job offers in unemployment takes time and is subject to the opportunity cost of foregone wages. Identical workers, therefore, accept a range of heterogeneous job offers.[1] The literature has come to call this *frictional wage dispersion*. Understanding how much of residual inequality results from search frictions opposed to unobserved worker heterogeneity is of first order importance for judging the efficiency of labor markets and designing appropriate social insurance schemes.

Structural models that seek to answer this question conclude that more than 40 percent of wage inequality can be explained by the search friction (see Postel-Vinay and Robin (2002) and Carrillo-Tudela (2012)). Hornstein et al. (2012) (henceforth referred to by HKV) show that on the job search is the key mechanism that generates large frictional wage dispersion. A high offer arrival rate on the job implies that workers are giving up less when moving out of unemployment. This makes them willing to accept relatively poor job offers. Moreover, they quickly move up the job ladder which means a larger share of workers with relatively high wages.

In this chapter, we provide evidence from the Survey of Income and Program Participation (SIPP) that an important share of job to job transitions is not value improving. Accounting for this, we calibrate a structural search model with worker

---

[1] See Mortensen (2003) and the references therein.

and job heterogeneity that replicates observed overall and residual wage inequality. It attributes less than 14 percent of overall wage inequality to the search friction. This result comes in spite of our inclusion of a number of important channels that enlarge the set of acceptable job offers to the worker: skill accumulation on the job, skill loss in unemployment and search on the job. The crucial novelty is the introduction of reallocation shocks that we calibrate to the share of wage losses after a job to job transition. Without them, in a recalibrated model, the variance of the wage offer distribution more than doubles and the contribution of the search friction jumps to over 38 percent, in line with the findings in the previous literature.

The basic intuition for our main result can be summarized in three steps. First, the share of residual inequality that the search friction can explain is an increasing function of the variance of the job offer distribution. This is an unobservable object that has to be identified from wage data. Here, we follow Low et al. (2010) and estimate if from the excess variance of wage growth for job switchers relative to job stayers.

Second, when all job to job transitions are value improving, workers quickly move into the high ranked jobs from which they are unlikely to accept further offers. Calibrated search efficiency therefore has to be high in order to replicate the size of observed job to job flows. This, in turn, means that workers are concentrating in the high ranked jobs even faster. Most observed wage changes between jobs are then small improvements such that a high excess variance of wage growth for job switchers can only be rationalized by a very disperse job offer distribution.

Third, as we demonstrate using a variation of the on the job search model studied by HKV, this causal chain is broken by the introduction of what Jolivet et al. (2006) label a *reallocation shock*: A fraction of the on the job offers leaves the worker only to decide between accepting the outside offer or moving into unemployment. The lower outside option decreases the inferred on the job offer arrival rate because workers are more likely to accept an offer resulting from a reallocation shock. These setbacks also mean that workers move into the high ranked jobs at a lower rate. This makes them more likely to accept outside offers. All in all, negative wage growth observations and a larger set of acceptable voluntary outside offers mean that the same excess variance of wage growth can be rationalized with a much less disperse wage offer distribution.

Can we find evidence for reallocation shocks in the data? Fujita (2011) using data from the UK Labour Force Survey shows that an important share of workers who search on the job do so to avoid unemployment. We extend his analysis using the SIPP employment data to show that reallocation shocks are an important driving force behind observed flows. About a third of all job to job transitions yield lower nominal wages for the worker and neither observable non-wage benefits nor higher expected wage growth can account for workers accepting these lower wages. Instead, workers who initially accept a wage cut are more likely to switch jobs again shortly afterwards. Our quantitative model allows us to map the share of losses into the size of reallocation shocks explicitly controlling for measurement error and stochastic innovations to workers' wages. We estimate reallocation shocks to be responsible for 60 percent of observed losses.

The remainder of the chapter is structured as follows. Section 1 gives an overview of related literature. In Section 2 we lay out the simple analytical model that highlights the importance of reallocation shocks. Section 3 provides empirical evidence for their presence in the data and highlights stylized facts of residual wage dispersion. We present our full model in Section 4. Section 5 discusses its parameterization. Section 6 presents and analyzes the results, and Section 7 concludes. Additional information on the analytical derivations, the empirical part and the numerical algorithm is relegated to the appendix.[2]

2.2 Further Related Literature

Burdett et al. (2011) and Ortego-Marti (2012) show that workers' reservation wages fall significantly in a job ladder model augmented by skill accumulation on the job and skill depreciation in unemployment, respectively. These models match the mean to minimum residual wage in the data, potentially rationalizing all residual inequality as frictional.[3] We incorporate these features into our model to give it a fair chance of generating substantial frictional inequality. We show that the inferred job offer distribution provides an upper bound for the share of residual inequality that can be

---

[2]All programs used for data analysis and model solution are available on the authors' web pages.

[3]Other recent papers that study conditions under which frictional wage inequality can explain all residual inequality are Papp (2012) and Michelacci et al. (2012). An earlier example is Bontemps et al. (2000).

thought of as frictional.

Another strand of related literature tries to decompose residual inequality from reduced-form specifications. Abowd et al. (1999) and Hagedorn and Manovskii (2010) find that search frictions explain between $7 - 25$ percent of the French inter-industry differential and 6 percent of US wages, respectively. These models rely on exogenous labor mobility and either a permanent component of worker heterogeneity (Abowd et al. (1999)), or a stationary shock process (Hagedorn and Manovskii (2010)). Our structural model allows us to explicitly model the selection of workers into matches. Moreover, we confirm findings from previous studies that residual wage inequality increases strongly over a worker's life-cycle. This suggests a permanent shock component in individual wage potential. Our model allows for such a non-stationary shock process and our decomposition of workers' wages over the life-cycle shows that a substantial part of heterogeneity is the result of different employment histories during working life.[4] Finally, also using the SIPP, Low et al. (2010) use a selection model to infer the wage offer distribution and the shock process of individual wage potential from US wage data. While we ask a different question and use a different empirical strategy, our estimates yield a comparable magnitude for the relative size of idiosyncratic and employment risk.

2.3 Intuition from a Simple Model

HKV show that the job offer arrival rate on the job is a key parameter determining the wage distribution, and thus the amount of frictional wage inequality, in job ladder models. The higher the on the job offer arrival rate is compared to in unemployment, the smaller is the option value the worker gives up by remaining unemployed and waiting for better offers. Consequently, the minimum wage accepted by workers decreases. Additionally, a high offer arrival rate on the job implies that workers quickly move up the job ladder. This leads to relatively many workers located at high paying jobs. The fact that 1 in 40 employees in the US labor market switches jobs every month seems to hint at high offer arrival rates on the job.

---

[4]Hagedorn and Manovskii (2010) assume transitory shocks to the worker component and attribute 6% of US wage dispersion to search frictions. Using their identification strategy on our non-stationary shock process, search frictions explain almost none of the variance of log wages in our simulated data.

Using an extension to the model studied by HKV, we now demonstrate that one can match high job to job transitions with substantially lower job offer arrival rates when introducing what Jolivet et al. (2006) label a *reallocation shock*: A fraction of all on the job offers do not allow the worker to stay with his current job, but only leave him to choose between accepting other employment or becoming unemployed. One may think of these shocks as both transitions within layoff notice period as well as those originating out of non-pecuniary motives such as moving in with one's spouse or closer to one's parents.[5] We show that these shocks crucially affect the wage distribution, both directly and indirectly by the lower inferred on the job offer arrival rate.

Our exposition here is parsimonious and focuses on a few key equations. Appendix 2.A provides a full characterization of the solution. There is a unit mass of homogeneous workers receiving wage offers at Poisson rate λ_u when unemployed and with rate λ when employed. Wage offers are random draws from a cumulative wage offer distribution $F(w)$ with upper support w_{max} that the worker can accept or reject. Time is continuous and workers discount the future at rate r. It is easy to see that the worker follows a reservation wage strategy where the minimum accepted wage is denoted w^*. The asset value of being employed with current wage w is:

$$rW(w) = w + \lambda(1 - \lambda_d) \int_w^{w_{max}} [W(z) - W(w)]dF(z)$$
$$+ \lambda\lambda_d \int_{w^*}^{w_{max}} [W(z) - W(w)]dF(z)$$
$$- (\omega + \lambda\lambda_d F(w^*))(W(w) - U).$$

The worker receives a "normal" on the job offer with probability $\lambda(1 - \lambda_d)$, where λ_d is the probability that an on the job offer is a reallocation shock. The second line is the value of accepting an outside offer after a reallocation shock. Note that now workers accept all wage offers above the reservation wage because they do not have the option to stay with their old jobs. The third line states the value of moving into unemployment which either happens with probability ω after exogenous job destruction, or when the worker refuses an offer after a reallocation shock which

---

[5]This is in distinction from a transition where the benefit might have been non-monetary but related to the new job like a more permanent work contract or employer provided health insurance.

occurs with probability $\lambda\lambda_d F(w^*)$. When setting $\lambda_d = 0$, the model reduces to the job ladder model studied by HKV. The asset value of unemployment reads:

$$rU = b + \lambda_u \int_{w^*}^{w_{max}} [W(z) - U]dF(z).$$

An unemployed worker receives benefits b and samples job offers at rate λ_u.

We now establish that a larger share of reallocation shocks decreases the job offer arrival rate inferred from employment transition data and reduces the share of workers with relatively high wages. We then demonstrate that this lowers the amount of wage dispersion implied by the model. The on the job offer arrival rate is typically identified by matching a fixed job to job transition rate, which we label JTJ, and which is given by:

$$JTJ = \lambda(1-\lambda_d) \underbrace{\int_{w^*}^{w_{max}} [1-F(z)]dG(z)}_{=:ANO} + \lambda\lambda_d \underbrace{[1-F(w^*)]}_{=:ARO},$$

where $G(w)$ is the realized distribution of wages. We define ANO as the average probability that a normal on the job offer is accepted and ARO as the probability that an offer is accepted after a reallocation shock. Solving for the implied on the job offer rate gives:

$$\lambda^* = \frac{JTJ}{(1-\lambda_d)ANO + \lambda_d ARO}.$$

Increasing the share of reallocation shocks λ_d decreases the inferred on the job offer rate λ^* for two reasons. First, job offers after a reallocation shock are accepted with probability ARO which is larger than the average probability of a normal on the job offer being accepted (ANO). Second, it indirectly affects the latter by changing the wage distribution $G(w)$ which we derive in Appendix 2.A:

$$G(w) = \frac{F(w) - F(w^*)}{1 - F(w^*)} \frac{\overbrace{\omega + \lambda^*\lambda_d}^{=:D}}{\underbrace{\omega + \lambda^*\lambda_d}_{=:D} + \underbrace{\lambda^*(1-\lambda_d)}_{=:C}[1-F(w)]}. \tag{2.1}$$

Reallocation shocks have two effects on the wage distribution. First, like exogenous destruction, they move workers into unemployment from which they subsequently

accept any offer above their reservation wage (D). In addition, C shows that they decrease the amount of regular job offers, and thus, the speed that workers climb up the job ladder. Consequently, $G(w)$ becomes steeper at low values, i.e., more workers have relatively low wages implying that the probability of a normal offer being accepted (ANO) rises.

In Section 2.6.2, we infer the wage offer distribution $F(w)$ from wage data and show that the mechanisms just outlined have large quantitative implications for the inference. To fix ideas, we here study the effects of changes in λ_d on wage dispersion for a given $F(w)$. HKV propose the ratio of the mean to the minimum wage (Mm-ratio: $\bar{w}/w^*$) as summary statistic to compare wage dispersion across different classes of search models.[6] The measure has become a popular statistic in the literature, and for comparability we use it as one summary statistic for wage dispersion later in the chapter.

Table 2.1: Parameterization Simple Model

Parameter	Value
b	$0.4\bar{w}$
λ_u	0.3
$F(w)$	$\ln \mathcal{N}(0, 0.04)$
JTJ	$2.5\ \%$
r	$0.33\ \%$

Unemployment benefits b are a fraction of the mean wage $\bar{w}$. JTJ designates the job to job transition rate.

In Appendix 2.A, we show that the reservation wage is characterized by:

$$w^* = b + (\lambda_u - \lambda^*) \int_{w^*}^{w_{max}} \frac{1 - F(z)}{r + \omega + \lambda^*\lambda_d F(w^*) + \lambda^*\lambda_d F(z) + \lambda^*[1 - F(z)]} dz. \quad (2.2)$$

It is the sum of the flow benefits in unemployment and the option value to keep searching in unemployment. As in a pure job ladder model ($\lambda_d = 0$), the latter is decreasing in the difference $\lambda_u - \lambda$, because workers are giving up less in terms of search efficiency when moving out of unemployment. Similarly, r and ω decrease the value of additional search because workers become more impatient and high wage

---

[6]In the models they study, this measure is independent of the wage offer distribution $F(w)$. This does not hold in the environment studies here (see Appendix 2.A for a proof).

offers have a lower duration, respectively. Using comparative statics, we demonstrate that changes in λ_d affect the minimum wage directly and indirectly via the implied search efficiency on the job:

$$\frac{dw^*}{d\lambda_d} = \underbrace{\frac{\partial w^*}{\partial \lambda_d}}_{<0} + \underbrace{\frac{\partial w^*}{\partial \lambda^*}\frac{\partial \lambda^*}{\partial \lambda_d}}_{>0}.$$

The direct effect of a reallocation shock can be directly read from (2.2): With probability $F(w^*)$, like exogenous job destruction, it decreases the expected duration of holding employment. Moreover, the further a worker moves up the job ladder, the more likely he will move into a lower ranked job, which decreases the difference in valuation between higher and lower ranked jobs. Both factors decrease the incentive to wait for better offers when moving out of unemployment.[7] However, the increase in reallocation shocks decreases λ^* which increases the reservation wage. Theoretically, the effect λ_d has on the minimum wage is therefore ambiguous and may change depending on parameter values.

The mean wage, is given by:

$$\bar{w} = \int_{w^*}^{w_{max}} w dG(z)$$

Provided our earlier discussion, it should be intuitive that it is a decreasing function of λ_d. More reallocation shocks imply a steeper $G(w)$ and hence a lower mean wage.

For the remainder of this section, to be able to supply graphical representations to our argument, we impose parametric assumptions on the model. Table 2.1 lists the parameter values. All of them are relatively common in the literature (HKV use similar parameter values in their exposition).

Figure 2.2a demonstrates how the wage distribution becomes steeper as λ_d increases. Figure 2.2b shows the drop in the inferred on the job offer arrival rate. The model estimate reacts particularly sensitive to changes at small values of λ_d. Regarding the reservation wage, Appendix 2.A shows that it rises up to $\lambda_d = 0.35$ and starts to decrease again slowly afterwards. The resulting Mm-ratio from varying λ_d given our parameter values is reported in Figure 2.2c. Especially for low values

---

[7]It is this effect which has Hornstein et al. (2007) conclude that reallocation shocks should unambiguously increase the Mm-ratio.

Figure 2.1: Parameterized Simple Model

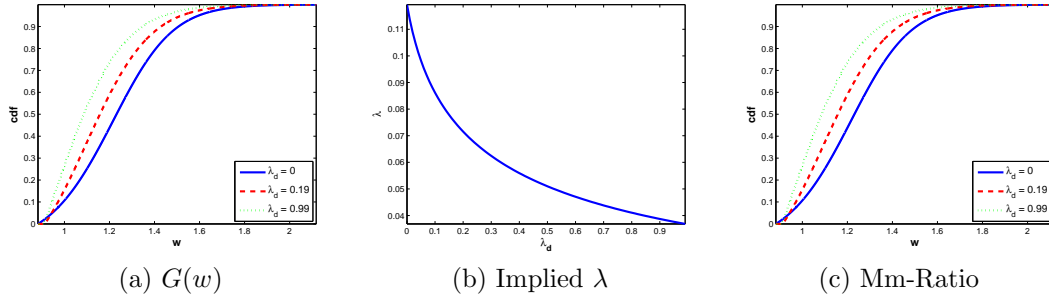(a) $G(w)$ (b) Implied λ (c) Mm-Ratio

Figure 2.2a shows the implied distributions of wages paid $G(w)$ for different reallocation shock probabilities λ_d using the parameterization reported in Table 2.1. Figure 2.2b reports the implied search efficiency λ for the same exercise, and Figure 2.2c reports the resulting Mm-ratio.

of λ_d, the Mm-ratio decreases quite sharply in the share of reallocation shocks.

2.4 Reallocation Shocks and Residual Wage Dispersion in the Data

In this section, we introduce our data set, the Survey of Income and Program Participation (SIPP), and discuss sample selection. We compile different pieces of evidence to show that reallocation shocks are an important feature of the data and link them to existing evidence in other studies. We also obtain the distribution of residual wages from a Mincerian wage regression. Residual inequality is large and shows a substantial increase with worker age.

2.4.1 Data Source and Sample Creation

Our analysis requires detailed longitudinal information on wages, worker and job characteristics at a very high temporal resolution. The data set most adequate for these requirements is the SIPP of which we employ the 1993 and 1996 panels.[8] It is a representative sample of the non-institutionalized civilian US population maintained

---

[8]Our data set is based on CEPR SIPP extracts available for download at *http://www.ceprdata.org/sipp/sipp_data.php*. We modify these abstracts to include further information contained in the SIPP files but not in the original abstracts. Appendix 2.B provides additional information on the differences between the two data sets and the steps we take to merge them.

by the US Census Bureau.[9] The level of detail it provides in individual records allows us to accurately identify an individual's main job and hourly wages on that job.[10] Our initial sample consists of 5,243,222 person/month observations.

Our data cover the years 1993-1995 (1993 sample) and 1996-1999 (1996 sample) providing us with up to 48 months of observations per individual. We use observations from individuals aged 23-55, for whom we require complete information on the individual's employment status, age and employer id. We only consider an individual's primary job[11] and drop workers that are recalled by former employers or have missing reporting months during a job spell.[12] Moreover, we drop workers reporting to be school enrolled, the self-employed, family-workers, members of the armed forces, workers at non-profit companies and anyone whose wage information was imputed by the SIPP.[13] Finally, we truncate the wage distribution at the top and bottom 1% to take care of outliers and top-coding.[14] These restrictions leave us with 2,039,345 person/month observations.

We identify job to job transitions as those transitions where the worker works in two consecutive months without reporting unemployment in between,[15] and either the worker's employer identification number or his two-digit occupational identifier changes.[16] Appendix 2.C provides a discussion for alternative measures of job to job transitions and compares our estimate to those obtained from CPS data.

---

[9]The 1996 panel oversamples poor households. We use population weights provided by the SIPP throughout our analysis.

[10]The survey reports at most two jobs for each 4-month recording period. In case an individual holds more than two jobs, the two jobs with most hours worked are reported.

[11]As primary job we consider the position where the largest share of hours worked is spent.

[12]In case of recall, we choose to exclude those observations because recalled workers likely possess a different search technology than what we include in our model specification.

[13]Since our investigation starts from the observation that wage predictions conditional on worker observables explain only a relatively small part of wages, it would seem odd to include wage observations which are mere predictions of these very models

[14]Earnings are topcoded at $33333 and $50000 for a four month period in the 1993 and 1996 sample, respectively.

[15]Theoretically, we could use the weekly employment status and count job to job transitions only, when a worker is employed in two consecutive weeks. However, it seems reasonable to assume that a few days in between jobs may be spent on a potential relocation or other pre-work sensitivities. Hence, we only discard observations where the worker reports to actively seek a job during non-employment.

[16]We think of job to job transitions as a change in the technology operated by the worker; therefore, we include both, changes in job ids (as in Fallick and Fleischman (2004)) and occupation (as in Moscarini and Thomsson (2007)).

2.4.2 Reallocation Shocks and On the Job Search

This section provides empirical evidence from previous studies and our own data that reallocation shocks are an important feature of employment transitions. While we cannot infer their size directly from the data, Section 2.5 uses a moment from the data together with an extended search model to quantify the share of these shocks.

The existing literature already highlights several shortcomings of a pure job ladder model. Fallick and Fleischman (2004) find for the CPS that a worker who reports to be actively searching on the job is more likely to be unemployed the next month. Fujita (2011) uses a question in the UK labor force survey that asks employees to state a reason for their engaging in on the job search. He finds that of those who report to be actively searching, 12 percent do so for fear of loosing their current job and another 27 percent because they are unsatisfied with their current job due to non-pecuniary reasons. Nágipal (2005) shows for a basic job ladder model that the job offer arrival rate on the job has to be higher than during unemployment in order to replicate observed flow rates. Jolivet et al. (2006) show that in the PSID 23.3 percent of job to job transitions are associated with nominal wage decreases. Including reallocation shocks into a Burdett and Mortensen (1998) model, they find that these shocks account for a third of all job to job offers. Using the SIPP, Connolly and Gottschalk (2008) find that 44.1 percent of all job to job transitions lead to lower real wages. They stress that a higher future expected wage growth may explain initial wage cuts and estimate that 64 percent of male and 81 percent of female wage cuts are truly transitions to lower valued jobs.[17]

Regarding our own data, the SIPP asks workers who terminate a job for their reason to do so. The answers further corroborate the evidence previously cited: Only 55 percent of those responding state that they *quit to take another job.* In contrast, 19 percent of jobs ended, because the previous job did not provide the possibility to continue.[18] Adding another 4 percent of cases which pertain to personal or family related issues, this yields up to 23 percent of transitions where, for one reason or another, staying with the old job may not have been an option. There are a number of caveats to the informativeness of this variable: Some of the possible answers are not

---

[17]Vice versa, they find that 1.3 percent of females' and 8.6 percent of males' transitions with wage improvements actually go into lower valued matches.

[18]This includes the answers *on layoff*, *job was temporary and ended*, *discharged/fired*, *employer bankrupt*, *employer sold business*, and *slack work or business conditions.*

Table 2.2: Wage Cuts after Job to Job Transitions

Sample Stratification	Share loss	Mean loss
Whole sample	0.344	-0.196
Job characteristics		
- Non-Union to Union	0.346	-0.196
- Health insurance	0.352	-0.196
- Education	0.352	-0.196
Old wage		
- Lowest 25%	0.232	-0.16
- 25-75%	0.352	-0.198
- Top 25%	0.457	-0.215

The Table shows the share of workers incurring a cut in nominal hourly wages after a job to job movement for our sample population as a whole as well as for several subsets. Mean loss reports the mean wage loss in log points conditional on suffering a wage cut upon movement. Under *Job characteristics*, the first line excludes workers from the sample who transit from non-unionized to unionized jobs, the second and third line additionally exclude workers who move from jobs without health insurance to an employer providing an insurance policy and movements where the new employer subsidizes expenses on education. The panel *Old wage* divides workers based on their wages on the old job.

mutually exclusive, or do not map directly into our interpretation of a reallocation shock. Even more problematic, in less than 30 percent of the cases we identify as job to job transitions, the worker provides an answer.[19] [20]

Instead of trying to infer search efficiency from this rather noisy variable, we follow a different strategy in combining employment flow data with accompanying wage dynamics. As we report in Table 2.2, a pervasive phenomenon in the data are job to job transitions resulting in nominal wage losses. In the whole population, roughly one third of all transitions result in workers earning lower hourly wages in the month after the transition compared to the last month on the previous job.[21] Conditional losses are substantial with workers on average receiving about 20 percent lower wages than previously.[22]

---

[19]For a negligible share the question is not applicable, because only the main job changed, but the worker stays with his old employer. See Appendix 2.C for a detailed discussion on how we identify job to job transitions.

[20]Nágipal (2008) discusses the same issue.

[21]As a robustness test, we also constructed three-month-averages of wages before and after a movement to mitigate other sources of reporting error in the months surrounding the transition. This did not affect our estimates.

[22]In Appendix 2.C, we report the same figures for real wage changes. In that case, the share of loss-making transitions increases to roughly one half with average losses of about 15 percent. In principle, the worker should only consider real wages. But in the presence of some wage rigidity the worker expects a wage loss on his current job as well and compares nominal wages.

More than one third of loss-making transitions may seem like a fairly large share at first glance. One possible objection is that wages do not accurately capture the full present value of the new job. As a robustness check, in the segment entitled *Job characteristics*, we exclude transitions from non-unionized to unionized jobs since the latter should have higher expected duration and therefore, potentially, higher present value. This does not materially affect our result. Neither does controlling for observable benefit payments such as moving from jobs without health insurance to jobs that provide insurance or into jobs which subsidize education.[23] Moreover, losses from job to job transitions are a frequent phenomenon across all segments of the wage distribution from top to bottom as can be seen in the segment *Old wage*. They are twice as likely to occur in the upper quartile of the distribution than in the bottom one, as might be expected given that higher wage earners also have more to lose. Still, even in the bottom part, more than 23 percent of transitions end up in lower paying jobs.

We perform a whole battery of further data stratifications to check whether a particular subgroup or time period is driving the results. Their results are reported in detail in Appendix 2.C. Share of losses and conditional changes do not materially change whether we split the sample by year to control for business cycle effects, by gender, age or tenure.[24]

In Appendix 2.C, we also give consideration to an alternative explanation put forward by Postel-Vinay and Robin (2002). They lay out a framework in which workers will accept wage cuts upon job to job transitions, if the option value of working at the other firm is sufficiently high. Indeed, Papp (2012) shows that this framework can rationalize a large amount of wage cuts and large frictional wage dispersion. The key operating mechanism in this class of models is that workers who experienced wage losses have on average steeper observed wage growth afterwards, i.e. wages are backloaded. As we show, there is no indication of that occurring in our data.[25]

---

[23]Given that e.g. Flinn and Mabli (2008) show, also using the SIPP, that wages and non-wage benefits are positively correlated, this should perhaps not be surprising.

[24]One exception occurs when we limit our sample to those individuals who report being paid by the hour. In that case, the share of losses drops to 23 percent and conditional losses to 7.8 percent. Still, this figure appears to understate the phenomenon for the population as a whole, because this group is a highly selective subsample of the population with relatively low wages.

