

Firm-to-firm Trade in Sticky Production Networks

Kevin Lim*

Princeton University

January 2016

Abstract

This paper studies the quantitative implications of frictions in the creation and destruction of firm-to-firm trading relationships for aggregate patterns of output and trade. I develop a structural model of trade between heterogeneous firms in which the network of firm-level input-output linkages is determined both dynamically and endogenously. The model generates rich predictions regarding firm connectivity, matching, and relationship dynamics, while remaining computationally tractable. Using both cross-sectional and panel data on trading relationships between US firms, I estimate the model's parameters and show that the model adeptly fits empirical regularities documented in the paper. I then study the model's predicted responses of trade