[25]This appears to contradicts the finding of Connolly and Gottschalk (2008) cited earlier. However,

As further piece of evidence that wage losses are the result of transitions into lower ranked jobs, we estimate a probit model conditioning the event of experiencing another subsequent job to job transition on the initial wage change upon movement. Workers who experience a loss making transition are significantly more likely to subsequently transit again. For example, someone having suffered a loss of 20 percent upon movement is 10.3 percent more likely to transit again then someone who experienced an increase of equivalent size and 5.6 percent more likely than someone whose wage remained unchanged.

These different tests lead us to conclude that most of the occurrences of loss-making transitions are not the result of some benefit not properly accounted for by reported compensation. However, we also cannot not conclude that they all result from reallocation shocks. Simple measurement error in wages is surely part of the story. Shocks to workers' idiosyncratic wage potential may be another contributing factor. In Section 2.5, we explicitly include these factors in our model specification in order to quantify the amount of reallocation shocks.

2.4.3 Residual Wage Dispersion in the SIPP

Table 2.3 summarizes measures of residual wage inequality from a regression of log hourly wages[26] on a constant, time dummies, a dummy for disabled workers, a dummy for gender, a dummy for marital status, dummies for race (White, Black, Hispanic, Other), dummies for education (Less than high school, High School, Some college, College), 45 regional dummies, the number of kids, experience and experience square. The mean R^2 of this regressions is 0.37 and the variance of log residual wages is 0.21 leaving a significant share of wage variance unexplained.[27]

---

the authors classify wages into only two categories (low,high) and subsequent wage growth into three categories (low,medium,high). In Appendix 2.C, we show using a continuous wage growth measure that the data suggest no correlation.

[26]See Appendix 2.B for details on how hourly wages are computed.

[27]In an earlier version of this chapter, we also controlled for unobserved individual worker fixed effect similar to Hornstein et al. (2007). The short observation period of 48 months means that many workers do not experience any job to job transition while they are in the sample. As a result, their individual effect captures the full firm effect in wages and the distribution of residual wages has a large mass point at one. We thank an anonymous referee and Tamás Papp for pointing out this issue to us. Nevertheless, we can compare our model results to this statistic when running the same regression on simulated data. Doing so does not change our conclusions drawn in Section 2.7.

Table 2.3: Residual Wage Inequality in the 1993/1996 SIPP

Mm-ratio		Mm-ratio by Age Cohort		Further measures	
Pctl.		Age	5th Percentile	var. log wages	Gini
1^{st}	3.02	25	1.95		
5^{th}	2.14	36	2.12	0.21	0.29
10^{th}	1.83	49	2.25		

The table reports summary measures of residual wage inequality in the merged 1993/1996 SIPP: the mean to minimum ratio, Gini-coefficient and variance of log wages after controlling for worker observables. Since the lowest wage observation in the data is likely the result of measurement error, we report several low percentiles as candidates for the actual minimum wage. Columns 3 to 4 report the Mm-ratio for different age cohorts using the 5th percentile as minimum wage.

The left part of Table 2.3 summarizes the Mm-ratio in the data. Since the lowest wages are likely the result of measurement error, we report a number of low percentiles as candidate points. Independent of the precise measure, the Mm-ratio, the variance of log wages or the Gini coefficient, residual wage dispersion is large and comparable to previous studies.

While regressions like the one above provide a measure for wage inequality among observationally equivalent workers, it is not clear that this should be interpreted as frictional inequality. Such an interpretation would e.g. falsely assign measurement error and unobserved stochastic innovations to individual wage potential to the search friction. The second column highlights a fact extensively analyzed in the incomplete markets literature, e.g. Storesletten et al. (2004), but not often addressed in the existing search literature on wage inequality: Cross-sectional residual inequality increases substantially over the life-cycle. Models with a fixed worker wage potential and no on the job search would imply that inequality does not change with age. Models with on the job search would even predict a decrease in inequality, because workers over time cluster at the higher paying jobs. Therefore, in our model specification, we follow the incomplete markets literature and allow for persistent stochastic innovations to workers' wage potential.

2.5 A Quantitative Model of Wage Dispersion

In this section, we extend our simple model studied in Section 2.3 by adding worker heterogeneity. We enrich the worker's decision problem by a number of empirically

relevant channels that imply larger frictional inequality.[28] We also add stochastic innovations to individual wage potential and measurement error in wages which allows us to disentangle wage losses resulting from reallocation shocks from those resulting from other sources.

The model is set in discrete time. Workers differ in their idiosyncratic log wage potential A_t and draw job offers from heterogeneous jobs with log wage contribution Γ.[29] When a worker of type A_t and a job of type Γ meet, the wage is given by $w_t = exp(A_t + \Gamma)$.[30] We assume that search is random, and unemployed workers contact job offers at rate λ_u in which case Γ is drawn from a distribution with cdf $F(\Gamma)$ on support $[\Gamma_m, \Gamma_M]$. Employed workers continue to sample job offers from the same distribution. Following our discussion in Section 2.3, we model some job to job transitions as the result of a reallocation shocks. An employed worker receives a job offer with probability λ and can in general decide to stay with his old match, or form a new one. However, in λ_d of those cases, the outside option becomes unemployment.

Unemployed workers receive unemployment benefits b_t and a value of leisure Z_t that both depend on the worker's idiosyncratic state:

$$b(A_t) = \min \left\{ b_{max}, rr_b \cdot \mathbb{E}\big[w_t(A_t, \Gamma)|A_t\big] \right\}$$
$$Z(A_t) = rr_Z \cdot \mathbb{E}\big[w_t(A_t, \Gamma)|A_t\big].$$

where b_{max} are statutory maximum UI payments. Averages are taken over the range of acceptable job offers, which themselves depend on A_t. In the case of unemployment insurance, the dependence on the worker's state capture the fact that benefits are a

---

[28]Our focus is on the decision problem of a worker, faces an exogenous job offer distribution. In an earlier version of this chapter, Tjaden and Wellschmied (2012), we used a general equilibrium approach with search and matching in the labor market and a Nash-Bargaining game played by workers and firms. We show that the resulting non-linear log wage schedule can be almost perfectly approximated by a linear one. For ease of presentation, we opt here for the partial equilibrium representation.

[29]Γ is the only source of job effects in our model. These can arise from different job specific productivities, match specific effects and, as Winfried Koeniger pointed out to us, differences arising from bargaining over quasi rents from capital.

[30]Following the existing literature, we assume that wages monotonically increase in the job component conditional on the worker component. Kircher and Eeckhout (2011) and Bagger and Lentz (2012) show that when job effects are independent of match specific effects and the production function has a non-zero cross-partial derivative, bargaining models imply a non-monotone wage schedule, and a specific sorting of workers over firms is an equilibrium outcome. If this was an important aspect of the data, our model would not control for it.

function of prior contributions and workers with higher wage potential contributed more before becoming unemployed. In the case of the value of leisure, we choose this as the closest analogy to the homogeneous agent world.[31]

Workers die with probability ϕ and are replaced by an unemployed labor market entrant whose idiosyncratic log wage potential is drawn from the distribution $N \sim N(\mu_N, \sigma_N^2)$. Burdett et al. (2011) show that introducing experience gains into an on the job search model increases the amount of frictional wage dispersion significantly. To allow for this feature, we let the evolution of workers' wage potential depend on the agent's employment status:

$$A_{t+1} = \begin{cases} A_t + \nu + \epsilon_t & \text{if } employed \\ A_t - \delta + \epsilon_t & \text{if } unemployed. \end{cases}$$

δ represents skill depreciation while being unemployed and ν represents learning on the job. ϵ is a stochastic shock with $\epsilon \sim N(0, \sigma_\epsilon^2)$. We think of shocks to wage potential as demand shocks for specific skills or health shocks. The assumption of a uni-root process in wage potential is in line with most of the labor literature.[32] A non-stationary stochastic specification for wages has also become a standard feature of the incomplete markets literature.[33] It has so far been less common in quantitative search models.

We summarize the worker problem by the value of employment W and the value of unemployment U. The value of employment depends on a worker's wage potential and a firm's wage contribution, the value of unemployment on the workers' wage potential alone. The value of employment reads:

$$W(A_t, \Gamma) = w_t(A_t, \Gamma) + \beta(1 - \phi)\mathbb{E}_t\big\{(1 - \omega)$$
$$\big[(1 - \lambda)H + \lambda[(1 - \lambda_d)\Omega_E + \lambda_d\Lambda]\big] + \omega U(A_{t+1})\big\}$$

$\mathbb{E}_t$ is the expectation operator given all information in period t and ω is an exogenous

---

[31]Furthermore, one can think of this as an, admittedly very stylized, reduced form for capturing wealth heterogeneity. High wage workers tend to have higher asset levels and unemployed workers deplete their assets over time.

[32]See Abowd and Card (1989), Topel (1991), Topel and Ward (1992), Meghir and Pistaferri (2004) and Low et al. (2010).

[33]See for example Krueger et al. (2010).

match destruction shock. For clarity of presentation, we defined the outcome of the choice whether to quit after a bad shock to wage potential as H, the upper envelopes for receiving a regular job offer on the job Ω_E and the upper envelope for receiving a reallocation shock Λ. Let Γ' be the job component at an outside job offer:

$$H = \max\{W(A_{t+1}, \Gamma), U(A_{t+1})\}$$

$$\Omega_E = \int_{\Gamma_m}^{\Gamma_M} \max\{W(A_{t+1}, \Gamma), U(A_{t+1}), W(A_{t+1}, \Gamma')\} dF(\Gamma')$$

$$\Lambda = \int_{\Gamma_m}^{\Gamma_M} \max\{W(A_{t+1}, \Gamma'), U(A_{t+1})\} dF(\Gamma').$$

The value of unemployment solves:

$$U(A_t) = b(A_t) + Z(A_t) + \beta(1 - \phi)\mathbb{E}_t\big\{(1 - \lambda_u)U(A_{t+1})$$
$$+ \lambda_u \int_{\Gamma_m}^{\Gamma_M} \max\{W(A_{t+1}, \Gamma), U(A_{t+1})\} dF(\Gamma)\big\}.$$

2.6 Parameterization

This section proceeds as follows: We first discuss our calibration regarding non-distributional parameters (preferences, institutions, flow rates) in Section 2.6.1. In Section 2.6.2, we discuss our calibration of distributional parameters. Table 2.4 summarizes our calibration.

2.6.1 Non-Distributional Parameters

The model period is one month. When comparing monthly wages in the model to hourly wages in the data, we assume an average of 160 work hours per month. The length of a period is of importance, because it puts an upper bound on the job offer probability λ_u and the minimum duration of an unemployment spell. A maximum of one offer per month is well supported by the data,[34] but the second constraint is likely to be binding.[35]

---

[34]Holzer (1988) reports based on NLSY data that 34 percent of the unemployed received at least one job offer and 12 percent received more than one offer per month.

[35]See Clark and Summers (1979). Our model cannot by construction match the high observed outflow rates within the first month. However, time disaggregation below one month is rather

We calculate the employment to unemployment and unemployment to employment flow rates in our SIPP sample. The exogenous job destruction rate ω is set such that the total job destruction rate, the sum of endogenous and exogenous movements from employment to unemployment, is 0.65 percent per month. We attach to λ_u a value that implies a monthly job finding rate of 12.3 percent.

Information on job to job movements and accompanying wage changes identify λ and λ_d. We adjust λ to imply that 1.43 percent of workers switch employers every period. Our identifying assumption for separating voluntary and involuntary movements is that voluntary movements always result in expected wage increases. Together with the losses due to stochastic idiosyncratic shocks to wage potential and measurement error, both of which are calibrated below, setting λ_d to 0.1 allows us to replicate that 34 percent of job to job movements result in nominal wage losses.[36]

The flow rates estimated from our sample are considerably lower than comparable estimates commonly found in the CPS. In Appendix 2.B, we discuss that this is largely explained by fact that our sample selection criteria lead us to focus on individuals with relatively stable employment histories. Estimated flow rates from our raw sample are considerably larger and comparable to those found in the CPS.[37]

Consistent with findings from Siegel (2002) for average bond and stock returns, we set β to imply a yearly interest rate of 4 percent. Next, we consider the flow value of unemployment. We set the replacement rate rr_b to 25 percent. As argued in Hall and Milgrom (2008) this provides a parsimonious description of the system. The maximum UI benefit payment is set to 1168 \$, which is the average across US states. The parameter determining the value of leisure rr_z is set to 15 percent which yields a total replacement rate of 40 percent when entering into unemployment as in

---

costly, because our numerical algorithm uses value function iteration, which converges at a rate of β.

[36]The share of realized job to job transitions that result from a reallocation shock is 28 percent, which compares nicely with our survey evidence presented in Section 2.4.2. In total, 60 percent of loss making transitions result from reallocation shocks. Our explicit modeling of measurement error and shocks to individual wage potential decrease the estimate of reallocation shocks considerably compared to the studies of Jolivet et al. (2006) and Connolly and Gottschalk (2008).

[37]Moreover, equation (2.2) highlights that for a worker's decision problem only the difference between the on and off the job offer arrival rates matters. Both are significantly lower in our study compared to the ones reported by e.g., Fallick and Fleischman (2004) based on CPS data, but the difference has a comparable size.

Shimer (2005).[38]

We choose an indirect inference approach in calibrating experience and depreciation.[39] In the data, we regress log hourly wages at zero tenure on individual fixed effects, time fixed effects and a quadratic polynomial in experience. The regression yields an average increase in annual wages of 3 percent per year of experience over a working life of 25 years.[40] We then use our model solution to simulate 30000 worker histories and draw a panel of the same length as the SIPP. We perform a similar regression[41] in our simulated data to control for selection and adjust ν to match this statistic. For skill depreciation δ we run a regression of log hourly wages after an unemployment to employment transition on the duration of the previous unemployment spell and worker observables. The results imply that an extra month of unemployment reduces wages by 0.39 percent. We then again replicate this regression in our data and adjust δ to match the regression statistic.

2.6.2 Distributional Parameters

We now describe the way we calibrate the variance of the wage offer distribution σ_F^2, idiosyncratic shocks to wage potential σ_ϵ^2, initial worker dispersion σ_N^2 and the measurement error process. None of the statistics is directly observable in the data because agents endogenously select themselves into and out of employment and into employment with jobs of specific wage offers in response to idiosyncratic productivity developments. Instead, we identify them from within our model.

Measuring Job Heterogeneity

Similar to Low et al. (2010), our identification of the job offer distribution rests on the excess variance of job switchers and job stayers in the data. Other than specifying an additive specification for log wages and assuming the firm contribution to be log normally distributed, this identification only relies on the assumption that

---

[38]The value of leisure is a much discussed object in the literature and Hall and Milgrom (2008) suggest a total replacement rate of 0.71. In Appendix 2.E we show that using this higher rate leaves our results virtually unaffected.

[39]We thank an anonymous referee for suggesting this approach to us.

[40]Altonji and Williams (1998) report very similar results.

[41]Experience is imperfectly measured in the SIPP. Workers are asked how many years they worker at least 6 full months since first entering the labor market. We construct the same measure for yearly experience in our simulated data.

Table 2.4: Calibration

Variable	Target
$\beta = 0.997$	4 percent annual interest rate
$rr_b = 0.25$	$\frac{b_{mean}}{w_{mean}} = 0.25$
$rr_Z = 0.15$	$\frac{Z_{mean}}{w_{mean}} = 0.15$
b_{max}	1168$
$\omega = 6.5 * 10^{-3}$	EU flow rate of 0.0065
$\lambda_u = 0.124$	UE flow rate of 0.123
$\lambda = 0.043$	JTJ flow rate of 0.0147
$\lambda_d = 0.096$	34 percent of wage cuts upon JTJ movements
$\nu = 2.5 * 10^{-3}$	3 percent yearly experience coefficient
$\delta = 2.3 * 10^{-3}$	0.39 percent monthly depreciation coefficient
$\phi = 0.04$	33 years of working life
$\sigma_F = 0.163, \Gamma \sim N(0, \sigma_F^2)$	Equation (2.4)=0.0397
$\sigma_\epsilon = 0.016, \epsilon \sim N(0, \sigma_\epsilon^2)$	Life-cycle wage profile
$\sigma_N = 0.293, N \sim N(\mu_N, \sigma_N^2)$	Life-cycle wage profile
$\sigma_\iota = 0.119, \iota \sim N(0, \sigma_\iota^2)$	Estimation
$\mu_N = 5.618$	Mean monthly wage 2139$

The left column states the calibrated variable with its value and the second states the relevant moment. EU stands for employment to unemployment, UE for unemployment to employment, and JTJ for job to job.

measurement error for job switchers is the same as for job stayers. Appendix 2.D provides evidence for this assumption.

In our SIPP data, we assume that wages are generated by:

$$ln(w_{i,t}) = \alpha_0 + \alpha_1 d_t + \alpha_2 Z_i + \beta_2 \Gamma_i + e_{i,t} \qquad (2.3)$$

where d_t captures aggregate states, such as TFP and Z_i is a vector of idiosyncratic components. We split the unobservable $e_{i,t}$ into two parts:

$$e_{i,t} = r_{i,t} + A_{i,t}$$

Like in the model $A_{i,t}$ is assumed to follow a random walk with drift and innovations $\epsilon_{i,t}$, and $r_{i,t}$ captures measurement error. For our present purpose, we have to make no further assumptions regarding the distributional properties of measurement error.

First-differencing eliminates the idiosyncratic wage components. As mentioned above, we only observe a self-selected subset of the realizations of Γ and ϵ as agents

can quit into unemployment after negative idiosyncratic shocks and refuse wage offers. The subsets of observed realizations Γ^{obs} and ϵ^{obs} are themselves random variables which follow distributions of unknown functional forms. However, we can use the workers' decision rules, which determine for each (A_t, Γ) combination whether to form or continue a match, to map these moments back into the structural parameters.

Define observed wage growth when a job to job transition takes place

$$\Delta ln(w_{i,t}^b) = \nu + \kappa_t + [\Gamma_i^{obs} - \Gamma_{i-1}^{obs}] + \epsilon_{i,t}^{obs} + \Delta r_{i,t}$$

and when no such transition takes place

$$\Delta ln(w_{i,t}^w) = \nu + \kappa_t + \epsilon_{i,t}^{obs} + \Delta r_{i,t}$$

where $\kappa_t = \alpha_1(d_t - d_{t-1})$. After regressing out a constant and time dummies, we obtain the residual excess variance of job movers relative to job stayers:[42]

$$Var\left[\Delta ln(\hat{w}_{i,t}^b)\right] - Var\left[\Delta ln(\hat{w}_{i,t}^w)\right]$$
$$= Var\left[\Gamma_i^{obs} - \Gamma_{i,-1}^{obs}\right] + Cov\left[\epsilon_{i,t}^{obs}(\Gamma_i^{obs} - \Gamma_{i,-1}^{obs})\right] \quad (2.4)$$

where we have invoked the assumption that measurement error is uncorrelated with the event of job switching.

Equation (2.4) also holds in our model and we use it as a calibration target for σ_F^2. The endogenous sorting that causes the observed distribution in the data to differ from the true one is also present in our model.

Calibrating Idiosyncratic Wage Potential

Similar to Storesletten et al. (2004), we calibrate the variance of idiosyncratic wage shocks to the life-cycle profile of cross sectional residual wage dispersion.[43] While we explicitly model initial worker heterogeneity and experience gains, the data possesses

---

[42]We delete the top and bottom 0.5% of the wage growth observations to get rid of reporting error.

[43]In principle, we could derive a moment condition similar to the one above to identify idiosyncratic wage uncertainty (see Meghir and Pistaferri (2004) for more details). Whereas the identification of the job component only required two consecutive wage observations, the maximum spell length of 48 months in the SIPP now becomes more of on an issue which is why we opt for a different approach.

well-known idiosyncratic wage components absent from our model that we regress out (gender, race, marriage, number of children, disability and time dummies).[44] We then choose σ_N^2 to match the initial variance of residual log wage inequality and σ_ϵ^2 to match its increase over the life cycle.

Lastly, wage fluctuations may result from measurement error. To accurately identify the share of reallocation shocks and to properly calibrate the innovations to individual wage potential, we require an explicit treatment for this source of wage fluctuations. At this point, we need to make further assumptions regarding its statistical properties. Appendix 2.D shows that the autocovariance function of within job wage growth goes to zero at longer lags. We therefore follow Meghir and Pistaferri (2004) and postulate an $MA(q)$ process (i.e. $r_{i,t} = \Theta(q)\iota_{i,t} = \iota_{i,t} - \sum_{j=1}^{q} \theta_j \iota_{i,t-j}$). The autocovariance function is close to zero after 12 lags, such that we fix q at 12. Assuming $\mathbb{E}(\epsilon_{i,t}^{obs}\epsilon_{i,t-j}^{obs}) = 0 \; \forall j \neq 0$, we obtain the parameters $\Theta(12)$ and σ_ι using Maximum Likelihood estimation and Kalman filtering.[45] Appendix 2.D supplies further detail on the procedure and shows that θ_{12} is indeed estimated close to zero.

2.7 Results

We now present the main results of this chapter. In Section 2.7.1 we demonstrate that our model generates residual wage dispersion of the size estimated in the data and that it matches its life-cycle profile. Moreover, the model provides a close fit to the shape of the overall wage distribution. Section 2.7.2 discusses the structurally inferred parameters of the wage offer distribution and of idiosyncratic wage uncertainty. We then go on to determine the relative contributions of job dispersion, development in workers' wage potential and the distribution of workers over jobs to overall wage dispersion. Our results attribute 13.7 percent of wage inequality to the presence of the search friction. Using an alternative model without reallocation shocks, the estimate jumps up to the size previously estimated in the data.

---

[44]We purify our data of these effects, which are well-known drivers of wages, because we think them inadequately represented by our model set-up. Gender and race biases are likely the result of discrimination. Marriage stands in for a joint labor supply decision absent from our model. Disability and the number of children likely do represent productivity, but not in a way adequately captured by our model.

[45]We thank Johannes Pfeifer for providing the Kalman filtering routine to us.

2.7.1 Empirical Fit

We simulate a cohort of 30000 workers over their life-cycle. From the resulting individual paths we sample 48 month observation spells to generate a data set of the same length as the SIPP. We then run a regression of log wages on a constant and experience to calculate the model counterpart to our measure of residual wages in the data. Table 2.5 summarizes our results.

Table 2.5: Residual Wage Dispersion

		Mean-Min Ratio		Gini		$Var(\log(\tilde{w}_{it}))$	
		Model	Data	Model	Data	Model	Data
	1^{st}	3.01	3.02				
Pctl.	5^{th}	2.21	2.14	0.24	0.29	0.18	0.21
	10^{th}	1.89	1.83				

The table compares the size of the residual wage dispersion generated by our baseline specification to the one found in the 1993/1996 SIPP. The first two columns report the Mm-ratio in the model and the data using the 1st, 5th, and 10th percentile as possible minimum wages. As further summary statistics, we compare the Gini coefficient and the variance of log wages.

The mean residual wage paid is 3.01 times the smallest observation evaluated at the first percentile. When looking at higher percentiles, model and data line up closely as well. Other summary statistics of inequality also indicate a good fit: the Gini coefficient and the variance of residual log wages are slightly smaller, but close to those found in our data set.

In Section 2.4.3, we discussed that a characteristic feature of residual inequality is its increase over the life-cycle and used the fact to motivate our stochastic wage potential process. Figure 2.3a compares the model to the data along that dimension. We closely match the magnitude of the increase over the life-cycle, while missing the concave shape at the end.

In our subsequent analysis, we use our model to compute the contribution of search induced wage inequality to overall wage inequality in the population cross-section. We therefore need to verify that our model fits the data along that dimension. As discussed previously, there are a few well-known wage determinants in the data that our model is not designed to include. Therefore, in what follows, we first regress log wages in our data on a constant and dummies for disability, gender, marriage status, the number of kids, time and race. These factors account for 13.3 percent

Figure 2.2: Model Fit

(a) Mm-ratio over the Life-Cycle
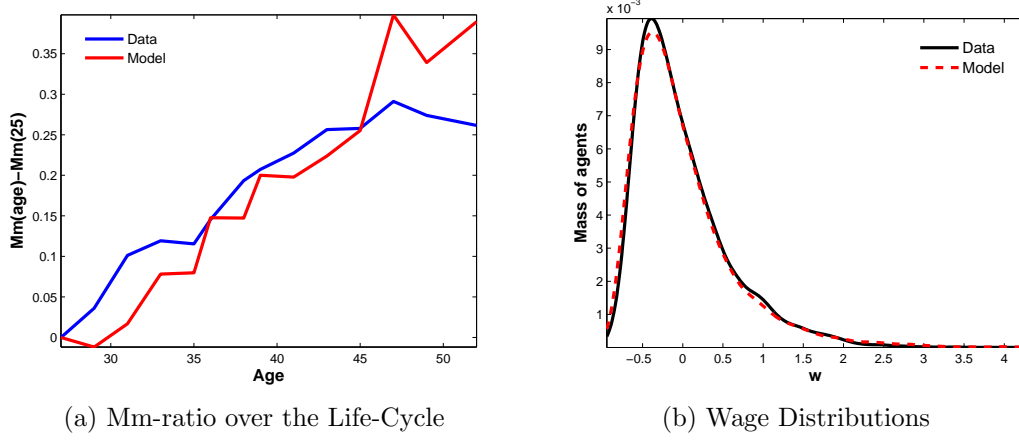
(b) Wage Distributions

Figure 2.3a plots the Mm-ratio by age in the model against the data. Figure 2.3b compares demeaned density functions of wages after applying a kernel smoother.

of log wage variation. We compare the wage distribution from our model to the resulting distribution. Figure 2.3b plots the kernel estimator of the density function of wages after transforming the data back to levels against its model counterpart. The two graphs match up almost perfectly well. There is substantial inequality and the distribution features the characteristic right skew.

2.7.2 Underlying Sources of Inequality

Confident that our model features the main determinants of wage inequality, we use it to infer the relative importance of differing initial abilities (σ_N, in our model), uncertainty of idiosyncratic wage potential (σ_ϵ), the search friction (σ_F) and a sorting term to be introduced below in explaining overall wage inequality. Our calibrated parameters are displayed in the first line of Table 2.6. The value for σ_ϵ implies an annual standard deviation for the permanent component of wages of 0.06. To put our results into perspective, Low et al. (2010), using the 1993 SIPP, estimate a standard deviation for the job offer distribution of 0.23 and of 0.103 for annual idiosyncratic innovations. Both our standard deviations are smaller, which can largely be explained by different sample selection criteria.[46] Their relative sizes, however,

---

[46]Most importantly, our exclusion of individuals with imputed wages reduces our estimates compared to theirs. The total amount of wage inequality is also lower in our data. Hourly wages

Table 2.6: Wage Offer Distribution and Idiosyncratic Risk

Specification	σ_F	σ_ϵ	σ_N	λ
Baseline	0.163	0.016	0.293	0.043
job ladder model $(\lambda_d = 0)$	0.296	0.017	0.117	0.1

The table displays the standard deviations of the wage offer distribution and of the idiosyncratic wage shock. The first line refers to the baseline specification and the second one to a calibration of a "pure" job ladder model.

are surprisingly similar given the very different identification strategies.

Our model implies a direct link between observed wage outcomes and these deep parameters. In order to map it out, we use our simulated data and consider the following variance decomposition, which we separately estimate for each age group in our simulated data

$$Var(ln(w_i)) = Var(A_i) + Var(\Gamma_i) + 2Cov(A_i, \Gamma_i) + Var(r_i).$$

The left panel of Figure 2.3 illustrates the results. For young workers, job heterogeneity explains about 24 percent of overall log wage variance but that number drops as workers' employment histories become more diverse. Our model identifies worker heterogeneity as the dominant factor in explaining variations in wages and this effect is increasing in age. Measurement error is responsible for about 2.4 percent of variation. Sorting of workers over job types has a mild positive effect. In a population weighted average, frictional wage dispersion accounts for 15.5 percent of wage inequality within our model. Given that we eliminated 13.3 percent of wage variation through our fixed effect regression, this implies frictional inequality to account for 13.7 percent of overall wage inequality present in our data.

2.7.3 On the Job Search and Structural Inference

Previous estimates from structural search models that try to disentangle the contributions of worker and job effects to wage variation imply a much lager role for the latter than we do (Postel-Vinay and Robin (2002) suggest numbers up to 50% and Carrillo-Tudela (2012) finds numbers around 40%.). In this section, we investigate whether the introduction of the reallocation shock alone can explain the large quan-

---

based on a yearly measure have a standard deviation of 7.66 in their data. In our case, this drops to 7.03.
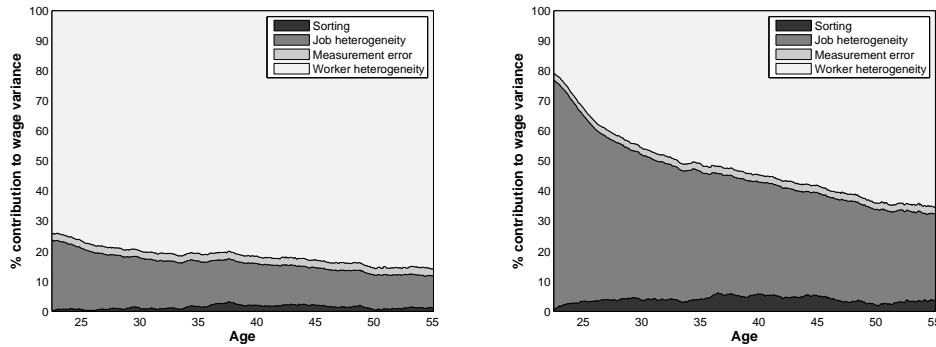
Figure 2.3: Contribution of Search Frictions to Overall Wage Dispersion Baseline v. job ladder Model

The graphs display the cumulative contribution of sorting (black area), firm effects (dark grey area) measurement error(medium grey area) and worker heterogeneity (light area) to the variance of log wages, conditional on age. The left panel is from our baseline specification, the right panel results from a job ladder model with idiosyncratic productivity risk.

titative discrepancy. We also highlight how the mechanisms outlined in Section 2.3 interact when we identify the variance of the job offer distribution.

We re-calibrate our baseline model to a more common job ladder model setting, $\lambda_d = 0$, and neglect wage losses upon transition as calibration target. With a Mm-ratio of 3.45 at the first percentile, the model yields a residual inequality of similar size as our baseline specification. To demonstrate that measurement error and stochastic wages alone cannot account for the stylized facts outlined in Section 2.4.2, we compare moments of wage dynamics upon job to job movement in the data to our our baseline specification and the job ladder-model. Table 2.7 displays the results.

In the data, job to job movements on average result in wage gains of 3.3 percent. Conditional on suffering a wage loss upon movement, workers lose 19.6 percent of

Table 2.7: Wage changes from job to job movements

Specification	Avg. gain	Avg. loss
Data	0.033	-0.196
Baseline	0.071	-0.186
job ladder model $(\lambda_d = 0)$	0.227	-0.09

The table compares the model baseline specification with a pure on the job search version in their implications for job to job transitions. Statistics are the resulting average wage gain upon job movement and the average wage loss, conditional on observing a loss. *Data* refers to computation from the SIPP for nominal wages.

their previous wages. Our baseline specification fares quite well in reproducing these statistics. Wage gains are too high, but the order of magnitude is comparable. The model does well in reproducing the large conditional wage losses. In Appendix 2.E, we show that our baseline specification is also in line with the large initial wage gains at job to job transitions reported by Topel and Ward (1992) and the convex decrease of these gains over experience. In the pure job ladder model, average wage gains at job to job transitions of 23 percent are much too large compared to the data. Since workers in this model only transit to higher ranked jobs, the wage losses are only observed as result of a negative shock to individual wage potential or due to measurement error. A conditional 9 percent average wage loss clearly fails in this respect. We come back to this fact below.

We now investigate what these differences imply for the inferred importance of difference sources of wage inequality. The right panel of Figure 2.3 shows that this model paints a much changed picture of the different sources of wage inequality, when compared to our baseline specification. The cross-sectional average for the contribution of frictional wage dispersion more than doubles to about 44 percent (38.8% percent of wage variation in the data) with values as high as 78 percent for the youngest workers. Closely related is an almost doubling in the inferred standard deviation of the wage offer distribution as can be seen in the second row of Table 2.6.

The reason for these results can be traced back to the role of reallocation shocks. Section 2.3 demonstrated that in the absence of reallocation shocks, the inferred job offer arrival rate on the job is higher and more workers are in the right tail of the job offer distribution. Table 2.6 shows that our recalibrated model implies an on the job offer arrival rate more than twice as large as our baseline calibration. Consequently, workers quickly move into very high ranked matches, accept further outside offers only infrequently and wage improvements are relatively small. Since they also do not experience large losses when moving, the implied wage offer distribution has to spread out substantially to reproduce the observed excess variance for job switchers. On the flip side, most initial dispersion is explained by job effects and the inferred initial worker heterogeneity drops by more than half in terms of its standard deviation. The two model versions tell rather different stories about the sources of life-time wage inequality. As a robustness analysis, we decrease the share of reallocation shocks

exogenously by a half in Appendix 2.E. The variance decomposition leads to results close to our baseline case, showing that already some reallocation shocks overturn the strong implications from the pure job ladder model.

2.8 Conclusion

We solve a rich structural model of job and worker heterogeneity to quantify the importance of the search friction in generating wage inequality. Our model features several major channels that expand the range of acceptable offers to the workers creating larger *frictional* inequality: skill accumulation on the job, skill loss in unemployment and search on the job. The baseline calibration reproduces both overall and residual wage inequality. Nonetheless, the search friction accounts for only 13.7 percent of total inequality.

The large quantitative difference to previous estimates stems from our introduction of reallocation shocks upon job to job transitions. These shocks allow our model to match a large job to job transition rate in the data with a relatively low on the job offer arrival rate. As a consequence, the endogenous wage distribution features few workers at high ranked jobs. The calibrated variance of the job offer distribution is relatively small and only a small share of wage variation can be explained by job differences.

Empirically, we provide various pieces of evidence to show that reallocation shocks provide a fitting description for about a quarter of observed job to job transitions. Most importantly, about one third of all job to job transitions end up with lower nominal wages than on the previous job. This finding is robust to both controlling for observed benefit payments as well as all kinds of data stratification.

Appendix to Chapter 2

2.A Solving the Simple Model

This section derives implicit solutions for the minimum wage, the mean wage, the wage distribution and the relationship between job to job transitions and the job offer rate for the model presented in Section 2.3.

Recall the worker problem:

$$rW(w) = w + \lambda(1 - \lambda_d) \int_w^{w_{max}} [W(z) - W(w)]dF(z)$$

$$+ \lambda\lambda_d \int_{w^*}^{w_{max}} [W(z) - W(w)]dF(z)$$

$$- (\omega + \lambda\lambda_d F(w^*))(W(w) - U)$$

$$rU = b + \lambda_u \int_{w^*}^{w_{max}} [W(z) - U]dF(z),$$

where $F(w)$ is the cdf of the wage offer distribution with upper support w_{max}, λ is the job offer arrival rate on the job, λ_d is the share of reallocation shocks, ω is the job destruction rate and λ_u the job offer arrival rate during unemployment. Evaluating the asset value of employment at w^* and setting it equal to the asset value of unemployment yields:

$$w^* = b + (\lambda_u - \lambda) \int_{w^*}^{w_{max}} W'(z)[1 - F(z)]dz.$$

Differentiating the asset value of employment with respect to w yields

$$W'(w) = \frac{1}{\omega + \lambda\lambda_d F(w^*) + r + \lambda\lambda_d + \lambda(1 - \lambda_d)[1 - F(w)]}$$
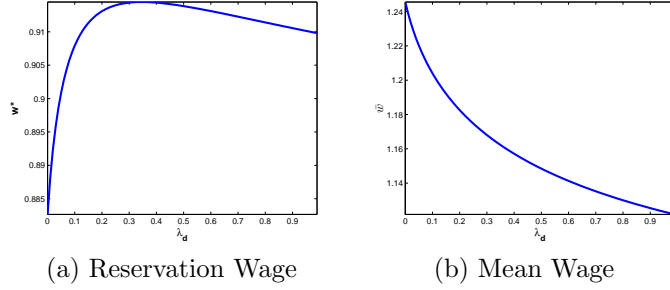
We therefore obtain an implicit solution for the reservation wage reported in Section 2.3:

$$w^* = b + (\lambda_u - \lambda) \int_{w^*}^{w_{max}} \frac{1 - F(z)}{r + \omega + \lambda\lambda_d F(w^*) + \lambda\lambda_d F(z) + \lambda[1 - F(z)]}dz. \quad (2.5)$$

Figure 2.5a highlights the non-monotone relationship between λ_d and w^* discussed

Figure 2.4: Parameterized Simple Model (2)

(a) Reservation Wage   (b) Mean Wage

The figure displays the relationship between the share of reallocation shocks, λ_d, the minimum wage and the mean wage for the calibration performed in Section 2.3.

in Section 2.3.

We now derive an implicit solution for the wage distribution $G(w)$. A stationary distribution of employment over wages implies:

$$(1-u)G(w)[\omega + \lambda\lambda_d F(w^*) + \lambda[1 - F(w)]] =$$
$$u\lambda_u[F(w) - F(w^*)] + (1-u)\lambda\lambda_d[1 - G(w)][F(w) - F(w^*)] \quad (2.6)$$

Rearranging yields

$$G(w) = \frac{u\lambda_u + (1-u)\lambda\lambda_d}{1-u} \frac{F(w) - F(w^*)}{\omega + \lambda[1 - F(w)] + \lambda\lambda_d F(w)}.$$

Evaluating (2.6) at w^{max} yields

$$\frac{u}{1-u} = \frac{\omega + \lambda\lambda_d F(w^*)}{\lambda_u[1 - F(w^*)]}.$$

Substituting into (2.6) gives the solution for $G(w)$:

$$G(w) = \frac{F(w) - F(w^*)}{1 - F(w^*)} \frac{\omega + \lambda\lambda_d}{\omega + \lambda\lambda_d F(w) + \lambda[1 - F(w)]}. \quad (2.7)$$

We now derive an implicit solution for the relationship between λ and the job to job transition rate that we omit in the main chapter for parsimony. Total job to job

flows are given by:

$$JTJ = \lambda\lambda_d[1 - F(w^*)] + \lambda(1 - \lambda_d)\int_{w^*}^{w_{max}}[1 - F(z)]dG(z).$$

Integrating the equation by parts yields

$$JTJ = \lambda\lambda_d[1 - F(w^*)] + \lambda(1 - \lambda_d)\int_{w^*}^{w_{max}} G(z)dF(z)$$

Substituting in $G(w)$ gives

$$JTJ = \lambda\lambda_d[1 - F(w^*)]$$
$$+ \lambda(1 - \lambda_d)\frac{\omega + \lambda\lambda_d}{1 - F(w^*)}\int_{w^*}^{w_{max}}\frac{F(z) - F(w^*)}{\omega + \lambda\lambda_d + \lambda(1 - \lambda_d)[1 - F(z)]}dF(z).$$

Replace $z = F(z)$ to obtain

$$JTJ = \lambda\lambda_d[1 - F(w^*)]$$
$$+ \lambda(1 - \lambda_d)\frac{\omega + \lambda\lambda_d}{1 - F(w^*)}\int_{F(w^*)}^{1}\frac{z - F(w^*)}{\omega + \lambda\lambda_d + \lambda(1 - \lambda_d)[1 - z]}dz. \quad (2.8)$$

Solving the integral yields:

$$\int_{F(w^*)}^{1}\frac{z - F(w^*)}{\omega + \lambda\lambda_d + \lambda(1 - \lambda_d)[1 - z]}dz =$$
$$\left| -\frac{\lambda(1 - \lambda_d)z + [\omega + \lambda]log(\omega + \lambda\lambda_d + \lambda(1 - \lambda_d)[1 - z])}{[\lambda(1 - \lambda_d)]^2}\right.$$
$$\left.+ \frac{F(w^*)log(\omega + \lambda\lambda_d + \lambda(1 - \lambda_d)[1 - z])}{\lambda(1 - \lambda_d)}\right|_{F(w^*)}^{1}.$$

Finally, we can derive a solution for the mean wage:

$$\bar{w} = \int_{w^*}^{w_{max}} wdG(z).$$

Integration by parts yields

$$
\begin{aligned}
\bar{w} =& w^{max} - \int_{w^*}^{w_{max}} G(z)dz \\
=& [w^{max} - w^*] + w^* - \int_{w^*}^{w_{max}} G(z)dz \\
=& w^* + \int_{w^*}^{w_{max}} [1 - G(z)]dz \\
=& w^* + \frac{\omega + \lambda - \lambda(1 - \lambda_d)F(w^*)}{1 - F(w^*)} \int_{w^*}^{w_{max}} \frac{1 - F(z)}{\omega + \lambda\lambda_d + \lambda(1 - \lambda_d)[1 - F(z)]}dz,
\end{aligned}
$$

which is an implicit solution for $\bar{w}$. Figure 2.5b shows the resulting downward sloping relationship between λ_d and λ. Upon inspection to the mean and minimum wage, it becomes apparent that their ratio is not a moment independent of $F(w)$ in our model with reallocation shocks.

2.B Creating the Data Set

2.B.1 Aligning the 1993 and 1996 SIPP

Our two samples from the SIPP differ regarding data collection and sample size. Unlike the 1993 sample, the 1996 SIPP uses *computer-assisted interviewing techniques* to increase data quality. The computer assures that employer identification numbers stay constant across interviewing waves. Moreover, the 1996 SIPP uses dependent interviewing across waves with respect to employer IDs asking: *"Last time we recorded that you worked for [Employer name]. Do you still work for [Employer name]?"*. Both features likely reduce misreporting in employer changes. In the 1993 sample, interviewers assign employer IDs manually for each wave and use no dependent interviewing across waves. To address the issue in the 1993 sample, we use employers IDs constructed by Stinson (2003) which combine the survey data with administrative records to accurately identify these changes.

The 1996 sample also has a considerably larger initial sample size, providing information on 95,402 sample members compared with 56,800 in the 1993 SIPP.[47]

---

[47]The 1993 SIPP was the last sample that published a *Full Panel Longitudinal Research File*. The imputation methods in the *Core files* do not use longitudinal information for imputation purposes and include records for individuals that did not respond in a given wave. We circumvent these problems by using only records that appear in the *Full Panel Longitudinal Research File*

Some information we use from both panels is grouped differently in the 1993 and 1996 SIPP. First, the grouping for the state of residence provides somewhat more detailed information for smaller states in the 1996 panel. Second, the 1993 panel contains monthly information on membership in the armed forces. This information is only available on a 4 months basis in the 1996 panel. We therefore have to drop entire individual waves from the 1996 SIPP when the individual reports to have been member of the armed forces during that time. For those readers interested in more of the details of sample creation and sample selection than are provided here and in the next section, STATA and Matlab codes for all our empirical work is available for download on the author's web pages.

2.B.2 Calculating Hourly Wages and Sample Selection

The SIPP asks respondents whether they are paid by the hour and their corresponding hourly pay rate in each month. We use this hourly pay rate whenever it applies. The SIPP also reports total monthly earnings per job, whether the job lasted the entire month and the number of hours worked per week. When computing monthly earnings of those workers that are not paid by the hour, we assume that workers do not alter their earnings response based on the length of a month and use smooth 4.3 weeks per months.[48] SIPP records starting date and end date of each job that does not last the entire month. We use this information to calculate hourly wages for those months.

As we also report in the main text, we select our sample from the original merged 1993 and 1996 SIPP by using only observations from individuals aged 23-55 (prime working age), for whom we require complete information on the individual's employment status, age and employer id. We only consider an individual's primary job and drop workers that are recalled by former employers[49] or have missing reporting months during a job spell. Moreover, we drop workers reporting to be school enrolled, the self-employed, family-workers, members of the armed forces, workers at non-profit companies and anyone whose wage information was imputed by the SIPP.

---

    as suggested by the CENSUS Bureau.

[48]Neither the reported wage, nor earnings or hours worked are reported by dependent interviewing across waves in either of the two samples.

[49]We chose to exclude those observations because recalled workers likely possess a different search technology than what is represented in our model specification.

Table 2.8: Comparing Sample data to Original Data

	Mean wage	EE rate	$Var(\ln(w_{it}) - \ln(w_{it-1})))$
SIPP	13.71	2.18%	0.1223
Sample	13.43	1.43%	0.0548

The table compares our sample data to the original SIPP data. The first column reports the nominal mean wage, the second the rate of job to job transitions and the third column the variance of log-wage growth.

Finally, we truncate the wage distribution at the top and bottom 1% to take care of outliers and top-coding.[50] These restrictions leave us with 2,039,345 person/month observations.

In particular our choices of excluding individuals who are recalled by former employers and those with imputed wages introduce changes to the data set which have bearing on some of our key calibration targets. Table 2.8 compares our final data set to the original, non-stratified SIPP samples. Mean hourly wages are almost identical. The aforementioned exclusions limit our data to a sub-sample of the population that has relatively stable work profiles. Job to job transition rates are considerably lower in the stratified data and lower than rates usually reported from CPS data. Moreover, the variance of log wages at job to job transitions is considerably lower in our final sample. Here, the exclusion of imputed wages along with the truncation of the wage distribution at the bottom are largely responsible for the more than fifty percent reduction.

2.C More on the Empirics of On the Job Search

2.C.1 Measuring Job to Job Flows

In order to calibrate the job offer arrival rate on the job, it is crucial to accurately identify job to job transitions in the data. One of the biggest advantages in working with SIPP data is that workers are asked to report an employment status for each week of the reporting period separately. This allows us to identify any unemployment spell lasting longer than one workweek.

In a given month we count as employed someone who reports holding a job for

---

[50]Earnings are topcoded at \$33333 and \$50000 for a four month period in the 1993 and 1996 sample, respectively.

the entire month. This definition includes paid as well as unpaid absences as result of vacations, illnesses or labor disputes. It does exclude, however, those who report having been on layoff for at least a week. There is no standard definition for job to job movements in empirical work. We therefore experiment with several different definitions. Our first measure is analogous to the definition in Fallick and Fleischman (2004) and equates job to job transitions with firm changes. We use a monthly employer identifier based on company names. We refer to this definition by $JTJ1$. Given that a firm is a match in our model and given that employees may transit between jobs within a given firm, we find it useful to somewhat broaden the concept beyond employer id changes. For $JTJ2$ we therefore follow Moscarini and Thomsson (2007) in identifying job to job movements by changes in the three digit occupational code. Moreover, we define $JTJ3 = JTJ1 \cup JTJ2$ as the union set from the two definitions.

Table 2.9: Different Definitions of JTJ Flow Rates

(a) Full 1993/1996 Merged SIPP Data Set

	$JTJ1$	$JTJ2$	$JTJ3$	CPS
JTJ	1.75	1.27	2.18	2.29
Share loss	36.2	37.2	35.3	
Ave. loss	23.8	23.6	22.6	

(b) Sample for Baseline Calibration

	$JTJ1$	$JTJ2$	$JTJ3$
JTJ	1.03	0.73	1.43
Share loss	36.4	34.4	34.4
Ave. loss	20.1	19	19.6

The Table shows percentage probabilities for job to job transitions based on SIPP data from 1993 to 1999. For reference we also cite monthly averages from Fallick and Fleischman (2004) using CPS data for the years 1994–2003 for workers between 25 and 54. The different flow definitions can be found in the text. Panel (a) reports figures for the whole merged 1993/1996 SIPP data set before applying any sample selection criterion. Panel (b) refers to the sample underlying our baseline calibration from which most other numbers in this chapter are computed. *Share loss reports* the percentage of *EE* transitions which result in lower hourly wages under the given flow definition. *Ave. loss* reports the corresponding conditional average wage loss.

Table 2.9 reports job to job flow rates based on the different definitions. For

comparison, we also report averages from monthly estimates for the years 1994-2003 taken from Fallick and Fleischman (2004), who use CPS data for individuals ages 24 to 54. As we noted in Appendix 2.B.2, our choice of excluding individuals with imputed wages and those who are recalled by former employers is equivalent to restricting the sample to people with relatively stable employment profiles, which lets EE-rates drop under any definition. For better comparability with the numbers based on the CPS, in Panel (a) we therefore first report estimates of flow sizes from our raw sample. Identifying job to job movements by either employer changes or changes in the occupational code alone yields roughly comparable flow sizes. However, only our broadest definition of job to job employment transitions comes close to the magnitude found in the CPS. Reassuringly, the share of transitions yielding lower hourly wages and their conditional average loss are very similar regardless of the definition used. This still holds true when repeating the estimation on our final sample in Panel (b). Our calibration is based on the 1.43 percent probability found when applying $JTJ3$ on our baseline sample. It is the only definition which, using the raw sample, yields estimates of comparable size to those from the CPS which many other studies use. Regarding the decrease after applying sample selection criteria, we are confident that this figure is more representative for the kinds of transitions included in our model environment.

2.C.2 Wages and On the Job Search

We argue in the main chapter that the magnitude of job to job flows in itself is insufficient to identify the on the job offer arrival rate. Instead, the question is how many of these job changes actually yield higher wages for the worker. In the main text, in Section 2.4.2, we established as stylized fact from our data that about one third of all job to job transitions yield lower nominal wages for the worker than his previous job. In this section, we establish the robustness of this result by considering a number of different data stratifications to demonstrate that those wage cuts are not driven by any subsample of the population but instead extend across workers of all kinds. All results are summarized in Table 2.13.

Wage Gains from Employment Changes

First, we use CPI data to deflate monthly wages. The share of losses in wage changes increases to 47 percent while the mean loss reduces to 14.6 percent. Evidently, many transitions result in unchanged nominal wages between jobs. Of course, the worker should only care about real wages in making his decision. Meanwhile, an argument can be made that in the presence of some wage rigidity, the worker expects a real wage loss on his current job as well and therefore compares nominal wages. Results are unaltered when trying to control for benefit payments at the new job. We subsequently exclude from the sample transitions from non-Union to unionized jobs, transitions into jobs that provide health insurance, and transitions into jobs that provide educational subsidies. None of these modifications significantly changes the moments of interest.

Next, we consider potential business cycle effects by splitting our sample into different years. The willingness of workers to accept a wage reduction upon transition might depend on the aggregate state of the economy. Between, 1993 and 1999, the time of our sample, the US economy experienced the longest uninterrupted expansion in post-WWII history which was to last until March 2001. Yet, we again observe no significant variation in the share of losses or the size of losses between years.

Women are known to have less stable work relationships than men and might therefore be responsible for an overproportional share of loss making job to job transitions. Nonetheless, in the data both sexes have an equal probability of experiencing a wage cut after moving. The same holds for stratifications by age groups. Young workers have a looser attachment to the labor market and may initially experiment with different career paths or search for jobs with higher non-monetary benefits. But none of these phenomena cause the youngest age group to experience markedly more job to job transitions with wage losses.

We also stratify our sample by earnings and tenure. We split the main sample into its lowest and highest quartile and the observations in-between. Again, we do not expect the outcome to be random, because high wage earners are more likely to incur a loss when they are forced to look for alternative employment. Nonetheless, low wage earners are far from insulated to wage losses when switching jobs and even in the lowest quartile, 23 percent of all job to job transitions result in nominal wage losses. Considering tenure is informative in two ways. For one, one might

hypothesize that a subsample of the population with a loser labor market attachment who never accumulate longer periods of tenure are disproportionately responsible for the observed job losses. Still, wage losses upon transition are a pervasive phenomenon across all of the tenure distribution. Alternatively, one might be assume that the observed losses are the result of losses in match-specific capital for high-tenured workers. In his case, they should be increasing in tenure on the previous job. This hypothesis, also, is not borne out in the data.[51]

Lastly, we restrict our sample to workers who report being paid by the hour. This might help to rule out potential measurement error resulting from hour calculations of hourly wages for people where only earnings are reported. In that case, the share of losses drops to twenty percent and conditional losses to seven. Still, this figure appears to understate the phenomenon for the population as a whole. First, the group of workers paid by the hour is a highly selective subsample of the population with relatively low wages. Mean hourly wages in the SIPP are $13.5, but drop to $11.1 within that group. Second, we are interested in total worker compensation. When workers are asked about their hourly pay rate, the question reads: "What was your regular hourly pay rate at the end of month X". Hence, respondents are unlikely to include any bonuses or performance payments. Contrary, when asking about total monthly earnings, the question explicitly states: "Be sure to include any tips, bonuses, overtime pay, or commissions".

Alternative Explanations

Postel-Vinay and Robin (2002) propose an alternative explanation for those wage losses. They lay out a model where wages can only be renegotiated by mutual agreement and the firm has all the bargaining power. Wage raises on the job occur as a result of counter-offers to bids by other firms. They demonstrate that in such a framework workers will accept wage cuts upon job to job transitions, if the option value of working at the other firm is sufficiently high. Workers only move to higher ranked firms than their current employer and very productive firms offer the potential

---

[51]Unfortunately, the tenure measure is of low quality in the SIPP. Of those being employed in the first month of the interview, more than 8 percent of workers report not having been employed with their current employer previously implying unrealistically high worker turnover rates. Moreover, the tenure variable is employer specific in the SIPP and not linked to a job as in our model.
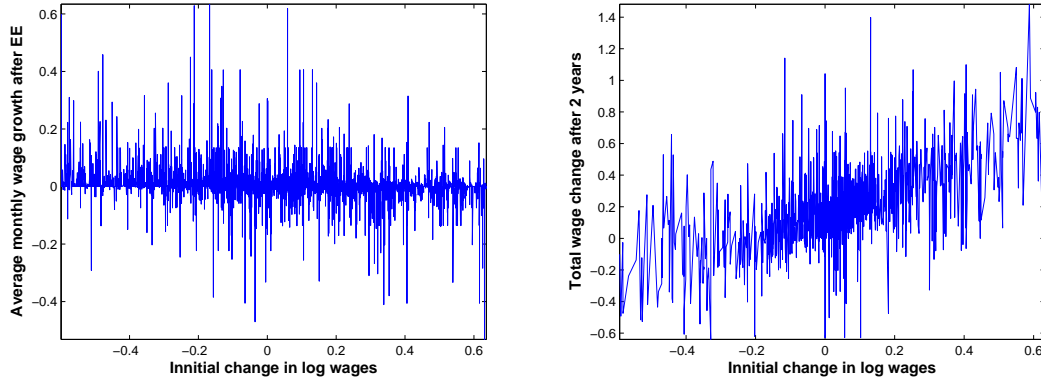
of large future wage gains.

Figure 2.5: Initial Wages Change and Subsequent Wage Growth

The left panel plots cumulative wage growth in the two months after a job to job movement against the initial wage change, excluding the latter from the calculation. The figure was generated using all observed job to job transitions. In the right panel, we only include job to job transitions where the worker was subsequently observed for at least 24 months. The cumulation of wage growth now includes the initial change upon transition.

A testable implication of these types of models is that expected future wage growth with the new employer should be an increasing function of the wage cut accepted. The left panel of Figure 2.5 plots cumulative wage growth with the new job against the initial wage change for our population of job to job transitions. There is no relationship between the initial wage change and consecutive wage growth. In the right panel, 2.5 we restrict the sample to agents whom we observe for at least two years with their new job (This time, the initial wage cut is included in the sum). We again find no evidence, that agents that accepted an initial wage cut are compensated by steeper wage profiles on the new job.[52] Hagedorn and Manovskii (2010) provide further evidence against the mechanism. They show that wage growth of job stayers in the US is uncorrelated to local labor market tightness whereas the model by Postel-Vinay and Robin (2002) would predict it to be an increasing function of the probability to receive a job offer.[53]

---

[52]It is of course possible that the higher expected wage increases lie further in the future than the two years we observe. Given that Dustmann and Meghir (2005) find wage-tenure profiles to be basically flat after two years, however, we find this not to be very likely.

[53]The same holds true for models that stress the importance of learning about match quality over time.

2.D Estimating the Measurement Error Process

For the identification of the amount of reallocation shocks, the identification of the wage offer distribution, the identification of innovations to individual wage potential and the amount of frictional wage dispersion, we need to identify the process of measurement error. Table 2.10 reports the results of regressing within job change in log wages (after taking out year dummies) on its lags. The regression indicates that the autocovariance of wage growth is falling at higher lags and close to zero after eleven months. Therefore, we follow Meghir and Pistaferri (2004) and postulate an $MA(q)$ process for measurement error (i.e. $r_{i,t} = \Theta(q)\iota_{i,t} = \iota_{i,t} - \sum_{j=1}^{q} \theta_j \iota_{i,t-j}$) fixing q at 12.

Table 2.10: Autocovariance Structure of Wage Growth

Lag	$Coef.$
1	-0.433
2	-0.232
3	-0.180
4	-0.362
5	-0.248
6	-0.147
7	-0.143
8	-0.171
9	-0.122
10	-0.143
11	-0.075
12	0.023
13	-0.070
14	-0.048

The table reports the coefficients from regressing within job wage growth (after controlling for time fixed effects) on its own lags.

In Section 2.6.2, we use the assumption that the measurement error process does not change upon a reported job to job transition to identify the size of job heterogeneity. One can think of two alternative assumptions, one leading to an over- and one leading to an underestimation of job heterogeneity. First consider the case where

measurement error is constant within job spells and upon each job to job movement the worker draws a new measurement error shock. In this case, the variance of measurement error and job dispersion are not separately identified. However, given the high autocovariance of within job wage growth displayed in Table 2.10, we find the assumption of constant measurement error to be at odds with the data.

Alternatively, one may assume that measurement error follow a MA process within job spells, but the process begins anew upon a job to job transition. In this case, the variance of wage growth of job stayers would be inflated relatively to job switchers and we would underestimate the variance of job heterogeneity. This theory would imply that the variance of within job wage growth is larger at later stages of the job than in the period directly after the job to job transition (when the MA process is still building up). To assess this implication more formally, we compare the variance of within job wage growth in the first ten months after a job to job transition and the later months. The respective numbers are 0.0111 and 0.0115, so virtually the same.

Finally, to simulate our model, we require estimates of $\Theta(12)$ and σ_ι. We obtain these by maximizing the sum of individual likelihoods of within job wage growth in the data. More specifically, we treat $\iota_{i,t}$ as unobserved state and obtain the individual likelihood for wage growth of individual i from the following state space representation:

$$
g_{it} = \begin{bmatrix} 1 \\ \theta_1 - 1 \\ \theta_2 - \theta_1 \\ \theta_3 - \theta_2 \\ \theta_4 - \theta_3 \\ \theta_5 - \theta_4 \\ \theta_6 - \theta_5 \\ \theta_7 - \theta_6 \\ \theta_8 - \theta_7 \\ \theta_9 - \theta_8 \\ \theta_{10} - \theta_9 \\ \theta_{11} - \theta_{10} \\ \theta_{12} - \theta_{11} \\ -\theta_{12} \end{bmatrix}' \rho_{it} + \beta_1 \epsilon_{it}
$$

$$
\rho_{it+1} = \begin{bmatrix} 0&0&0&0&0&0&0&0&0&0&0&0&0&0 \\ 1&0&0&0&0&0&0&0&0&0&0&0&0&0 \\ 0&1&0&0&0&0&0&0&0&0&0&0&0&0 \\ 0&0&1&0&0&0&0&0&0&0&0&0&0&0 \\ 0&0&0&1&0&0&0&0&0&0&0&0&0&0 \\ 0&0&0&0&1&0&0&0&0&0&0&0&0&0 \\ 0&0&0&0&0&1&0&0&0&0&0&0&0&0 \\ 0&0&0&0&0&0&1&0&0&0&0&0&0&0 \\ 0&0&0&0&0&0&0&1&0&0&0&0&0&0 \\ 0&0&0&0&0&0&0&0&1&0&0&0&0&0 \\ 0&0&0&0&0&0&0&0&0&1&0&0&0&0 \\ 0&0&0&0&0&0&0&0&0&0&1&0&0&0 \\ 0&0&0&0&0&0&0&0&0&0&0&1&0&0 \\ 0&0&0&0&0&0&0&0&0&0&0&0&1&0 \end{bmatrix} \rho_{it} + \begin{bmatrix} \iota_{it+1} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}
$$

Our calibration imposes the following moment restriction: $\sigma_\epsilon = 0.016$. Table 2.11 reports our estimation results.

Table 2.11: Estimates for Measurement Error

Parameter	Estimate
σ_ι	0.119
θ_1	0.538
θ_2	0.468
θ_3	0.376
θ_4	0.093
θ_5	0.109
θ_6	0.115
θ_7	0.070
θ_8	0.058
θ_9	0.056
θ_{10}	-0.008
θ_{11}	0.021
θ_{12}	0.093

The table shows the estimation results for the measurement error process.

2.E Robustness Exercises

This section performs three robustness exercises for our quantitative model. We show that our results are almost unaffected by calibrating to a higher replacement rate. Second, we show that our main conclusions are unchanged when reducing the share of reallocation shocks by half. Third, we show that our modeling of reallocation shocks is not in contrast to the large average wage gains of young workers reported by Topel and Ward (1992).

HKV show that the amount of frictional wage dispersion in the standard job ladder model depends crucially on the replacement rate in unemployment. Indeed, equation (2.2) shows that the minimum wage is an increasing function in this parameter. In our calibration, we follow Shimer (2005) and choose a total replacement rate of 0.4. However, the literature has not settled on an appropriate value yet, and substantially higher values have been suggested. Therefore, we follow Hall and Milgrom (2008) and adjust the value of leisure upwards to imply a replacement rate of 0.71. Neither the identified deep parameters, nor the sources of wage inequality over the life-cycle change economically significant as a result from this experiment.[54] The reason for

---

[54]The Mm-ratio changes by less than 10^{-3} at each wage quintile.

the robustness is that the ability to search on the job, the process of individual wage potential and the option value generated by stochastic innovations to wage potential make workers accept almost all wage offers in our baseline calibration. Put differently, rejected offers have a very low probability of realization. In consequence, shifting the threshold somewhat to the right of the distribution still leads to very low probabilities of jobs being declined. Therefore, the equilibrium distributions are almost unaffected and the inferred parameters are almost unchanged. In fact, we observe substantial changes only with replacement rates close to one.

Given its crucial quantitative importance, we also perform a robustness exercise for the reallocation shocks. We drop the share of loss making job to job transitions as calibration target and reduce the share of reallocation shocks from 10 to only 5 percent. As expected, this calibration performs worse in matching moments of wage growth reported in Table 2.7. Yet, our main conclusions are quite robust to this exercise. The share of model implied wage inequality attributed to the search friction rises by 14 percentage points, and the share explained by initial worker heterogeneity drops by 3 percentage points. Also, the variance of the job offer distribution and innovations to individual wage potential are almost unchanged. The finding reflects our result from Section 2.3 that these types of models are mostly affected at the margin of introducing reallocation shocks.

Table 2.12: Comparing Wage Growth in the Model to Topel and Ward (1992)

Market Experience (years)	Average wage gain at job transition (model)	Average wage gain at job transition (TW)
$0 - 2.5$	0.132	0.145
$2.5 - 5$	0.099	0.099
$5 - 7.5$	0.079	0.064
$7.5 - 10$	0.068	0.046
$0 - 10$	0.102	0.094

The table compares for different ranges of labor market experience the average change in log wages at a job to job transition in our model to those reported in Topel and Ward (1992) (TW).

Topel and Ward (1992) find that early in their careers, workers experience large wage gains resulting from job to job movements. One may conjecture that this finding is in contrast to our modeling of reallocation shocks. However, our model lines up closely with their data regarding the profile of wage gains at job to job transitions as a function of experience. Table 2.12 shows that we match both the

average wage increase during the first ten years of labor market experience as well as its decreasing profile with labor market experience.[55] To understand this fact, one should keep in mind that in our model inexperienced workers are located at rather low paying jobs, implying that reallocation shocks are more likely to yield wage gains than for older workers at better jobs.

2.F Numerical Algorithm

The numerical algorithm consists of two nested loops followed by simulations. Codes are available on the authors' webpages.

- We start the algorithm by guessing functions for $b(A_t)$ and $Z(A_t)$.

- Next, we discretize the workers' log wage potential by 1500 grid points. We find 15 to be a non-binding upper bound. The distribution of the log job component is discretized into 100 equi-likely grid points.

- Given the initial guesses, we can start the inner loop, which calculates the value functions using value function iteration. Expectations regarding next period's idiosyncratic wage potential are calculated using Gaussian quadrature with 10 nodes for evaluating the innovations and linear interpolation[56] between grid points.

- The value functions of the workers allow us to to obtain policy rules for match formation. Using these, we compute the stationary distribution of the economy by distribution function iteration. For the distribution function we use a finer grid for workers' wage potential of 5000 grid points. We then update policy functions.

- Next, we update $b(A_t)$ and $Z(A_t)$ and iterate until convergence.

---

[55]The model also does a good job in matching potential labor market experience into actual labor market experience. The model performs less well in matching the large amount of job holdings after the first five years. One should keep in mind that both their data set and sample selection is different from ours, and their sample includes workers that are still in education.

[56]We opt for linear interpolation at this step, as it considerably decreases the computational burden and does not appear to alter the results compared to spline interpolation. Also, spline extrapolation is known to be unreliable.

– The last step are the simulations, that employ the policy functions and equilibrium job offer rates. We use linear inter and extrapolation on the worker and job grid.

Table 2.13: Wage Cuts after Job to Job Transitions

Stratification		Share loss	Mean loss
Whole sample			
	Nominal	0.344	-0.196
	Real	0.471	-0.146
Job characteristics			
	Union	0.345	-0.196
	+ Health insurance	0.352	-0.196
	+ Education	0.351	-0.196
Year			
	1993	0.346	-0.196
	1994	0.36	-0.216
	1995	0.364	-0.211
	1996	0.33	-0.192
	1997	0.343	-0.187
	1998	0.324	-0.196
	1999	0.345	-0.178
Gender			
	Male	0.352	-0.202
	Female	0.334	-0.189
Age			
	23-34	0.348	-0.198
	35-43	0.336	-0.196
	44-55	0.345	-0.192
Income			
	Lowest 25%	0.233	-0.16
	25-75%	0.353	-0.198
	Top 25%	0.454	-0.214
Tenure			
	Less than 6 months	0.357	-0.209
	6-12 months	0.309	-0.185
	1-3 years	0.344	-0.182
	3-10 years	0.35	-0.173
	10 and more years	0.343	-0.161
Paid by the hour		0.232	-0.078

The Table shows the share of workers incurring a cut in hourly wages after a job to job movement for the whole population and different subsamples in the 1993/1996 SIPP. Mean loss reports the mean wage loss in log points conditional on suffering a wage cut upon movement. All figures refer to nominal wages, except in the row labeled *Real*.

Chapter 3

Foreign Customer Accumulation and Export Dynamics

3.1 Introduction

New plants, after setting up production and entering a market, typically lag behind their industry competitors in terms of sales for a number of years. This holds true even for highly commoditized products where entrants and incumbents produce very similar products. An intuitive explanation for this phenomenon is that new producers are simply less efficient than their experienced competition and take a long time to catch up in terms of process and organizational know-how. A recent line of literature questions this supply side explanation and presents evidence that hints at demand side forces as determinants of the fate of young plants. Foster et al. (2012) use price information from the Census of Manufactures to show that new plants actually possess a small advantage in physical productivity compared to incumbent plants. Instead, their lower sales volumes seem to be the result of an insufficient number of customers to sell to.

If indeed new customers are so hard to come by, this micro friction has potentially important macroeconomic implications. Entrants are by no means the only plants facing the problem of finding new customers to sell to. Any plant, after a positive productivity shock would have to invest time and resources into building marketing and distribution capacities to exploit productive potential. In international economics, the presence of such market expansion friction might provide an explanation for what has been termed the *elasticity puzzle* - the discrepancy between high estimated elasticities of substitution between goods produced in different countries from trade liberalization episodes and the low elasticities needed to reproduce the co-movement

of exports and real exchange rate at business cycle frequency. Persistent tariff reductions should induce exporters to larger investments into their export demand leading to larger trade reactions than transitory exchange rate variations.

This chapter explores the macroeconomic implications of slow and active demand accumulation within the context of a dynamic model of plant exporting behavior. It introduces the notion of "customer capital" into a set-up in which plants differ in terms of revenue productivity and exporting is subject to sunk entry and fixed costs that has become the workhorse for empirical studies of export participation. I structurally estimate the model on a new panel data set of German manufacturing firms between 1995 and 2008. The substantial expansion in exporting that German manufacturing experienced during this time serves as an important case study on which to test the model's empirical predictions. The estimated model implies that during the observed time period the average plant spends between 26 and 38 percent of export revenue on building and maintaining a customer stock in export markets. The estimated demand elasticity in the export market of 1.42 is well within the range typically calibrated in international business cycle models. The model predicts a low elasticity of aggregate exports with respect to real exchange rate movements. The predicted much larger trade gains after a tariff elimination are in the same order of magnitude as those predicted by a more standard fixed costs model.

The model setup builds on the large literature of estimated dynamic models of export participation. The paper most closely related is Das et al. (2007). To their model I introduce the notion of a consumer base in export: Firms have to accumulate consumer-capital in order to sell in the foreign market. I assume that there are DRS in generating demand from investing in a consumer base. This technology may be broadly interpreted as follows: tt encompasses advertising expenditure on building brand reputation as well as the establishment of a network of local buyers and distribution channels. The unifying feature of these activities is that they take time and money to complete and additional benefits become increasingly more expensive. Customer capital, like physical capital, depreciates over time and is subject to adjustment costs. Exports are subject to an ad valorem tariff which may vary over time and exporters have to form expectations over aggregate export demand and the real exchange rate which vary stochastically.

The paper uses the AFiD Panel of Industrial Establishments, a plant level panel

maintained by the German national statistical agency to estimate the model parameters. With more than 50,000 plant observations per year, it provides extensive coverage of the German manufacturing sector and spans the years 1995 to 2008. During the observed time period, the sector experiences a strong expansion in export activity. Export participation by plants in the sample rises from 54 to 65 percent and total export revenue doubles in real terms. This expansion was the result of a drop in worldwide tariffs after the conclusion of the Uruguay round in 1995, a strong expansion in demand especially in transition economies, and favorable exchange rate movements in the initial years of the sample.

I use the data to structurally estimate the model parameters using a simulated method of moments (SMM) estimator. After solving the plant problem by value function iteration, I simulate the export behavior of plants to obtain a panel data set of the same size as the underlying data set. During the simulation, I feed in the observed time series for aggregate tariffs, export demand and real exchange rate. I obtain estimates by minimizing a quadratic form criterion function in a vector of data moments.

This chapter is the first to obtain an estimate from plant level data of the marketing costs plants incur to maintain and increase their customer stock in export markets. The estimated costs are large and account by far for the largest share of export costs. In 1995, the beginning of the sample for example, the average exporter spent 2.4 million euros on maintaining customer capital and a further 0.96 million on expanding it. Estimated entry costs into exporting of 33,467 1995 euros are relatively small when compared to other estimates in the literature.

The estimation procedure succeeds in matching the chosen target moments well, in particular when comparing growth rates and survival probabilities for new exporters. Its predictions on export participation are well in line with the data up to 2006 after which it misses a further surge in export participation. It closely matches the shape of the growth in total export revenue while overpredicting its absolute size. Importantly, when compared to an estimated restricted model version in which plants face fixed costs of exporting only, the baseline correctly predicts a further surge in exporting revenue after 2003 while the fixed cost model predicts exports to slightly decline.

I use three different scenarios to compare the dynamic implications of the baseline model to a more standard sunk/fixed cost model of exporting: a real exchange rate

depreciation, an increase in worldwide demand and an elimination of all export tariffs. Aggregate exports in the fixed cost model react almost twice as strong to a real exchange rate depreciation compared to the in baseline model with estimated elasticities of 3.4 compared to 1.9. Meanwhile, both models predict gains of total exports of 10 to 11 percent when moving to free trade. In case of the baseline model, however, these take almost ten years to fully realize. So, unlike the fixed cost model, the estimated baseline model can reconcile large predictions of trade gains after a tariff reduction with a subdued reaction of aggregate exports to a real exchange rate depreciation. Slow and active demand accumulation therefore quantitatively manages to provide an answer to the elasticity puzzle.

This chapter builds on a long line of literature in empirical trade that estimates structural and reduced form discrete choice models of export participation on plant or firm level data. Some of the most prominent examples are Roberts and Tybout (1997), Bernard and Wagner (2001), Bernard and Jensen (2004) and the aforementioned study by Das et al. (2007). Willis and Ruhl (2009) and Arkolakis (2010) point out that the predictions of these models are at odds with the observed increasing survival rates for new exporters and the above average export growth rates among small exporters following trade liberalizations.

A number of recent contributions has therefore started to extend this framework by broadening the focus beyond mere participation. Fitzgerald and Haller (2012) use plant level information on export destinations and sales by six digit tariff line from Irish manufacturing firms to show that lagged export sales are an important predictor of future trade participation. They also interpret this as evidence of some market specific demand factor that plants have to accumulate over time. This chapter differs in its use of a structural estimation approach which allows me to quantify the costs associated with foreign customer accumulation and to conduct counterfactual experiments. Eaton et al. (2010) combine Columbian plant level data with U.S. Customs data to establish patterns in individual sales relationships between the Columbian plants and their U.S. buyers. They then propose a continuous time search model with heterogeneous buyers. The evidence they provide is mostly consistent with the model and data of this chapter. One may think of the search friction as micro foundation for the demand technology I assume. Arkolakis (2010) also proposes a static model of costly costumer accumulation to reconcile theory and data.

The chapter proceeds as follows. Section 3.2 presents the model. In Section 3.3, I introduce my data set, discuss sample selection, and establish some stylized facts which indicate a slow process of demand accumulation for new exporters. I also report how I calculate series for aggregate tariffs, export demand and the real exchange rate. Section 3.5 presents parameter estimates, evaluates model fit and discusses what implications they have for the cost of exporting. Section 3.6 contrasts the dynamic behavior of the model to that of a more standard fixed cost model. Section 3.7 concludes.

3.2 A model of export participation and intensity choice

The section presents a dynamic model of plant export participation and costly customer accumulation. Like in much of previous literature, exporting is subject to sunk entry and fixed participation costs. I extend the framework by the notion of *customer capital* in the export market. After entering into exporting, plants first have to spend resources on acquiring a stock of customers to sell to. This includes advertising expenditures on building brand reputation as well as the establishment of a network of local buyers and distribution channels. In the model, this concept takes the form of a capital good that the firm has to spend resources on in acquiring and which generates demand through a decreasing returns to scale technology[1]. Like physical capital, it depreciates geometrically and is subject to convex adjustment costs in the size of gross investment. The latter imply that plants take several years after entering the market to reach their desired market size. They also induce a decreasing exit hazard over time as plants have more customers to sell to. When a firm chooses not to export in any year, it starts its next exporting episode having to build up its entire customer base anew. The sunk component in exporting costs therefore increases as firms grow bigger.

In the part of the export market that a plant has acquired access to, it behaves as monopolistic competitor which rules out strategic considerations in pricing. Other than on its price, a plant's profits from exporting depend on idiosyncratic shocks to export demand and production costs and aggregate movements in real exchange rate, tariffs, and total income in export markets. Idiosyncratic profitability, real exchange

---

[1]Given my assumption of linear costs in accumulating customers, decreasing returns imply an interior solution for the target market size.

rate and aggregate demand follow known Markov processes over which the plant has to form expectations. The model focuses on the export market to keep the already highly dimensional model tractable and focused. The abstraction from the home market implies that home and foreign market are independent in terms of consumer base.

3.2.1 Export Revenues and Profits

When plant i chooses to export in a given period t, it faces the following demand schedule in its export market

$$q_{it}^D = \epsilon_{it} p_{it}^{*-\eta} D_{it}^\alpha D_t^W. \tag{3.1}$$

ϵ_{it} is a shock that shifts the idiosyncratic demand schedule. p_{it}^* is the price in foreign currency terms that the plant sets for its product . D_{it} is the size of the customer base that the firm has accumulated in foreign markets. I will discuss the details of how a firm accumulates market share below. The assumption that $0 < \alpha < 1$ implies that *customer capital* is subject to decreasing returns to scale. D_t^W represents aggregate demand in export markets. It, too, evolves stochastically[2].

Export sales are subject to an ad valorem tariff τ_t, which may evolve over time. The real exchange rate in period t is denoted by RER_t. The foreign demand schedule therefore implies the following revenue function dependent on goods sold q_{it}:

$$R\left(q_{it}\right) = \frac{RER_t}{1 + \tau_t} q_{it}^{\frac{\eta-1}{\eta}} \epsilon_{it}^{\frac{1}{\eta}} \left(D_{it}^\alpha D_t^W\right)^{\frac{1}{\eta}}.$$

Plants face variable costs of production c_{it} such that gross profits from exporting are given by:

$$\pi_{it} = R\left(q_{it}\right) - c_{it}q_{it}.$$

Profit maximization implies that plants choose prices as a fixed mark-up over variable production costs: $p_{it}^* = \frac{\eta}{\eta-1} \frac{1+\tau}{RER_t} c_{it}$. In my data I observe a plant's export revenue denoted in terms of domestic currency but no information on production costs and

---

[2]Equation (3.1) would follow naturally from the assumption of a CES consumption aggregator in the export markets. In that case, D_t^W would be total consumption and D_{it}^α would measure the intensity, with which plant i has penetrated the market, i.e. the amount of customers it can reach.

therefore no direct measure of gross profits from exporting. However, given the previously made assumptions, optimal pricing implies that profits are a fixed fraction of revenues: $\pi_{it} = \frac{1}{\eta} R_{it}$. Per period gross profits from exporting evaluated using optimal prices are therefore given by:

$$\pi_{it} = \frac{1}{\eta} \left(\frac{\eta - 1}{\eta} \right)^{\eta - 1} \epsilon_{it} c_{it}^{1-\eta} D_{it}^{\alpha} D_t^W \left(\frac{RER_t}{1 + \tau_t} \right)^{\eta}.$$

I normalize constants to one and define $z_{it} = \epsilon_{it} c_{it}^{1-\eta}$ as a composite state for idiosyncratic export profitability. For the estimations, this allows me to work with the following relatively simple equation for potential export profits:

$$\pi_{it} = z_{it} D_{it}^{\alpha} D_t^W \left(\frac{RER_t}{1 + \tau_t} \right)^{\eta}. \tag{3.2}$$

It is also z_{it} for which I make distributional assumptions and over which the firm forms expectations. More specifically, I assume idiosyncratic export profitability to be the sum of a fixed component χ_i and a persistent component ϕ_{it}:

$$z_{it} = \phi_{it} + \chi_i.$$

The permanent component χ_i is a realization of a log-normal distribution $\ln N \left(0, \sigma_{\chi}^2 \right)$. The logarithm of ϕ_{it} follows an AR(1) process with innovations ε_{it}^{ϕ} drawn from a normal distribution with zero mean and variance σ_{ϕ}^2:

$$\ln (\phi_{it}) = \rho_{\phi} \ln (\phi_{it-1}) + \varepsilon_{it}^{\phi}.$$

3.2.2 Costs of Exporting

A firm that wants to export in any given period faces two types of costs: a stochastic fixed overhead and costs of maintaining or increasing its foreign customer stock. The fixed cost is on average higher, when the firm has not been an exporter in the previous period and therefore has a sunk component. It represents administrative costs and costs of complying with foreign regulations and customs procedures. Those are unlikely related to the amount of exporting a firm does. In addition, penetrating the export market means having to acquire customers via marketing and potentially building up a network for distribution and sales[3]. I treat foreign demand potential D_{it} as a capital good like in Arkolakis (2010), Drozd and Nosal (2012), and Fitzgerald and Haller (2012). It depreciates over time and its accumulation is costly and subject to convex adjustment costs. Convex adjustment costs mean that a firm, after entering into exporting, will take a number of years to build up its desired stock of foreign customers.

Costs of accumulating a customer base

A plant that did not export previously enters the year without any customers in foreign markets. A plant that served D_{it-1} customers last year, retains a fraction of $(1 - \delta)$. Other supply relations dissolve, because partners in foreign markets go out of business, and built up reputation from previous marketing campaigns becomes less valuable. Before determining its export volume for the current year, the plant has the opportunity of investing into its customer stock. Accumulating additional units of $I_{it} = D_{it} - (1 - \delta)D_{it-1}$ costs $c^{lin}I_{it}$ in terms of current profits. Additionally, it has to pay quadratic adjustment costs $\frac{c^{conv}}{2}\left(\frac{I_{it}}{D_{it}}\right)^2 D_{it}$ on its investment. The total costs of investing $c\left(D_{it}, D_{it-1}\right)$ are:

$$c\left(D_{it}, D_{it-1}\right) = c^{lin}\left(D_{it} - (1 - \delta)D_{it-1}\right) + \frac{c^{conv}}{2}\left(\frac{I_{it}}{D_{it}}\right)^2 D_{it}.$$

---

[3]Alternatively, the firm can outsource marketing and distribution to local subcontractors. In this case, it probably has to find different subcontractors for different regions such that the costs of distribution network are still increasing in its size.

Fixed Costs of Exporting

A firm that did not export in the previous year needs to pay a sunk entry cost $\gamma_E - \xi_{it}^E$ where $\xi_{it}^E \sim N\left(0, \sigma_E^2\right)$. If it has exporting experience from the last year, it pays a fixed cost of $\gamma_F - \xi_{it}^F$, where $\gamma_F < \gamma_E$ and $\xi_{it}^F \sim N\left(0, \sigma_F^2\right)$. Letting $y_{it} \in \{0, 1\}$ denote the export state of plant i in period t, per period net profits from exporting can be summarized by:

$$u\left(\phi_{it}, \chi_i, D_{it}, D_{it-1}, RER_t, D_t^W, \tau_t\right) = \begin{cases} \pi_{it} - \gamma_F + \xi_{it}^F - c(D_{it}, D_{it-1}) & y_t = 1 \wedge y_{t-1} = 1 \\ \pi_{it} - \gamma_E + \xi_{it}^E - c(D_{it}, 0) & y_t = 1 \wedge y_{t-1} = 0 \\ 0 & y_t = 0 \end{cases}$$

3.2.3 Bellman Equations

In any given year t, a plant observes the current realizations of $\xi_{it}^X, \phi_{it}, RER_t, D_t^W$. It then decides whether to participate in the exporting business this period. If it does, it also decides how much to invest into its customer stock and sets prices thereafter. Prior to 1995, plants expect the current tariff level to persist forever. In 1995, they learn the whole tariff sequence up to 2008. They assume that from 2008 on, tariffs will stay at that level forever. Dropping time subscripts for all variables except τ and denoting future values of a variable x by x', I summarize the dynamic problem for the firm using two Bellman equations:

– *For plants that exported in the previous period*

$$V^1(\xi^F, \phi, D_{-1}, RER, D^W, \tau_t) = \max\Big[\max_D\big\{\pi - \gamma_F + \xi^F - c(D, D_{-1}) \\ + \beta\mathbb{E}\left[V^1\left(\xi^{F\prime}, \phi', (1-\delta)D, RER', D^{W\prime}, \tau_{t+1}\right)\right]\big\}, \\ \beta\mathbb{E}\left[V^0\left(\{\xi^{E\prime}, \phi', RER', D^{W\prime}, \tau_{t+1}\right)\right]\Big]$$

(3.3)

– *For plants that did not export in the previous period*

$$V^0(\xi^E, \phi, RER, D^W, \tau_t) = \max\Big[\max_D\big\{\pi - \gamma_E + \xi^E - c(D, 0) \\ + \beta\mathbb{E}\left[V^1\left(\xi^{F\prime}, \phi', (1-\delta)D, RER', D^{W\prime}, \tau_{t+1}\right)\right]\big\} \\ \beta\mathbb{E}\left[V^0\left(\xi^{E\prime}, \phi', RER', D^{W\prime}, \tau_{t+1}\right)\right]\Big]$$

(3.4)

3.3 Data

This section introduces the data set, the *AFiD Panel of Industrial Establishments* maintained by the German Federal Statistical Office. It contains on average more than 50,000 establishments per year and covers the years 1995 to 2008. During the observed time period, German manufacturing experienced an exporting boom. Export participation in the sample rose from 54 to 65 percent and the (unweighted) average revenue share of exporting increased from 22 to 28 percent. As a result, aggregate real export revenue doubled.

After briefly describing sample selection, in Section 3.3.3 I provide further evidence from the data set that plants after entering the export market face important demand side frictions. In the first four years of an exporting spell, entrants have on average higher growth rates than incumbents. While the year of entry is associated with a healthy increase in domestic revenues, pointing to an initial increase in productivity, it is not sustained in the following years. This is evidence against upgrading in physical productivity as result of exporting. Also, survival probabilities are an increasing function of exporting tenure.

In Section 3.3.4 I discuss how I calculate tariffs. The year 1995, the beginning of the data set, saw the conclusion of the Uruguay round of trade negotiations which resulted in the creation of the WTO and a commitment to substantial tariff reductions for the trade in manufacturing goods from joining nations. Those were to be gradually phased in until 2000 for developed and until 2004 for developing countries. Also, in 2004 ten mostly Eastern European countries joined the EU which meant that German firms could now access those neighboring markets tariff free. As a result, the trade weighted average value added tariff for German manufacturing exports dropped from about 3.3 percent in 1995 to about 2.1 percent in 2008.

Simultaneously, real per capita income in many emerging economies started growing strongly. Market size in those countries relative to the German home market increased by more than thirty percent. Section 3.3.5 discusses how I calculate a time series for real aggregate export demand. Finally, in Section 3.3.6 I present the real exchange rate series. Germany's effective real exchange rate in 1995 had just strongly appreciated after the calamities in the European Exchange Rate Mechanism that forced out Britain and caused Italy to depreciate against the German Mark. In the years up to 2000, Germany devaluated by more than 20 percent against its trade

partners. Until the end of the sample, it appreciated again by about 10 percent.

3.3.1 AFiD - Administrative Firm Data for Germany

My data set is the *AFiD Panel of Industrial Establishments.* It is an annual, administrative plant-level panel maintained by the German Federal Statistical Office. It samples from the universe of German manufacturing establishments with 20 or more employees. In sectors with predominantly smaller firms, this cut-off can be substantially lowered. Participation for the sampled plants is mandated by law. A plant is counted as an individual unit if it is locally separated from other establishments belonging to the same firm. Ownership is recorded, but does not influence sampling. Establishments owned by German firms in other countries are not included.

The sample covers the years 1995 to 2008. Variables collected include total revenue, total export revenue, employment, hours worked and investment. One of the appeals of the data set is its coverage. It comprises an average of about 50,000 plants per year. The panel is unbalanced but plants tend to stay in the sample quite long such that an uninterrupted series of 14 observations exists for 26,522 plants and about 10,000 more have at least 10 years worth of observations. Sectorial classifiers allow to group plants into *NACE rev. 1.1* sectors.[4] All NACE sectors related to manufacturing and the extractive industries are covered (NACE 10-36). While there are clear sectoral differences with regards to export participation and intensity, they share a common upwards trend which is why I opt not to do estimations separately by sector.[5]

3.3.2 Sample selection and summary statistics

I first delete every observation from the data set that has missing information on employment, total revenue and export revenue. In order to eliminate plant-year observations which are potentially the result of misreporting, I also delete for every year separately observations which are in the top and bottom percentile for employ-

---

[4]For further information on the methodology behind the NACE classification and its relation to other systems of industry classification, you may consult Eurostat.

[5]The interested reader finds more information on the data set in Appendix 3.A. Tables 3.6 - 3.7 reports sample splits by sectoral classification, firm type and employment size categories. All of them use the year 1999 as an example.

Figure 3.1: The extensive and the intensive margin of exporting in AFiD
(1995-2008)

(a) Export participation   (b) Growth in aggregate export revenue
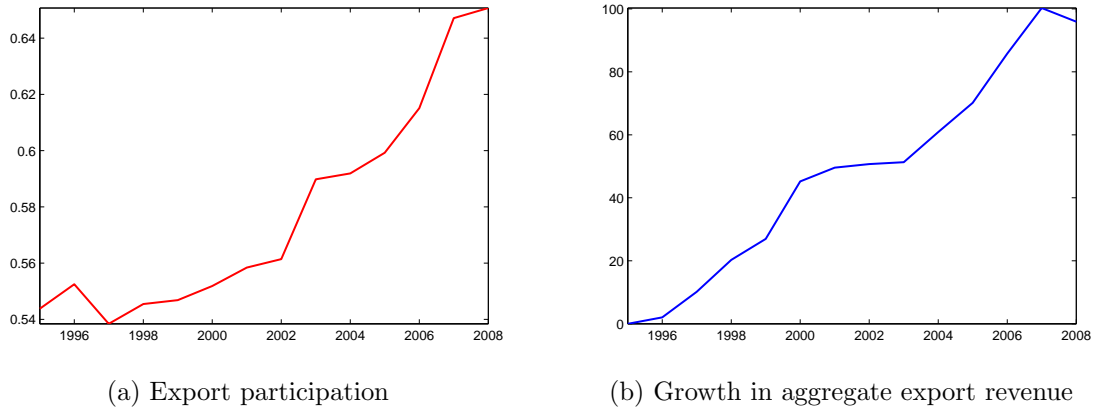
Note: The figure illustrates trends in the exporting behavior of the plants in the data set during the observed time period. Panel (a) shows the percentage of plants with positive export revenues. Panel (b) reports the growth in real aggregate export revenue.

ment or total revenue growth[6]. The resulting sample contains 667,601 plant-year observations from 81,913 different establishments.

Historically always export oriented, the German manufacturing sector grew even substantially more so during the time of our sample. As Figure 3.2a shows, the percentage of firms in the sample engaged in some form of exporting activity grew by more than 10 percentage points from about 54 percent 1995 to about 65 percent in 2008. Conditional on being an exporter, export revenue grew also more important as a share of total revenue. The average export revenue share rose from less than 22 percent in 1995 to more than 28 percent in 2008. Figure 3.2b shows that as a consequence of this expansion on the extensive and intensive margin, real total exports doubled between 1995 and 2007.

As illustrated in Figure 3.7 in the appendix, the expansion in export participation was the result of high entry rates and a steady decline in exit rates from exporting. Nonetheless, even though export participation as a whole increased substantially during the sample period, turnover between exporting and non-exporting remained high. Even in the boom year of 2006, more than two percent of plants with positive

---

[6]This criterion is not applicable for export revenue growth. Because of the way I measure growth rates, the top and bottom percentile of export revenue growth rates would largely contain episodes of exit and entry

export revenue in 2005 did not report any exporting in 2006.

3.3.3 Evidence for slow demand accumulation

In a recent contribution, Foster et al. (2012), using data from the US Census of Manufacturers, show that entrants into the US domestic market are much smaller than their established industry competitors, and that they may take over a decade to close the gap in sales. Using price data, they can show that these size differences are not the result of lower productivity but rather of a lack customers to sell their products to. Established firms entering into exporting should face a similar problem. They, too, first have to establish distribution channels, explore and penetrate markets and build a reputation with customers. These are both time- and resource consuming activities that imply that it takes young exporters a number of years to reach their desired export volumes. These entry episode therefore serve as another good test case on which to evaluate the hypothesis of demand side impediments to plant growth.

Table 3.1: Sales growth rates for new exporters

Year after entry	Sales growth	
	Export	Domestic
0	2.000	0.093
1	0.288	-0.011
2	0.060	-0.017
3	0.059	-0.016
4	0.040	-0.015

Note: The table displays growth rates for plants who enter into exporting for export revenue (column 1) and domestic revenue (column 2) conditional on surviving as exporter. I count the year when the plant enters into exporting as year zero. Growth rates are calculated as $\frac{x_t - x_{t-1}}{.5(x_t + x_{t-1})}$. This statistic is bounded between -2 (exit) and 2 (entry).

As documented in the previous subsection, due to high turnover and a substantial net expansion in export participation, entrance into exporting is a very frequent occurrence in my data set. In total, there are 14,814 observations where a plant shows a switch from non-exporter to exporter. I exploit this fact to test hypothesis on sales behavior during the years after entry. Table 3.1 reports mean sales growth

rates for entrants into exporting in foreign and domestic markets. Throughout the chapter, I calculate growth rates as $(x_t - x_{t-1})/(.5(x_t + x_{t-1}))$. This measure, first introduced by Davis et al. (1998), has the advantage of being bounded between -2 and +2, which allows to include market entry and exit.

There are two observations one can make from the table. For one, conditional on surviving as an exporter, in the first four years after entering into exporting, new exporters show on average higher export sales growth rates than incumbent exporters. The average export sales growth rate over all plants and years is 0.051. Second, while entering into exporting is also associated with a healthy increase in domestic sales which points to a positive productivity innovation, this cannot be said for the subsequent years which all display negative growth rates (which are consistent with mean reversion in productivity). In combination, these two observations suggest that while entering into exporting is associated with a productivity increase, subsequent growth in export sales is not the result of further improvements in productivity. Instead, improved demand conditions seem to be responsible for the strong export sales growth in the year following entry.

Table 3.2: Survival probabilities for new exporters

Year after entry	Probability of survival
0	67.9
1	82.6
2	87.3
3	89.7
4	92.0
5	91.7
6	93.1
7	94.4
8	95.5

Note: The table displays survival probabilities as exporter in the years after entering into exporting. For example, the first line shows that for firms who report export revenue after having had not exported in the first year, the chance of again exporting in the next year is 67.9 percent.

Table 3.2 illustrates another fact that hints at the accumulation of some export market specific factor: the probability of surviving as an exporter is an increasing

function in the number of years already spent exporting. While sunk fixed costs in entering exporting would explain hysteresis in the exporter status, their combination with mean reversion in revenue productivities still implies a decreasing survival probability as more plants revert below the continuation threshold.

3.3.4 Tariff data

The AFiD data set does not record plant level information regarding export destination country. The only variable recorded is total export revenue. In order to empirically identify the effect of variations in tariff duties on export participation and sales, I construct a yearly average measure of tariff duties that German manufacturing exports were subject to between 1995 and 2008. Export data by sector are available from the Eurostat Comext data base[7] and recorded according to the Harmonized System (HS) classification as is the norm with trade data. I manually construct matches between NACE (rev 1.1) sectors and HS(2) chapters and obtain dissaggregate export data for the relevant sectors. In order to construct a stable world market aggregate for the time period, I then rank export destinations by trade value for every year. The union set of the top twenty partners in every year constitutes the world market. The set stably accounts for about 80 percent of total German manufacturing exports. Please consult Appendix 3.D for more details on this procedure and for a list of the export destinations included.

I then obtain six digit tariff line data from the WTO[8] and match them with the export data from Comext to calculate an average trade-weighted measure for ad valorem export duties during the sample period. The resulting time series is displayed in Figure 3.3a. With an average of only 3.3 percent, tariff duties on German exports were already low at the beginning of the sample period. The trade liberalizations described in the introduction to this section are well visible in the series. Between 1995 and 2000, the tariffs drop by a full percentage point to 2.2 percent as a result of the implementation of the Urugay tariff reductions and abolitions. In 2004, the accession of ten mostly Eastern European countries to the European Union means exports to those destinations are now tariff-free which has aggregate tariffs drop by another .3 percentage points.

---

[7]For more details see `http://epp.eurostat.ec.europa.eu/newxtweb/`
[8]`http://www.wto.org/english/tratop_e/tariffs_e/tariff_data_e.htm`

3.3.5 Aggregate export demand

In my model, there are four reasons why export revenue for a plant grows. The plant may expand its penetration of world markets by investing into its stock of customers. It may experience an increase in revenue productivity. The value of exports in domestic currency terms may shift because of favorable exchange rate movements. Tariff reductions increase the share of sales revenue the plant retains. Finally, aggregate import demand in the part of the world that the plant has penetrated grows with aggregate income.

I use the same set of countries from the calculation of the tariff series to obtain a world demand series. Total demand in country i at time t in real terms is given by

$$D_{i,t} = GDP_{i,t} + IM_{i,t} - EX_{i,t}$$

where $GDP_{i,t}$ is real GDP and $IM_{i,t}$ are $EX_{i,t}$ are real aggregate imports and exports. The time series for these variables are taken from the World Bank Development Report[9] and the IMF Global Economic Outlook[10] . The resulting aggregate demand series by country are then averaged and weighted by total German manufacturing exports. Finally, I make the assumption that the costs of exporting (production costs and other) are growing at the same rate as domestic demand. The relevant series for evaluating export profitability is therefore $\frac{D_t^W}{D_t}$, world demand relative to domestic demand. Normalizing relative aggregate export demand in 1995 to 1.0, the resulting time series is displayed in Figure 3.3b. Between 1995 and 2005, total export demand relative to domestic demand grows by 27 percent, a level at which it stays more or less constant until the end of the observed time period.

3.3.6 Real exchange rate

The real exchange rate data are from the Bank for International Settlements (BIS). Ideally, one would want to calculate the real exchange rate for the same set of countries used for calculating tariffs and aggregate demand series. Meanwhile, especially for the emerging market economies which are are part of the world aggregate, con-

---

[9]Available at `http://databank.worldbank.org/Data/Home.aspx`

[10]Available at `http://www.google.com/publicdata/explore?ds=k3s92bru78li6_&hl=en&dl=en`

Figure 3.2: Aggregate series

(a) Trade weighted tariff series   (b) Aggregate export demand series (1995=1)
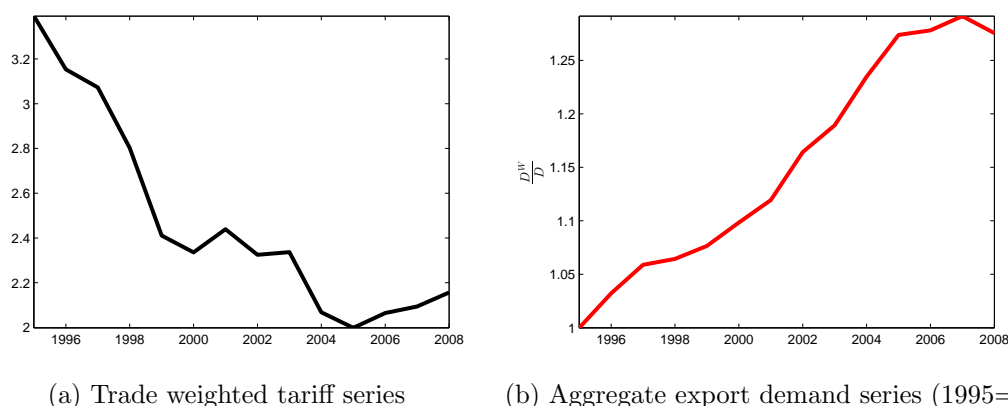
Note: Panel (a) displays the trade-weighted average tariff series for German manufacturing exports between 1995 and 2008. To calculate the series, I match average across HS(6) tariff lines for manufacturing exports and weight them by trade volume for export destination and tariff line. Panel (b) displays world demand in the export market for German manufacturing relative to home market demand. The series is normalized to 1 in 1995. The export market is an aggregate of Germany's 23 most important trade partners weighted by trade volume.

sistent inflation data for this time period is hard to obtain. I therefore use the BIS data as the closest approximation. The BIS time series is based on a trade basket of 51 partner countries going back to 1994. There also exists a longer time series going back to 1964 based on 22 trade partners. My estimations of the real exchange rate process below are based on a series where I splice the series based on the more narrow basket until 1993 to the series based on the broader basket. In the case of Germany, until about 1997, these two series almost coincide so there should be no problem of continuity.

The real exchange rate is calculated as the geometric weighted average of bilateral nominal exchange rates adjusted with the corresponding relative consumer prices. The trade based weighting methodology has its theoretical underpinnings in Armington (1969). The weights capture both direct bilateral trade and third market competition by double-weighting.[11] Please consult Appendix 3.D for a display of Germany's real exchange rate during the sample period.

---

[11]For more details on the methodology, please consult Klau and Fung (2006).

3.4 Estimation

This section describes how I estimate the model presented in Section 3.2 on the plant level data from the previous section. In a first stage, I make assumptions on the processes for the real exchange rate (RER_t) and world demand (D_t^W) whose parameters I estimate independently. This leaves $\Omega = \left\{ \beta, \eta, \alpha, \delta, \rho_\phi, \sigma_\phi, \sigma_\chi, c^{lin}, c^{conv}, \gamma_E, \sigma_E, \gamma_F, \sigma_F \right\}$, a set of 13 parameters to be determined. I set β to an annual interest rate of 5 percent and estimate the other parameters using a Simulated Method of Moments (SMM) approach as developed by McFadden (1989), Lee and Ingram (1991), and Duffie and Singleton (1993). The twelve estimated parameters are collected in the parameter vector θ. Further information on how I solve for the establishment policy functions used in the simulation and on SMM itself can be found in Appendix 3.E.

3.4.1 Parameters estimated outside of the Model

I obtain parameter estimates for the real exchange rate process and aggregate export demand prior to estimating the remaining model parameters. Table 3.3 summarizes the estimated parameters and gives standard errors.

World demand process parameters

In my model simulation below, I assume D_t^W to follow a bounded random walk of the form

$$\log \left(D_t^W \right) = \max \left[\min \left\{ \log \left(D_{t-1}^W \right) + \sigma_{DW} \epsilon_t^{DW}, \log(D_{max}^W) \right\}, \log \left(D_{min}^W \right) \right].$$

Section 3.3.5 explained how I calculate a times series for aggregate export demand for German manufacturing products between 1995 and 2008. I estimate σ_{DW} from the standard deviation of the growth rate of that series assuming that the random walk is not at its bounds. The boundedness assumption helps in assuring stationarity of the dynamic programming problem. After normalizing D^W to 1.0 in 1995, I pick $D_{min}^W = 0.75$ and $D_{max}^W = 1.5$ as lower and upper bounds. These turned out large enough to ensure that further expanding them did not affect estimation results[12].

---

[12]It should be noted that discounting by the plants and the fact that I am simulating the model over a finite number of periods mean that my results are unlikely to differ much from a model where D^W follows an unbounded random walk.

Real exchange rate parameters

I assume the annual real exchange rate to follow an AR(1) process in logs:

$$\log(RER_t) = \rho_{RER} \log(RER_{t-1}) + \sigma_{RER}\epsilon_t^{RER} \quad , \; \epsilon_t^{RER} \sim N(0,1)$$

The p-value of an Augmented Dickey-Fuller Test for non-stationarity on the real exchange rate series from 1973 (the end of Bretton Woods) to 2008 is 0.0139, so non-stationarity is rejected at the 5, but not the 1 percent confidence level. The point estimate of ρ_{RER} is 0.9073 with a standard deviation of 0.0968. In order to make my analysis comparable to that in previous studies, I stick to the assumption of a stationary exchange rate.[13] Table 3.3 summarizes the estimates for the processes for aggregate demand and real exchange rate. I obtain the standard errors for the innovation variances from bootstraps with 2,000 repetitions.

Table 3.3: Estimated parameters for the aggregate stochastic processes

Parameter	σ_{DW}	ρ_{RER}	σ_{RER}
Estimate	0.0149	0.9073	0.0366
	(0.0025)	(0.0968)	(0.004)

Note: The table reports the estimated standard deviation of innovations to aggregate export demand (σ_{DW}), the persistence of the real exchange rate (ρ_{RER}) and the standard deviation of real exchange rate innovations (σ_{RER}). Standard errors are in parenthesis.

3.4.2 A Simulated Method of Moments Approach

A direct evaluation of the model's likelihood function in the parameter vector θ is infeasible. I therefore use a Simulated Method of Moments procedure. Its underlying idea is that, under the null hypothesis of a parameter estimate $\hat{\theta}$ being the true parameter vector θ_0, simulated panels of the same number of units N and of the same length T will on average yield the moments $\mu(\hat{\theta})$ observed in data set. Variations around the mean are the result of simulation uncertainty for the simulated data and sampling uncertainty for the observed moments. Simulations are therefore repeated S

---

[13]As also discussed in Das et al. (2007), studies of real exchange rates dynamics fail to reject a random walk because of limited test power, while studies that exploit long time series or pool countries like Frankel and Rose (1995) are often able to do so.

times over a fixed set of different stochastic draws and $\mu(\hat{\theta})$ is estimated by averaging over these draws. The challenge then is to identify a vector of moments $\mu(\theta_0)$ which is informative about the underyling parameters in the sense that $E_0 \frac{\partial \mu}{\partial \theta} \gg 0$ for as many entries as possible, i.e. the moments are responsive to changes in underlying parameters.

I obtain parameter estimates by minimizing the quadratic form criterion function

$$\left[\mu(\theta_0) - \frac{1}{S} \sum_{s=1}^{S} \mu(\hat{\theta})_s \right] W^{*-1} \left[\mu(\theta_0) - \frac{1}{S} \sum_{s=1}^{S} \mu(\hat{\theta})_s \right] \tag{3.5}$$

where W^* is the optimal weighting matrix $Var\left(\mu(\theta_0) - \frac{1}{S} \sum_{s=1}^{S} \mu(\theta)_s \right)$. As shown by Lee and Ingram (1991), under the null $Var\left(\frac{1}{S} \sum_{s=1}^{S} \mu(\theta)_s \right)$ is equal to $\frac{1}{S} Var\left(\mu(\theta_0) \right)$. Independence of simulated and data moments then implies

$$Var(\hat{\theta}) = \left(1 + \frac{1}{S} \right) \left[E_0 \frac{\partial \mu'}{\partial \theta} W^{*-1} E_0 \frac{\partial \mu}{\partial \theta} \right]^{-1} . \tag{3.6}$$

I obtain $Var\left(\mu(\theta_0) \right)$ from (block-)bootstrapping the data 1,000 times with replacement. Setting $S = 20$ means that the standard error of $\hat{\theta}$ is increased by 5 % as result of simulation uncertainty.

Given a guess for the parameter vector $\hat{\theta}$, I solve the firm problem using value function iteration. I then simulate a panel of $N = 50,000$ plants S times over a fixed set of random draws feeding in the empirically observed time series for tariffs, aggregate demand and real exchange rate. During an initial period, I first let the economy settle into its stochastic steady state in which aggregate demand and tariffs are fixed at their 1995 values. The real exchange rate varies randomly during the initial periods following the AR(1) process estimated in the previous subsection. In the 31 years leading up to the estimation window, I then fix the exchange rate to follow its observed path.[14] I also assume that, during the initial periods, the plants are unaware of the tariff changes that start taking place in 1995, the first year of the simulation window. In 1995, they learn the tariff transition path that leads to the new, low-tariff stochastic steady state. Real exchange rate and aggregate

---

[14]I have real exchange rate data for the years 1964 to 2008 such that I can use the observations for the years 1964 to 1994 during the initialization period.

export demand follow their observed paths and plants form expectations according to the estimated processes. For more details on the numerical solution method for solving the plant problem and the global optimization algorithm used in finding the minimum of the objective function you may also consult Appendix 3.E.

Table 3.4: Data moments used in estimation

Export revenue growth rate			**Growth rates entrants**			**Survival rates entrants**		
Moment	Data	Model	Period	Data	Model	Period	Data	Model
Mean	0.051	0.056	1	0.276	0.281	0	0.667	0.643
Std	0.730	0.634	2	0.062	0.127	1	0.826	0.760
Skewness	0.002	0.061	3	0.063	0.065	2	0.877	0.847
Kurtosis	5.326	6.682	4	0.041	0.028	3	0.902	0.903
Autocorr(1)	-.1958	-.151				4	0.923	0.935
Autocorr(2)	-0.040	-0.086				5	0.922	0.950
						6	0.936	0.960
						7	0.949	0.968
						8	0.960	0.967

Export revenue distribution			**Export participation**		
Quintile	Data	Model	Moment	Data	Model
2	0.006	0.002	$Corr(Exp_t, Exp_{t-2})$	0.881	0.878
3	0.019	0.008	$Corr(Exp_t, Exp_{t-4})$	0.826	0.832
4	0.060	0.043	$NE - E$	0.056	0.052
5	0.915	0.946	$E - NE$	0.030	0.029
			1995 Exporters	0.544	0.560

Note: The table reports the data moments that identify the model parameters via a Simulated Method of Moments estimation. They can be classified into five subcategories: 1) Statistics of the export revenue growth rate 2) Export revenue growth rates for entrants into exporting in the years after entry 3) Survival rates for entrants in the years after entry 4) Quintiles of the export revenue distribution 5) Moments on export participation, and transitions. Growth rates are calculated as $(x_t - x_{t-1})/(.5(x_t + x_{t-1}))$. This statistic is bounded between -2 (exit) and 2 (entry). Exp_t is an indicator variable that equals one if a plant has positive export revenue in period t and zero otherwise. Data refers to data moments. Model are simulated counterparts from the model's baseline specification.

I use a total of 28 moments from the AFiD data to estimate the model. See Table 3.4 for a summary. They can be devided into five broad categories. The first category contains moments of the export revenue growth distribution, the first four centered moments and one and two year autocorrelation. Revenue growth is again defined by $\frac{Rev_t - Rev_{t-1}}{.5(Rev_t + Rev_{t-1})}$ for it to include entry and exit. Export revenue growth rates for

entrants in the first four years after entry constitute the second subcategory.[15] The third set of moments is made up of the survival rates for entrants into exporting after the initial year and the eight subsequent years. The fourth subcategory describes the export revenue distribution and contains the shares in total export revenue (averaged over the sample period) for the second to fifth quintile.[16] Finally, the last set of moments concerns the extensive margin of exporting: persistence in exporting status over two and four years, average entry $(NE - E)$ and exit $(E - NE)$ rates into and from exporting as well as total export participation at the beginning of the sample.

3.5 Results

This section presents my estimation results. Table 3.5 reports the parameter estimates for the baseline model and for a version in which all accumulation of export specific demand has been shut down. Overall, the baseline model performs well in matching the target moments and in accounting for the behavior of entrants into exporting in the year after entry. The model predicts an even larger growth in total exports than what is observed, even though the estimated price elasticity of demand in the export market of 1.42 is well within the low range commonly calibrated in open economy business cycle models. While both the baseline version and the restricted fixed costs version somewhat underpredict the growth in export participation towards the end of the sample, the baseline model clearly does a better job in replicating the data. Most importantly, while the fixed cost model predicts aggregate exports basically flat after 2002 in the absence of favorable macroeconomic shocks, the baseline model reproduces well the 50 percent surge between 2002 and 2008 as plants are still building up their foreign demand after the favorable shocks which occurred during the first half of the sample. Estimated entry costs of exporting are low compared with other estimates in the literature. It turns out, however, that by far the largest cost of exporting is spent on building up and maintaining a foreign demand base which constitutes a large sunk investment as well.

Table 3.5: Parameter estimates

	Baseline		**No customer capital**	
Parameter	Estimate	95% Conf. Int.	Estimate	95% Conf. Int.
Demand elasticity				
η	1.432	[1.248, 1.615]	3.472	[3.238, 3.706]
Foreign demand technology				
α	0.526	[0.508, 0.543]	–	–
δ	0.207	[0.191, 0.222]	–	–
Revenue productivity distribution				
ρ_ϕ	0.871	[0.862, 0.881]	0.992	[0.992, 0.992]
σ_ϕ	0.307	[0.303, 0.311]	0.369	[0.364, 0.373]
σ_χ	1.776	[1.704, 1.849]	3.346	[3.258, 3.343]
Costs of foreign demand accumulation				
c^{lin}	1.166	[0.994, 1.339]	–	–
c^{conv}	2.851	[2.448, 3.254]	–	–
Fixed costs of exporting				
γ_E	1.040	[0.990, 1.339]	1.679	[1.513, 1.844]
σ_E	0.347	[0.316, 0.379]	0.767	[0.679, 0.853]
γ_F	0.625	[0.582, 0.667]	0.610	[0.585, 0.634]
σ_F	1.241	[1.163, 1.319]	0.355	[0.310, 0.399]

Note: The table displays the estimated components of the parameter vector $\hat{\theta}$ along with the respective 95 percent confidence intervals. Columns two and three report results for the model baseline specification. The two rightmost columns reports results for a restricted model version in which I set α to zero such that firms do not accumulate export market specific demand. In consequence, the parameters δ, c^{lin}, c^{conv} are not part of the estimation either.

3.5.1 Parameter estimates

The second and third column of Table 3.5 report the estimated parameter vector $\hat{\theta}$ along with its 95 percent confidence interval for the model baseline specification. I also estimate a restricted model version with α set to zero which reduces the set-up to a pure fixed cost model like the one estimated in Das et al. (2007). This allows me to investigate in how far they differ in terms of predictive power and dynamic implications. Most parameters are fairly tightly estimated. I estimate a

---

[15]The growth rate in the year of entry is 2 by construction.

[16]The first quintile is linearly dependent on the other four quintiles and its inclusion into the estimation would add no further information.

price elasticity of demand of 1.432.[17] With confidence bounds at 1.248 and 1.615, this estimate is unfortunately not very precise. Still, the estimation procedure puts it well in the low range of about 0.9 to 2.0 typically calibrated in open economy macro models. At the same time it manages to account for the strong growth in average export revenue as result of the favorable shocks to the aggregate export environment.

3.5.2 Identification

Table 3.10 in the appendix reports a stylized Jacobi matrix evaluated at the baseline estimate of the parameter vector. As one can see there, many of the above moments are affected by a large subset of the parameter vector such that the estimate of any parameter cannot be attributed to one moment alone. Here, I want to give some more intuition for what subset of moments identifies which parameter.

- **Parameters of the foreign demand technology (α, δ)**

  The curvature of the foreign demand technology has very strong positive effects on the autocorrelation of foreign revenue growth as well as on the growth rates for plants that newly enter into exporting. The higher α, the higher the target level of foreign demand and the longer the plant will show positive revenue growth after entry or after a positive productivity shock. This also means that the distribution of foreign revenue is more spread out and a larger share of total revenue is generated in the top quintile. Finally, α decreases transitions into and out of exporting by increasing persistence in exporting. In that sense, the accumulation of a demand base acts much like a sunk entry cost spread out over time. Depreciation δ in many ways has an opposing effect. Larger attrition rates for foreign sales relationships mean less of an incentive to invest into them and lower growth rates after entry. This also implies a more compressed sales distribution. δ also decreases persistence in the export status and and increases transitions.

---

[17]An estimate of 1.432 for η implies a mark-up over marginal costs of 231 percent and a gross profit share in revenues of over 81 percent. One should keep in mind however, that an important part of German manufacturing exports constitute medium to high tech products and investment goods such that fixed operating and development costs make up an important part of overall costs. Also, these profits have to pay for entry and fixed costs of exporting and costs of accumulating and maintaining distribution channels which, as reported below, make up between 20 to 40 percent of revenues alone.

– **Demand elasticity (η)**

As mentioned above, the foreign demand elasticity is not estimated with a very strong precision. This also shows in the Jacobi matrix. η affects mean growth rate overall as well as for entrants.

– **Costs of foreign demand accumulation (c^{lin}, c^{con})**

A higher proportional profit loss in building a foreign customer stock decreases growth rates for entrants and means a more compressed revenue distribution overall. Higher convex investment costs have plants spread out their investment over more years after entry or a positive productivity shock. They do, however, also decrease the value of investment overall which, evaluated at the baseline estimate, actually means lower growth rates in the years three and four after entry. Experimentation showed this effect to be non-monotonic throughout the larger parameter space.

– **Revenue productivity distribution ($\rho_\phi, \sigma_\phi, \sigma_\chi$)**

Both ρ_ϕ and σ_ϕ have effects on the centered moments of the overall growth rates and they both induce positive autocorrelation as productivity innovations become larger and more persistent. More persistent innovations in particular also mean higher growth rates for entrants as they can expect to stay productive for a longer time and accumulate more demand. Increases in all three parameters spread out the sales distribution. In some sense σ_χ picks up residual dispersion after α and the parameters of the revenue productivity process have been identified.

– **Fixed costs of exporting ($\gamma_E, \sigma_E, \gamma_F, \sigma_F$)**

Sunk entry and continuation costs have an impact on almost all aspects of the model. They increase the higher moments and autocorrelation of export revenue growth by introducing stronger selection on productivity in entering exports. They make exit more likely for entrants which reduces initial revenue growth. The more compressed revenue distribution is also a consequence of stronger selection and less export participation. Finally, a higher sunk component in fixed costs introduces more persistence in exporting, a fact that has been exploited in previous estimations of sunk cost models[18].

---

[18]See for instance Roberts and Tybout (1997).

Figure 3.3: Export participation and total export revenue

|(a) Export participation|(b) Total exports|(c) Q_t|
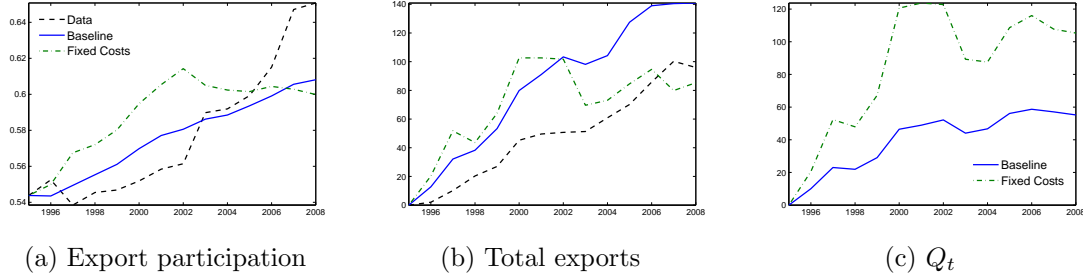
Note: Panels (a) and (b) display export participation and the cumulative growth in total export revenues in the data sample used for estimating the model (dash-dotted line) and compares it to the respective series in the baseline model (solid line) and the restricted model version (dashed line) where α is restricted to zero. Panel (c) displays a summary measure for aggregate export profitability, $Q_t \equiv D_t^W \left(\frac{RER_t}{1+\tau_t} \right)^\eta$, for the two model specifications.

Higher variability in sunk and fixed costs in many ways has the opposite effect off a higher mean. It makes export participation more volatile thereby increasing exit and (re-)entry and decreasing persistence in exporting. Higher variance in entry costs means more firms are going to enter who are unlikely to survive long. More variable fixed costs increase exit probability for all firms and in consequence also for entrants. This introduces selection which means higher average growth rates for those entrants that survive.

3.5.3 Model fit

Coming back to Table 3.4, the columns entitled $\mu(\hat{\theta})$ report the model generated moment vector evaluated at the baseline estimate $\hat{\theta}$ next to the target vector $\mu(\theta_0)$. Overall, the model provides a close fit to the data. Regarding the export revenue growth rate, the model reproduces the centered moments well and gets right the increase in one and two year autocorrelation, even though the increase is steeper in the data. Like in the data, entrants into exporting grow strongly in the years after entry and their survival probabilities are increasing in exporting tenure, though the data feature a step increase after the initial year, where the model produces a smoother profile in survival rates. The export revenue distribution in the model is more disperse than in the data and the model overestimates initial export participation.

3.5.4 Model comparison

I now want to turn to the question in how far the baseline model's implications differ from those of a model with stochastic sunk and fixed cost as only costs of exporting like the one estimated in Das et al. (2007). I therefore re-estimate the model setting α to zero and have $\delta, c^{lin}, c^{conv}$ take some arbitrary positive values. The parameter estimates of this restricted specification are reported in the two rightmost columns of Table 3.5. I then simulate the two estimated model versions for the duration of my estimation sample feeding in the observed series for aggregate tariffs, aggregate export demand and the real exchange rate and record the resulting series for aggregate export participation and cumulative growth in aggregate exports which were not specifically targeted in the estimation. The first two panels of Figure 3.3 compare the results to the data.

It follows from equation (3.2) that aggregate export revenue at time t in the model is given by

$$\frac{1}{\eta} \sum_{i=1}^{N} z_{it} D_{it}^{\alpha} D_t^{W} \left(\frac{RER_t}{1 + \tau_t} \right)^{\eta}.$$

Let me now define $Q_t \equiv D_t^{W} \left(\frac{RER_t}{1+\tau_t} \right)^{\eta}$ as summary measure of aggregate export profitability which moves the incentive to enter/stay in exporting and to increase foreign market presence. Again normalizing the sample beginning in 1995 to zero, I plot the cumulative growth in aggregate export profitability for both model versions in the rightmost panel (c) of Figure 3.3. While both series use the same aggregate price and demand series in their calculation, they differ due to the much higher elasticity of substitution η of 3.47 compared to only 1.43 in the baseline version. Both series experience a strong increase up to 2000 as tariffs drop and the German real exchange rate experiences a 20 percent depreciation vis-à-vis its trading partners. Afterwards the increase is much more subdued in case of the baseline model and even reversed to a degree as rising overall demand is counteracted by a rise in the real exchange rate of about 10 percent.

The much higher estimate for the demand elasticity η in the fixed cost model implies that aggregate exports are quite responsive to movements in the real exchange rate. This constitutes a version of what has previously been called the *elasticity puzzle* in international economics. In order to reconcile the high observed mean

growth rate in plant export revenue in response to the favorable changes in the aggregate export environment at the time, export demand has to be relatively elastic.

The two model versions differ markedly in their implications for both export participation and total export growth. They both underpredict the rise in export participation at the end of the sample because they do not catch the growth spurt after 2005. Neither do they catch the slight initial drop after 1996. Importantly, whereas the baseline model has export participation increase slowly after 1996, the fixed cost version predicts a steep increase which actually peaks in 2002 to then level off again. This is a consequence of the discussed strong sensitivity of aggregate export profitability to real exchange rate movements and clearly counterfactual to what is observed in the data.

Turning to the growth in aggregate exports, there is an important difference between the model versions. The baseline model predicts stronger growth, especially in the first six years, than is observed. Setting η a bit lower to a value still within confidence bounds moves the series closer to the data. It is however, a deficiency that it shares with the alternative specification. More importantly, the fixed cost model has aggregate exports pretty much flat after 2002 and even sees a decline which closely tracks the behavior of aggregate export profitability Q_t. Meanwhile, in the model with costly demand accumulation, export participation increases and firms slowly continue increasing their customer stock in reaction to the previous shocks. In consequence, aggregate exports rise by an additional 40 percent after 2003 which accords well with what we see in the data. As can one can also see in Figure 3.7, the baseline model matches the data well in another dimension by predicting the export expansion to be accompanied by a steady decline in exit rates whereas the fixed cost model predicts a rise in exit rates after 2002. The two models make similar predictions regarding entry.

3.5.5 Estimated Costs of Exporting

My structural approach allows me to quantify the size of sunk entry and continuation costs that plants face and the amount of resources they spend on maintaining and expanding their sales network. An estimate of 1.04 for γ_E implies that entry costs are on average .33 percent of the average plants export revenue in 1995 which translates

into 33,467 1995 euros. Analogously, mean continuation costs are 20,108 1995 euros[19]. Compared with an entry cost estimate of between \$344,000 and \$412,000 in terms of 1986 dollars for a sample of Columbian manufacturing plants in Das et al. (2007), these estimates seem surprisingly low. It turns out that by far the largest share of sunk expenditures in exporting arises from costly investment into the customer network. As one may also see in Figure 3.8, during the time of the sample, the average exporter spends about 19 percent of export revenues on maintaining its customer base and between 8 and 19 percent on further expanding it. In 1995, the beginning of the sample, this translates into an average expenditure of 2.4 million euros on maintaining customer capital and a further 0.96 million on expanding it.

The low entry and fixed cost estimates are also well in line with the findings in Willis and Ruhl (2009) who, using the same data set as Das et al. (2007) to estimate a two-country heterogeneous firm model with endogenous export participation, find entry costs of \$104,242 in terms of 1986 dollars. This estimate drops by a factor of 8 when exogenously imposing the slow export revenue growth for new exporters. Their fixed costs estimate of \$11,316 puts their results for the extended model version in the same ballpark as mine. There are as yet not many figures in the literature to which to compare the substantial estimates for money spent on marketing and maintaining and building a sales network. One exception is Arkolakis (2010) who reports that, depending on the definition used, American firms spend up to 7.7 percent of GDP on marketing activities. Taken together with the estimates presented here, this shows expenditures on building up demand, in the domestic or the foreign market, to be substantial and an important factor behind the slow revenue growth of plants after entering a market.

3.6 Macroeconomic Implications

This section investigates the macroeconomic implications of introducing costly demand accumulation into a model of endogenous export participation. I trace out the effects of three kinds of shocks for the behavior of aggregate exports, export participation and entry and exit. First, I look at a persistent real exchange rate

---

[19]The stochastic nature of entry and continuation costs implies that average costs paid are lower. On average, plants pay about 15.923 1995 euros to enter into exporting and 19.077 euros to continue.

depreciation. Second, I introduce a permanent stochastic innovation to aggregate export demand. Finally, I consider the dynamic effects of moving the economy into a free-trade environment by removing all tariffs.

3.6.1 Real exchange rate depreciation

Figure 3.4 displays the dynamic response of aggregate exports, export participation and export entry and exit to a one standard deviation real exchange rate depreciation. A first observation is that exports in the fixed cost model are much more responsive to the positive real exchange rate innovation than they are in the baseline version. The more than 12 percent increase in aggregate exports translates into an elasticity of 3.4. That is more than 70 percent higher than the elasticity of 1.93 implied by the baseline estimate. Also, whereas exports in the fixed cost model follow the real exchange rate in monotonically reverting to their long-run mean, exports take until the second year after the shock to reach their maximum response when firms take time to react to the favorable conditions by investing into new sales relationships.

Export participation reacts stronger in the fixed cost model as well. This is the result of both a more pronounced increase in entry and a much steeper initial decline in exit from exporting. Nonetheless, the participation response in the baseline model turns out to be more persistent, because exit stays below average for much longer. As the exchange rate returns to its mean plants stop investing into the export market and let their additional sales relations slowly dissolve. Since that takes time, they remain less likely to exit for many years after the shock. This does not, however, make aggregate exports more persistent than in the fixed cost model since the effect is mostly on smaller firms.

3.6.2 Positive innovation to export demand

Next, I look at the response to a positive innovation to aggregate export demand in Figure 3.5. In period zero, demand increases by 1.49 percent (one standard deviation) and stays at that level throughout the simulation. The real exchange rate remains fixed at its long-run average and tariffs are at the level of 1995. Again, the two models differ markedly in their behavior. Already the initial reaction of exports is much stronger in the baseline than it is in the restricted model, because firms immediately start building more market share. In the subsequent years, exports rise
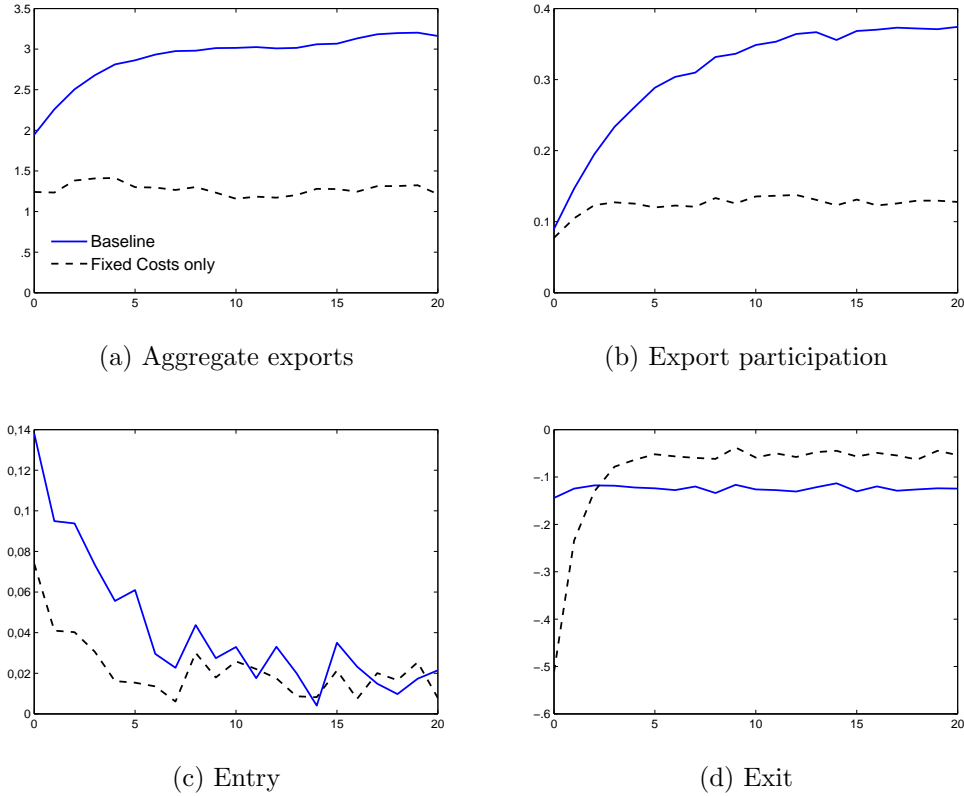
Figure 3.4: Dynamic response to a real exchange rate depreciation
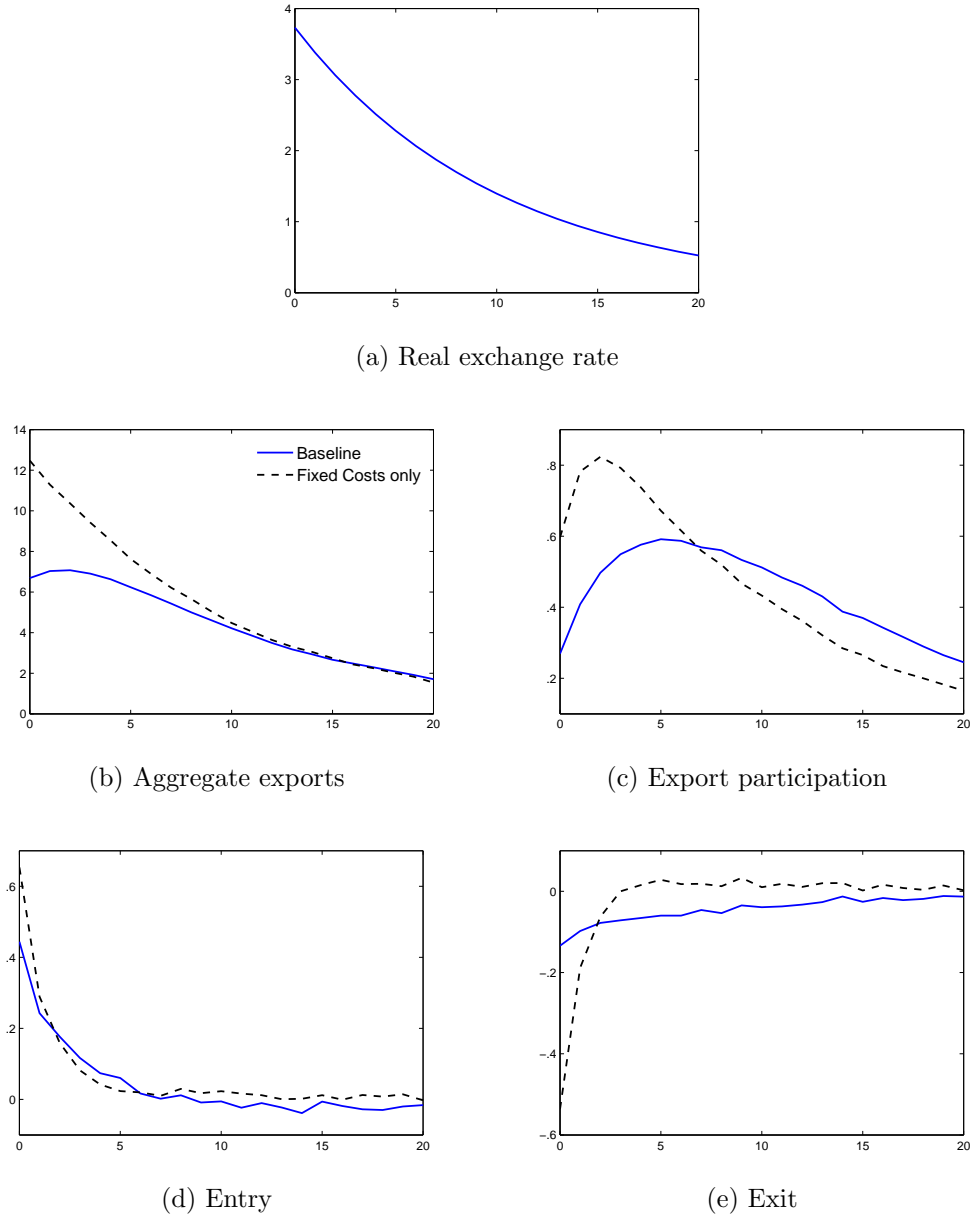
(a) Real exchange rate

(b) Aggregate exports

(c) Export participation

(d) Entry

(e) Exit

Note: This figure displays the dynamic response of aggregate exports, export participation and export entry and exit after a one standard deviation real exchange rate depreciation. The shock hits the economy in period zero and the graphs show the mean response in that year and the twenty years afterwards. The solid lines represent responses in the baseline model and the dashed lines represent the model with fixed entry and continuation costs only. The responses of real exchange rate and aggregate exports (Panel a and b) are in percentage deviations from steady state. The unit for participation, entry and exit (Panel c-e) is percentage point differences from steady state. Aggregate export demand is fixed at one. Tariffs are at the level of 1995.

Figure 3.5: Dynamic response to persistent increase in export demand

(a) Aggregate exports

(b) Export participation

(c) Entry

(d) Exit

Note: This figure displays the dynamic response of aggregate exports, export participation and export entry and exit after a one standard deviation innovation to aggregate export demand. The shock hits the economy in period zero and the graphs show the mean response in that year and the twenty years afterwards. The solid lines represent responses in the baseline model and the dashed lines represent the model with fixed entry and continuation costs only. The response of aggregate exports (Panel a) is in percentage deviations from steady state. The unit for participation, entry and exit (Panel b-d) is percentage point differences from steady state. The real exchange rate is fixed at one. Tariffs are at their level of 1995.

by another percentage point and it takes about five years for them to get there. While initially, participation in both cases increases by only .1 percent, this gain increases fourfold in the 10 years following the shock for the baseline case, but there is no comparable development in the fixed cost model.

The possibility to subsequently accumulate more demand increases the value of exporting more in the baseline model than it does in the fixed cost version as can be seen by an initial reaction in entry which is about twice as high. The explanation for the much larger increase in participation is similar to that for the more persistent reaction to the real exchange rate shock. Exit rates in the fixed cost model return to their mean two years after the positive innovation. They stay permanently lower

in the baseline model. Larger idiosyncratic customer capital again implies a larger sunk investment in exporting and in consequence lower exit probabilities.

3.6.3 Moving to free trade

Figure 3.6: Dynamic response to trade liberalization
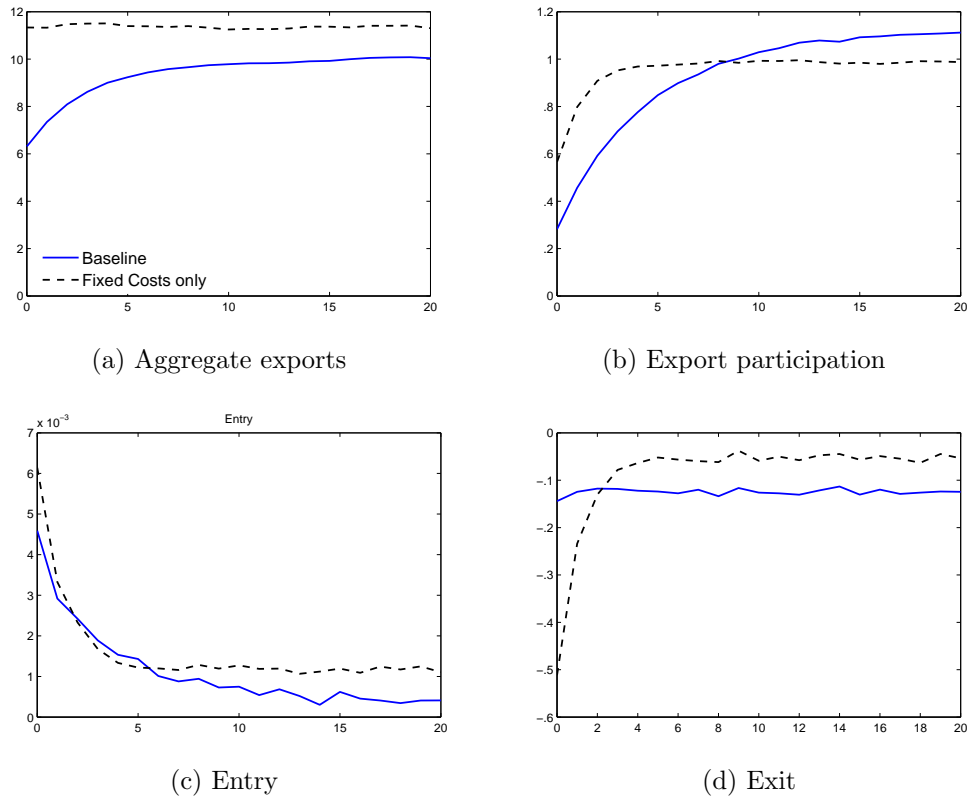
(a) Aggregate exports

(b) Export participation

(c) Entry

(d) Exit

Note: This figure displays the dynamic response of aggregate exports, export participation and export entry and exit after moving from the level of 1995 level of tariffs to free trade in period zero. The tariff reduction hits the plants unexpectedly in period zero. The solid lines represent responses in the baseline model and the dashed lines represent the model with fixed entry and continuation costs only. The response of aggregate exports (Panel a) is in percentage deviations from steady state. The unit for participation, entry and exit (Panel b-d) is percentage point differences from steady state. The real exchange rate and aggregate export demand are fixed at one.

Finally, I consider moving the economy from its 1995 tariff level to free trade in Figure 3.6. This implies a 10 percent long-run gain in aggregate exports in the case of the baseline model and an 11.5 percent gain for the alternative version. Export participation rises by 1.13 and 1.00 percent in the long run respectively. While these numbers may seem rather small, one should keep in mind that with a trade-weighted average of 3.3 percent, tariffs were already quite low in 1995. The larger expansion

in export participation in the baseline case again comes from a permanent drop in exit rates.

There are two important observations to make. First, as consequence of the demand friction that plants face when expanding export sales after the tariff drop the full realization of trade gains takes almost ten years. This is an important consideration to take into account when making welfare predictions about the reform. Second, both models predict long-run gains of similar size. Only the baseline model, however, is able to reconcile its prediction with a relatively low elasticity of exports to a real exchange rate depreciation as observed empirically. It therefore provides a solution to the elasticity puzzle in international economics by predicting both modest export reactions to real exchange rate movements and much larger gains after trade liberalizations.

3.7 Conclusion

In this chapter, I introduce a capital theoretic concept of the customer stock into a dynamic structural model of export participation. I thereby contribute to a recent literature in the areas of macroeconomics and international trade which investigates the implications of slow demand accumulation for the propagation of aggregate shocks and the distribution of export participation and intensity across international destinations. The model builds on the framework introduced in Das et al. (2007) and extends it by the introduction of a decreasing returns to scale technology which converts an export market specific factor into export demand. Plants have to spend resources on acquiring this factor which I call *customer capital.* A plant starts every new exporting spell without any customer capital. In consequence, investments into export demand constitute a type of sunk cost which causes firms to be more reluctant to exit the market when faced with bad shocks. Adjustments to the customer capital stock are subject to convex adjustment costs such that entrants spread investments over several years until they reach their desired level. Just like physical capital, export demand depreciates in the absence of new investment as trade partners go out of business and brands lose salience with consumers. Maintaining and expanding the customer stock therefore imposes steady costs on exporters.

I structurally estimate the model on a panel data set of German manufacturing

plants using a Simulated Method of Moments procedure. The data set spans the years 1995 to 2008, a time period that saw a strong increase in exporting activity among the plants in the sample. Export participation rose from 58 percent in the year 1995 to 66 percent in 2008 and total exports doubled in real terms. As consequence of the richness of the data set, I observe more than 14,000 export entry episodes in the sample. In the first four years of an export spell, plants display above average revenue growth rates and exit hazard rates are strongly declining with export tenure. These observations are consistent with a steady accumulation of export specific demand. While the entry into exporting is associated with strong growth in domestic revenues which hints at a positive productivity innovation, strong export revenue growth rates do not coincide with equally strong growth in the domestic market in the following years. This is an important indication that continued productivity growth is not responsible for the high export growth rate, at least insofar as productivity is not market specific.

I provide a first estimate of the costs related to export market demand accumulation. Estimated costs are substantial with plants on average having to spend between 26 and 38 percent of export revenues on maintaining and expanding their customer stock. On the other hand, estimated entry costs into exporting of 33,467 1995 euros are small compared to other estimates in the literature. The model fits the data well. Unlike a pure fixed cost model of export participation, it correctly predicts a strong growth in aggregate exports in the years after 2003 even though aggregate export profitability was flat in those years.

Regarding dynamic behavior, the model predicts total exports to keep expanding up to ten years after the hypothetical elimination of all tariffs for manufacturing exports. The total predicted rise in trade after eliminating an average of 3.3 percent ad valorem tariffs is 10 percent. This number is comparable in magnitude to the 11.5 percent a restricted alternative model without customer capital would predict. Unlike the fixed cost version, however, it is able to reconcile this prediction with a relatively subdued reaction of aggregate exports to real exchange rate movements. It therefore offers a possible answer to what has previously been called the *elasticity puzzle* in international economics - the discrepancy between high estimates for elasticities of substitution between goods from different countries from trade liberalization episodes and low estimates at business cycle frequency.

Appendix to Chapter3

3.A More information on the data set

Table 3.6: Summary statistics by sector

Sector (NACE rev. 1.1)	% of obs	Employment		Export	
		Mean	Total	Part. (%)	% in Rev.
Mining (10-14)	2.9	86	129,093	15.2	25.8
Food, Bev., Tobacco (15-16)	12.7	90	571,013	24.3	16.8
Textile, App. Leath (17-19)	4.7	96	226,591	71.7	23.2
Wood products (20)	4.1	57	115,840	39.4	17.1
Paper, Printing (21-22)	7.9	104	411,524	52.8	14.8
Chemicals (24)	3.5	268	480,703	81.1	34.4
Rubber, Plastics (25)	6.3	113	356,709	71.5	21.4
Non-metallic min. (26)	7.7	65	251,323	29.4	20.8
Metal, Metal prod. (27-28)	17	102	863,567	54.7	20.1
Machinery (29)	14	141	992,867	73.9	33.5
Electr. machinery (30-33)	11.4	152	871,488	60.1	30.4
Transport equip. (34-35)	3.1	569	908,038	66.3	26.0
Other manufacturing (36)	4.4	103	228,067	68.3	20.5

Note: The table displays summary statistics by sector for AFiD the sample of German manufacturing firms in 1999. Sectors are defined by broad NACE sectoral classifications. Statistics are the percentage of plant observations in the sector, average firm employment, total sectoral employment in the underlying population, export participation in percent and average share of exports in revenue for exporters. The 1999 sample contains a total of 50,154 plant observations.

Table 3.7: Employment size distribution

| Size class | % share | Employment | | Revenue | Revenue share (%) |
		Total	Share (%)	(Billion 1999 €)	
1 to 49	52.82	723,667	11.2	84.6	7.1
50 to 99	21.56	757,004	11.7	101.8	8.5
100 to 249	15.83	1,232,709	19.1	191.8	16.0
250 to 499	5.76	992,980	15.4	169.5	14.1
500 to 999	2.58	871,948	13.5	195.7	16.3
Over 1000	1.45	1,857,817	28.9	452.9	37.9
Total	100	6,436,127	100	1,196	100

Note: The table displays the distribution of plants in the sample over 6 employment size classes for the year 1999. Columns (3) to (6) additionally report total employment, share in total employment, total revenue and share in total revenue for each size class.

3.B Export entry and exit

Figure 3.7: Export transitions in the data and the two model specifications (1995-2008)

(a) Entry   (b) Exit
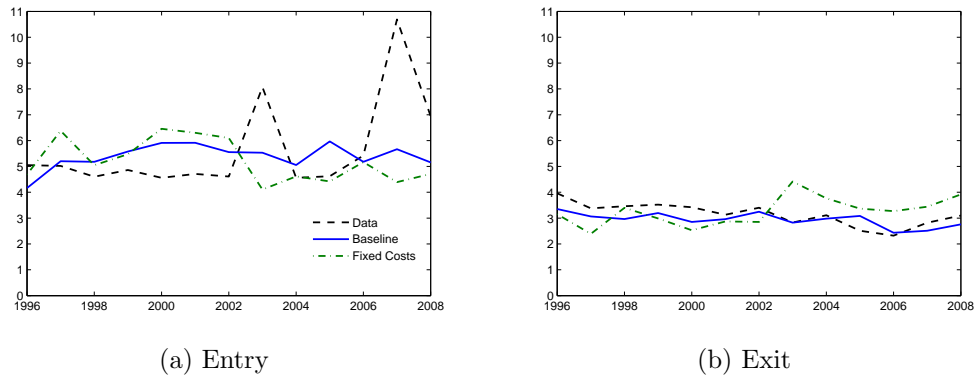
Note: Panels (a) and (b) display export entry and exit rates in the data sample (dashed lines) used for estimating the model and compares it to the respective series in the baseline model (solid lines) and the restricted model version (dash-dotted lines) where α is restricted to zero.

3.C Costs of customer capital accumulation

Figure 3.8: Average costs of demand base accumulation in terms of revenue
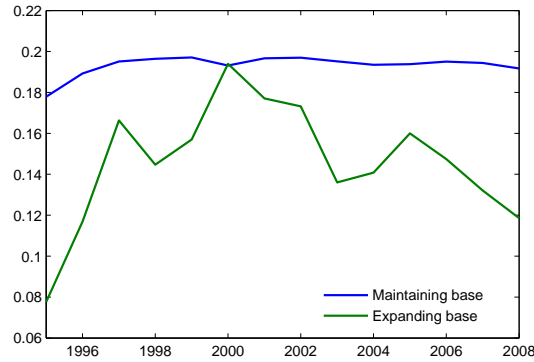
Note: The figure displays the average share of total export revenue that plants spent on maintaining and expanding their foreign customer base in a simulated version of the baseline model for the time of the sample (1995 to 2008) given the observed series for real exchange rates, export demand and tariffs.

3.D Aggregate variables

3.D.1 Tariffs

This section describes how I calculate the average tariff series used in simulating and estimating the model. The procedure consists of three steps. First, I calculate German manufacturing exports by year and by country using export data from the Eurostat Comext data base. I use those to identify the most important export destinations and to construct a stable aggregate "rest of the world" export market. Finally, I obtain disaggregated tariff data applicable to imports coming from the European Union for countries that are part of the aggregate. Weighting them by export volume (across destinations and within destinations across product lines) yields an average tariff measure.

The first challenge consists in identifying which export flows recorded in the trade data correspond to the sectors in my plant data base. The trade data are organized by the Harmonized System (HS) classification whereas sectors in the plant data are assigned based on the NACE rev. 1.1 classification. I match those sectoral classifications manually which requires some judgment based on the precise description

of the sectoral definitions. I choose to construct matches between two digit NACE
sectors and HS(2) chapters. Fortunately, in most cases identifying correspondences
is relatively straightforward. Table 3.8 provides some examples.

Table 3.8: NACE to HS matches - examples

NACE rev. 1.1 sector		Matched HS(2) chapter	
16	Manufacture of tobacco products	24	Tobacco Manuf. Tobacco substitutes
21	Manufacture of pulp, paper,	47	Pulp of wood, waste & scrap of paper
	and paper products	48	Paper & paperboard,
			articles of paper pulp
29	Manufacture of machinery	84	Nuclear reactors, boilers,
	and equipment n.e.c.		machinery & mechanical appliances,
30	Manufacture of office		computers
	machinery and computers		

Note: The table reports three three matches between the NACE (rev 1.1) sectoral classifi-
cation and Harmonized System chapters that serve to construct an average tariff measure
for German manufacturing exports between 1995 and 2008.

For the HS chapters identified as being manufacturing related, I then aggregate
up exports and rank export destinations by export value for every year separately.
The union set of the top twenty export destinations constitutes the world market
for German manufacturing exports between 1995 and 2008. Table 3.9 lists the coun-
tries included and their share in total exports. Throughout the sample period, the
aggregate stably accounts for 80 percent of total export volume.

For each member country I obtain six digit add valorem tariff data from the WTO.
Exports to EU countries and to Switzerland are tariff free. I first compute tariff
measures by country (value-weighted and unweighted) and then weight them again
by export volume by year to obtain an aggregate series of for applicable tariff duties.
Figure 3.9 displays the resulting series. The left panel displays the series used in
the estimations, where value-weights are used in computing the tariff measures by
country. The right panel is based on unweighted averages.

Figure 3.9: Tariff series

<table>
<tr><td>(a) Trade-weighted average</td><td>(b) Unweighted average</td></tr>
</table>
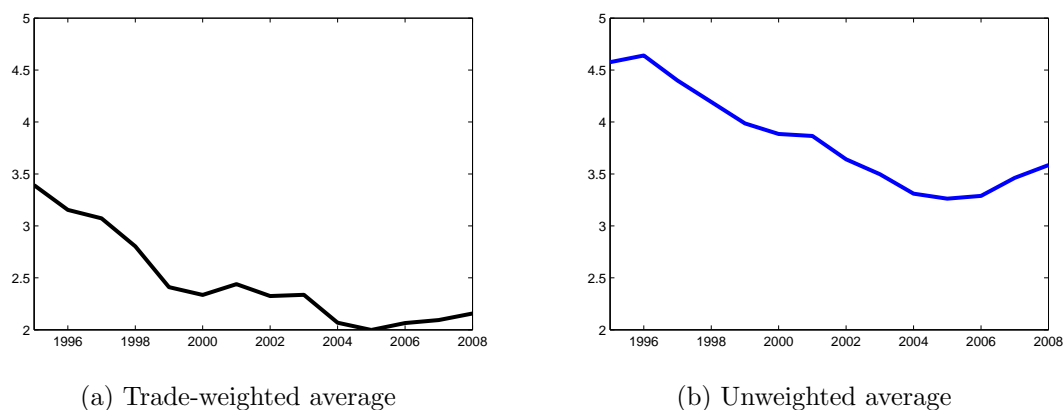
Note: The figure displays the average tariff series applicable to German manufacturing exports between 1995 and 2008. In the left panel, individual tariff lines have been value-weighted by their share in exports to some destination before weighting these averages by export value to destination. The right panel is based only on a value weighted average across countries. Individual tariff lines at some destination are not further weighted.

3.D.2 Real exchange rate

Figure 3.10: Real exchange rate series for Germany (1995-2008)
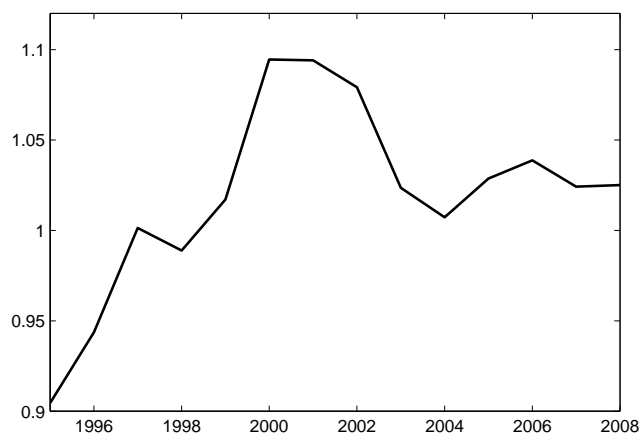
Note: The figure displays the trade-weighted real exchange rate for Germany between 1995 and 2008.

3.E Numerical Solution Method

This section of the appendix provides more details on the numerical solution algorithm. Source code for replication of the results will be available for download from

the author's web page at `http://www.uni-bonn.de/~s6votjad/`. I carry out the calculations of the time series for aggregate tariffs and aggregate demand in Stata. The code for estimating the processes and generating the graphs in the main chapter are written in the MATLAB® language. Finally, the estimation algorithm is written in FORTRAN 90. The program makes use of the NAG Numerical Library for numerical integration and some statistical calculations. To greatly speed up the solution of the value function iteration, some of the main work is carried out on an NVIDIA Tesla GPU using the CUDA Fortran set of extensions to the FORTRAN language from PGI. For an exposition to the benefits of using graphics processors to solve dynamic program problems see Aldrich et al. (2011). For an introduction to programing in CUDA Fortran you should consult Ruetsch and Fatica (2011). For readers who wish to replicate the results without access to an NVIDIA card, there is a non-GPU version of the program available from the author upon request.

3.E.1 Simulated Method of Moments Procedure

In order to obtain parameter estimates, the problem consists of minimizing the quadratic form given by equation (3.5) in the parameter vector $\hat{\theta}$. $\hat{\theta}$ has twelve elements, many of them without reasonable prior guesses to start from. Equation (3.5) has no known derivatives and is not guaranteed to be continuous. It therefore employ a Particle Swarm global optimization algorithm. This method tries to improve on a set of candidate solutions, called particles, by randomly moving them around the search-space where the size of changes depends on the particle's position and velocity. Movements are also guided by the global and local memory of the best solutions visited so far[21]. Upon convergence to the global minimum, the variance-covariance matrix of the estimate is given by equation (3.6), where $E_0 \frac{\partial \mu}{\partial \theta}$ is again calculated using numerical differentiation.

3.E.2 Value Function Iteration

During the estimation procedure, for any given set of parameters θ, firm policies for export participation and customer capital investment solve the dynamic programming problem described by equations (3.3) and (3.4). I approximate the continuous

---

[21]For an introduction to the method see Poli (2008). For a description of the algorithm I employ, you may also consult `http://www.nag.co.uk/numeric/fl/nagdoc_fl23/xhtml/E05/e05saf.xml`.

variables D, ϕ, χ, RER and D^W using discrete sets of grid points. I then iterate on equations (3.3) and (3.4) until convergence.

The customer capital grid is equi-spaced in $\log(D)$ with an upper bound of ωD_{stat}^{max}, where $D_{stat}^{max} = \left(\alpha \left(RER^{max}/\tau^{low} \right)^{\eta} DW^{max} \phi^{max} \chi^{max} \delta^{-1} \right)^{\frac{1}{1-\alpha}}$ is the optimal customer stock choice in the absence of uncertainty and convex adjustment costs evaluated at the largest grid points of the aggregate and idiosyncratic state variables and ω is a fixed fraction. Experimentation yielded $\omega = 0.6$ as a suitable choice for D_{stat}^{max} never to be a binding constraint on plant choices. In discretizing the AR(1)-processes for the persistent component of idiosyncratic revenue productivity ϕ and the real exchange rate RER, I use Tauchen (1986). The grid points for the permanent component of idiosyncratic productivity $\log(\chi)$ are the mean values of a set of equi-likely bins of the normal distribution $N(0, \sigma_{\chi})$. Finally, in discretizing the bounded random walk for $\log(D^W)$, I use an equi-distant set of grid points and obtain transition probabilities from a mixture of linear interpolation and Gaussian quadrature. During the estimation procedure I use 160 grid points for D , 10 grid points for ϕ and χ each, 7 grid points for RER and 5 grid points for D^w.

Export participation policies take the form of threshold realizations $\hat{\xi}^E$ and $\hat{\xi}^F$ and are therefore straightforward to calculate. The large dispersion in export revenues across plants in the data implies that customer capital policies span a large interval in $\mathbb{R}_0^+$. The challenge then consists of finding an accurate solution using only a relatively small number of grid points. Instead of optimizing over the discrete set of grid points, I use a golden section search optimization algorithm coupled with cubic spline interpolation of the value function[22].

The solution of the plant problem itself is subdivided into three steps. First, I solve for plant policies for the stochastic steady state that persists after 2008 in which tariffs are at a permanently low level. I then iterate backwards from 2008 to 1995 using the tariff series calculated in Section 3.3.4. Plants perfectly forecast the tariff series and they form expectations over D_{t+1}^W and RER_{t+1} given their estimated stochastic processes. Finally, I calculate policies that persist during the initial stochastic steady state ante 1995 in which plants are ignorant of the tariff changes

---

[22]For an introduction to the calculation and use of splines in numerical analysis you may consult de Boor (1978). In calculating spline coefficients, I also make use of two routines that come as open source code with that book. They are available for download at `http://orion.math.iastate.edu/burkardt/f_src/f_src.html`.

that will start taking place in that year.

3.E.3 Transition Simulation

The policy functions that solve the plant's dynamic problem serve as input into $S = 20$ Monte Carlo Simulations with $N = 50,000$ plants each. During the iterations over different parameter guesses, the random draws are fixed to assure convergence of the global minimization algorithm. The simulation consists of 150 initial periods during which global demand is fixed and tariffs are at their high initial level and of the transition period lasting 14 years between 1995 and 2008. I do not limit plant choices regarding D to be on-grid and evaluate them using spline interpolation. Policy functions for export entry and continuation policy I interpolate linearly. During the transition period, when I feed in observed realizations for export demand D^W and the real exchange rate RER, I first linearly interpolate policies in these two dimensions. The estimated parameter vector $\mu(\theta)$ is the average over the S simulations

3.E.4 Stylized Jacobi matrix

Table 3.9: Countries in world market aggregate

Country	Share in total exports (%)	
	1995	2008
Austria	5.49	5.51
Belgium[20]	6.49	5.10
Brazil	0.99	0.90
China	1.48	3.56
Czech Republic	1.63	2.84
Denmark	1.85	1.56
Finland	0.92	0.99
France	11.65	9.62
Hungary	.95	1.79
Italy	7.49	6.30
Japan	2.60	1.32
Netherlands	7.33	6.28
Poland	1.72	4.11
Portugal	0.87	0.84
Romania	0.35	0.91
Russia	1.38	3.28
South Korea	1.18	0.90
Spain	3.43	4.42
Sweden	2.45	2.06
Switzerland	5.41	3.88
Taiwan	1.01	0.50
Turkey	1.11	1.58
United States	7.60	7.42
United Kingdom	8.27	6.61
Total	80.56	79.15

Note: The table reports the set of countries I use to calculate aggregate series for tariffs and export demand. I select them by ranking German manufacturing export destinations for every year between 1995 and 2008 by value. The set of countries is the union set of the top twenty destinations in every year.

Table 3.10: Stylized Jacobi matrix at baseline estimate

	η	α	δ	ρ_ϕ	σ_ϕ	σ_χ	c^{lin}	c^{conv}	γ_E	σ_E	γ_F	σ_F
$Mean\left(\frac{\Delta R_t}{R_t}\right)$	+	+	o	−	-	o	o	o	o	o	o	o
$Std\left(\frac{\Delta R_t}{R_t}\right)$	o	-	+	o	+	o	o	o	−	+	-	+
$Skew\left(\frac{\Delta R_t}{R_t}\right)$	o	++	+	−	-	+	+	+	++	−	++	−
$Kurtosis\left(\frac{\Delta R_t}{R_t}\right)$	o	o	-	-	−	o	o	o	++	-	+	-
$Autocorr(1)$	o	++	o	++	++	o	o	o	++	−	++	−
$Autocorr(2)$	o	++	-	++	++	-	-	-	++	−	+	−
$GRE(1)$	o	++	o	++	o	o	o	o	-	+	o	+
$GRE(2)$	o	++	-	++	-	o	-	+	−	+	−	+
$GRE(3)$	+	++	-	++	++	-	-	-	−	++	−	++
$GRE(4)$	-	-	−	++	−	−	−	-	−	+	−	++
$SR(0)$	o	o	o	o	o	o	o	o	++	-	o	-
$SR(1)$	o	o	o	o	o	o	o	o	+	-	+	-
$SR(2)$	+	o	o	o	o	o	o	o	+	-	+	-
$SR(3)$	o	o	o	+	o	o	o	o	+	-	+	-
$SR(4)$	o	o	o	+	o	o	o	o	o	o	o	-
$SR(5)$	o	o	o	o	o	o	o	o	o	o	o	o
$SR(6)$	o	o	o	o	o	o	o	o	o	o	o	o
$SR(7)$	o	o	o	o	o	o	o	o	o	o	o	o
$SR(8)$	o	o	o	o	o	o	o	o	o	o	o	o
$Q2$	o	−	+	−	−	−	+	+	+	-	+	-
$Q3$	o	−	+	−	-	−	+	o	+	-	+	-
$Q4$	o	−	+	−	-	−	o	o	+	-	+	-
$Q5$	o	+	o	o	o	o	o	o	o	o	o	o
$NE-E$	o	−	+	-	o	-	o	o	−	++	−	++
$E-NE$	o	−	+	o	o	-	+	o	−	++	−	++
$Corr(Exp_t, Exp_{t-2})$	o	o	o	o	o	o	o	o	+	-	o	-
$Corr(Exp_t, Exp_{t-4})$	o	+	o	o	o	o	o	o	+	-	+	-
$Exp95$	o	o	-	+	o	o	o	o	-	o	-	o

Note: The table displays a stylized version of the Jacoby matrix evaluated at the baseline estimate of the model. The Jacoby matrix provides a good guide as to what elements of the moment vector μ are strongly affected by any given parameter and in turn, help identifying this parameter. − signifies an estimated elasticity of below -1 (strong negative effect), - means an estimated elasticity between -1 and -.2 (negative effect). o marks no important effect (though entries are scarcely zero). + represents an estimated elasticity of between .2 and 1 (positive effect). All elasticities larger than 1 are marked by ++.

Bibliography

ABOWD, J. M. AND D. CARD (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57, 411–445.

ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67, 251–334.

AIYAGARI, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 109, 659–84.

ALDRICH, E. M., J. FERNÁNDEZ-VILLAVERDE, A. RONALD GALLANT, AND J. F. RUBIO-RAMÍREZ (2011): "Tapping the supercomputer under your desk: Solving dynamic equilibrium models with graphics processors," *Journal of Economic Dynamics and Control*, 35, 386–393.

ALESSANDRIA, G. AND H. CHOI (2007): "Do Sunk Costs of Exporting Matter for Net Export Dynamics?" *The Quarterly Journal of Economics*, 122, 289–336.

——— (2011): "Establishment heterogeneity, exporter dynamics, and the effects of trade liberalization," Working Papers 11-19, Federal Reserve Bank of Philadelphia.

ALGAN, Y., O. ALLAIS, AND W. J. D. HAAN (2010): "Solving the incomplete markets model with aggregate uncertainty using parameterized cross-sectional distributions," *Journal of Economic Dynamics and Control*, 34, 59 – 68.

ALTONJI, J. G. AND N. WILLIAMS (1998): "The Effects of Labor Market Experience, Job Seniority, and job Mobility on Wages," *Research in Labor Economics*, 17, 233–276.

Bibliography

ARKOLAKIS, C. (2010): "Market Penetration Costs and the New Consumers Margin in International Trade," *Journal of Political Economy*, 118, 1151 – 1199.

———— (2011): "A Unified Theory of Firm Selection and Growth," NBER Working Papers 17553, National Bureau of Economic Research, Inc.

ARMINGTON, P. S. (1969): "A Theory of Demand for Products Distinguished by Place of Production," *Staff Papers - International Monetary Fund*, 16, pp. 159–178.

ATKESON, A. AND A. BURSTEIN (2008): "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, 98, 1998–2031.

BACHMANN, R. AND C. BAYER (2011a): "Investment Dispersion and the Business Cycle," NBER Working Papers 16861, National Bureau of Economic Research, Inc.

———— (2011b): "Uncertainty Business Cycles - Really?" NBER Working Papers 16862, National Bureau of Economic Research, Inc.

BACHMANN, R., R. J. CABALLERO, AND E. M. ENGEL (2010): "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model," Mimeo, Cowles Foundation Discussion Paper No. 1566R.

BACHMANN, R. AND L. MA (2012): "Lumpy Investment, Lumpy Inventories," Mimeo, RWTH Aachen University.

BACKUS, D. K., P. J. KEHOE, AND F. E. KYDLAND (1992): "International Real Business Cycles," *Journal of Political Economy*, 100, 745–75.

———— (1994): "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?" *The American Economic Review*, 84, 84–103.

BAGGER, J. AND R. LENTZ (2012): "An Equilibrium Model of Wage Dispersion and Sorting," Mimeo, https://sites.google.com/a/wisc.edu/rasmus-lentz/research.

BAXTER, M. AND M. J. CRUCINI (1993): "Explaining Saving–Investment Correlations," *The American Economic Review*, 83, pp. 416–436.

BAYER, C. AND V. TJADEN (2013): "Large Open Economies and Fixed Costs of Capital Adjustment," Mimeo, University of Bonn.

BERGER, D. AND J. VAVRA (2010): "Dynamics of the US Price Distribution," Mimeo, Yale University.

BERNARD, A. AND J. WAGNER (2001): "Exports and success in German manufacturing," *Review of World Economics (Weltwirtschaftliches Archiv)*, 133, 134–157.

BERNARD, A. B. AND J. B. JENSEN (2004): "Why Some Firms Export," *The Review of Economics and Statistics*, 86, 561–569.

BLOOM, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

BONTEMPS, C., J.-M. ROBIN, AND G. J. VAN DEN BERG (2000): "Equilibrium Search with Continuous Productivity Dispersion: Theory and Nonparametric Estimation," *International Economic Review*, 41, 305–58.

BURDETT, K., C. CARRILLO-TUDELA, AND M. COLES (2011): "Human Capital Accumulation and Labour Market Equilibrium," *International Economic Review*, 52, 657–677.

BURDETT, K. AND D. T. MORTENSEN (1998): "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39, 257–73.

CABALLERO, R. J. AND E. M. R. A. ENGEL (1999): "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach," *Econometrica*, 67, 783–826.

CABALLERO, R. J., E. M. R. A. ENGEL, AND J. C. HALTIWANGER (1995): "Plant-Level Adjustment and Aggregate Investment Dynamics," *Brookings Papers on Economic Activity*, 26, 1–54.

CARRILLO-TUDELA, C. (2012): "Job Search, Human Capital and Wage Inequality," Working Paper No. 6949, IZA.

CENTER FOR ECONOMIC AND POLICY RESEARCH (2010): "SIPP Uniform Extracts, Version 2.1.5," .

Bibliography

CHANEY, T. (2011): "The Network Structure of International Trade," NBER Working Papers 16753, National Bureau of Economic Research, Inc.

CHANG, Y., S.-B. KIM, AND F. SCHORFHEIDE (2010): "Labor-Market Heterogeneity, Aggregation, and the Lucas Critique," NBER Working Papers 16401, National Bureau of Economic Research, Inc.

CHARI, V. V., P. J. KEHOE, AND E. R. MCGRATTAN (2002): "Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates?" *Review of Economic Studies*, 69, 533–63.

CLARK, K. B. AND L. H. SUMMERS (1979): "Labor Market Dynamics and Unemployment: A Reconsideration," *Brookings Papers on Economic Activity*, 10, 13–72.

COLEMAN, T. F. AND Y. LI (1996): "An Interior Trust Region Approach for Nonlinear Minimization Subject to Bounds," *SIAM Journal on Optimization*, 6, 418–445.

CONNOLLY, H. AND P. GOTTSCHALK (2008): "Wage Cuts as Investment in Future Wage Growth," *Labour*, 22, 1–22.

COOPER, R., J. HALTIWANGER, AND L. POWER (1999): "Machine Replacement and the Business Cycle: Lumps and Bumps," *American Economic Review*, 89, 921–946.

COOPER, R. W. AND J. C. HALTIWANGER (2006): "On the Nature of Capital Adjustment Costs," *Review of Economic Studies*, 73, 611–633.

DAS, S., M. J. ROBERTS, AND J. R. TYBOUT (2007): "Market Entry Costs, Producer Heterogeneity, and Export Dynamics," *Econometrica*, 75, 837–873.

DAVIS, S. J., J. C. HALTIWANGER, AND S. SCHUH (1998): *Job Creation and Destruction*, vol. 1 of *MIT Press Books*, The MIT Press.

DE BOOR, C. (1978): *A Practical Guide to Splines*, vol. 27, Springer, 2001, revised, illustrated ed.

152

DEN HAAN, W. J. (2010): "Assessing the Accuracy of the Aggregate Law of Motion in Models with Heterogeneous Agents," *Journal of Economic Dynamics and Control*, 34, 79–99.

DEN HAAN, W. J., K. L. JUDD, AND M. JUILLARD (2011): "Computational Suite of Models with Heterogeneous Agents II: Multi-Country Real Business Cycle Models," *Journal of Economic Dynamics and Control*, 35, 175 – 177.

DEN HAAN, W. J. AND P. RENDAHL (2010): "Solving the incomplete markets model with aggregate uncertainty using explicit aggregation," *Journal of Economic Dynamics and Control*, 34, 69 – 78.

DOMS, M. E. AND T. DUNNE (1998): "Capital Adjustment Patterns in Manufacturing Plants," *Review of Economic Dynamics*, 1, 409–429.

DROZD, L. A. AND J. B. NOSAL (2012): "Understanding International Prices: Customers as Capital," *American Economic Review*, 102, 364–95.

DUFFIE, D. AND K. J. SINGLETON (1993): "Simulated Moments Estimation of Markov Models of Asset Prices," *Econometrica*, 61, 929–52.

DUSTMANN, C. AND C. MEGHIR (2005): "Wages, Experience and Seniority," *Review of Economic Studies*, 72, 77–108.

EATON, J., M. ESLAVA, C. J. KRIZAN, M. KUGLER, AND J. TYBOUT (2010): "A Search and Learning Model of Export Dynamics," .

FALLICK, B. AND C. A. FLEISCHMAN (2004): "Employer-to-employer flows in the U.S. labor market: the complete picture of gross worker flows," Finance and Economics Discussion Series 2004-34, Board of Governors of the Federal Reserve System (U.S.).

FIORI, G. (2012): "Lumpiness, capital adjustment costs and investment dynamics," *Journal of Monetary Economics*, 59, 381–392.

FITZGERALD, D. AND S. HALLER (2012): "Exporters and Shocks," Mimeo, Stanford University.

Bibliography

FLINN, C. AND J. MABLI (2008): "On-the-Job Search, Minimum Wages, and Labor Market Outcomes in an Equilibrium Bargaining Framework," Carlo Alberto Notebooks 91, Collegio Carlo Alberto.

FOSTER, L., J. HALTIWANGER, AND C. SYVERSON (2008): "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98, 394–425.

FOSTER, L., J. C. HALTIWANGER, AND C. SYVERSON (2012): "The Slow Growth of New Plants: Learning about Demand?" NBER Working Papers 17853, National Bureau of Economic Research, Inc.

FRANKEL, J. A. AND A. K. ROSE (1995): "A Panel Project on Purchasing Power Parity: Mean Reversion Within and Between Countries," NBER Working Papers 5006, National Bureau of Economic Research, Inc.

FUJITA, S. (2011): "Reality of on-the-job search," Working Paper 10-34/R, Federal Reserve Bank of Philadelphia.

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

GOURIO, F. AND A. K. KASHYAP (2007): "Investment spikes: New Facts and a General Equilibrium Exploration," *Journal of Monetary Economics*, 54, 1–22.

GOURIO, F. AND L. RUDANKO (2011): "Customer Capital," NBER Working Papers 17191, National Bureau of Economic Research, Inc.

HAGEDORN, M. AND I. MANOVSKII (2010): "Search Frictions and Wage Dispersion," Mimeo, http://economics.sas.upenn.edu/∼manovski/research.html.

HALL, R. E. AND P. R. MILGROM (2008): "The Limited Influence of Unemployment on the Wage Bargain," *American Economic Review*, 98, 1653–1674.

HANSEN, G. D. (1985): "Indivisible labor and the business cycle," *Journal of Monetary Economics*, 16, 309 – 327.

HEATHCOTE, J. AND F. PERRI (2002): "Financial autarky and International Business Cycles," *Journal of Monetary Economics*, 49, 601 – 627.

HOLZER, H. J. (1988): "Search Method Use by Unemployed Youth," *Journal of Labor Economics*, 6, 1–20.

HOPENHAYN, H. AND R. ROGERSON (1993): "Job Turnover and Policy Evaluation: A General Equilibrium Analysis," *Journal of Political Economy*, 101, 915–38.

HOPENHAYN, H. A. (1992): "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60, 1127–50.

HORNSTEIN, A., P. KRUSELL, AND G. L. VIOLANTE (2007): "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," NBER Working Papers 13674, National Bureau of Economic Research, Inc.

——— (2012): "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," *American Economic Review*, 101, 2873–98.

HOUSE, C. L. (2008): "Fixed Costs and Long-Lived Investments," NBER Working Papers 14402, National Bureau of Economic Research, Inc.

HUGGETT, M. (1993): "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of Economic Dynamics and Control*, 17, 953–969.

JOLIVET, G., F. POSTEL-VINAY, AND J.-M. ROBIN (2006): "The empirical content of the job search model: Labor mobility and wage distributions in Europe and the US," *European Economic Review*, 50, 877–907.

KHAN, A. AND J. K. THOMAS (2003): "Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities matter?" *Journal of Monetary Economics*, 50, 331–360.

——— (2008): "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics," *Econometrica*, 76, 395–436.

KIM, S. H., R. KOLLMANN, AND J. KIM (2010): "Solving the incomplete market model with aggregate uncertainty using a perturbation method," *Journal of Economic Dynamics and Control*, 34, 50 – 58.

Bibliography

Kircher, P. A. and J. Eeckhout (2011): "Identifying Sorting–In Theory," *The Review of Economic Studies*, 78, 872–906.

Klau, M. and S. S. Fung (2006): "The new BIS effective exchange rate indices," Bis quarterly review, Bank for International Settlements.

Krueger, D., F. Perri, L. Pistaferri, and G. L. Violante (2010): "Cross Sectional Facts for Macroeconomists," *Review of Economic Dynamics*, 13, 1–14.

Krusell, P. and A. A. Smith (1997): "Income And Wealth Heterogeneity, Portfolio Choice, And Equilibrium Asset Returns," *Macroeconomic Dynamics*, 1, 387–422.

——— (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *The Journal of Political Economy*, 106, 867–896.

Kydland, F. E. and E. C. Prescott (1982): "Time to Build and Aggregate Fluctuations," *Econometrica*, 50, 1345–70.

Lee, B.-S. and B. F. Ingram (1991): "Simulation estimation of time-series models," *Journal of Econometrics*, 47, 197–205.

Low, H., C. Meghir, and L. Pistaferri (2010): "Wage Risk and Employment Risk over the Life Cycle," *American Economic Review*, 100, 1432–1467.

Lucas, R. J. (1972): "Expectations and the neutrality of money," *Journal of Economic Theory*, 4, 103–124.

Malin, B. A., D. Krueger, and F. Kubler (2011): "Solving the multi-country real business cycle model using a Smolyak-collocation method," *Journal of Economic Dynamics and Control*, 35, 229–239.

McFadden, D. (1989): "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration," *Econometrica*, 57, 995–1026.

Meghir, C. and L. Pistaferri (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72, 1–32.

MIAO, J. AND P. WANG (2011): "A Q-Theory Model with Lumpy Investment," Mimeo, Boston University.

MICHELACCI, C., J. PIJOAN-MAS, AND H. RUFFO (2012): "Inequality in Unemployment Risk and in Wages," Mimeo, https://sites.google.com/site/hernanruffo/research.

MORTENSEN, D. (2003): *Wage Dispersion: Why are Similar Workers Paid Differently*, MIT Press.

MOSCARINI, G. AND K. THOMSSON (2007): "Occupation and Job Mobility in the US," *Scandinavian Journal of Economics*, 109, 807–836.

NÁGIPAL, E. (2005): "On the Extend of Job-to-Job Transitions," Mimeo, http://faculty.wcas.northwestern.edu/~een461/research.html.

———— (2008): "Worker Reallocation Over the Business Cycle: Importance of Employer-to-Employer Transitions," Mimeo, http://faculty.wcas.northwestern.edu/~een461/research.html.

ORTEGO-MARTI, V. (2012): "Unemployment History and Frictional Wage Dispersion," Mimeo, http://faculty.ucr.edu/~victorom/.

PAPP, T. (2012): "Frictional Wage Dispersion with Bertrand Competition: An Assessment," Mimeo, http://elaine.ihs.ac.at/~tpapp/.

POLI, R. (2008): "Analysis of the publications on the applications of particle swarm optimisation," *J. Artif. Evol. App.*, 2008, 3:1–3:10.

POSTEL-VINAY, F. AND J.-M. ROBIN (2002): "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, 70, 2295–2350.

REITER, M. (2010): "Approximate and Almost-Exact Aggregation in Dynamic Stochastic Heterogeneous-Agent Models," Economics Series 258, Institute for Advanced Studies.

ROBERTS, M. J. AND J. R. TYBOUT (1997): "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs," *American Economic Review*, 87, 545–64.

Bibliography

RUETSCH, G. AND M. FATICA (2011): *CUDA Fortran for Scientists and Engineers*, NVIDIA Corporation, 2701 San Tomas Expressway, Santa Clara, CA 95050.

RUHL, K. J. (2008): "The International Elasticity Puzzle," Working Papers 08-30, New York University, Leonard N. Stern School of Business, Department of Economics.

SARGENT, T. J. AND N. WALLACE (1976): "Rational expectations and the theory of economic policy," *Journal of Monetary Economics*, 2, 169–183.

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

SHIMER, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95, 25–49.

SIEGEL, J. (2002): *Stocks for the Long Run*, McGraw-Hill.

STINSON, M. (2003): "Technical Description of SIPP Job Identification Number Editing in the 1990-1993 SIPP Panels," Mimeo, U.S. Census Bureau.

STORESLETTEN, K., C. I. TELMER, AND A. YARON (2004): "Consumption and Risk Sharing over the Life Cycle," *Journal of Monetary Economics*, 51, 609–633.

SUSTEK, R. (2011): "Plant-level nonconvex output adjustment and aggregate fluctuations," *Journal of Monetary Economics*, 58, 400 – 414.

SYVERSON, C. (2004): "Product Substitutability and Productivity Dispersion," *The Review of Economics and Statistics*, 86, 534–550.

TAUCHEN, G. (1986): "Finite state markov-chain approximations to univariate and vector autoregressions," *Economics Letters*, 20, 177–181.

TJADEN, V. (2013): "Foreign Customer Accumulation and Export Dynamics," Mimeo, University of Bonn.

TJADEN, V. AND F. WELLSCHMIED (2012): "Exploring the Causes of Frictional Wage Inequality," IZA Discussion Papers 6299, Institute for the Study of Labor (IZA).

158

———— (2013): "Quantifying the Contribution of Search to Wage Inequality," Mimeo, University of Bonn.

Topel, R. H. (1991): "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority," *Journal of Political Economy*, 99, 145–176.

Topel, R. H. and M. P. Ward (1992): "Job Mobility and the Careers of Young Men," *The Quarterly Journal of Economics*, 107, 439–79.

Willis, J. L. and K. J. Ruhl (2009): "New Exporter Dynamics," 2009 Meeting Papers 111, Society for Economic Dynamics.

Young, E. R. (2010): "Solving the Incomplete Markets Model with Aggregate Uncertainty using the Krusell-Smith Algorithm and Non-Stochastic Simulations," *Journal of Economic Dynamics and Control*, 34, 36–41.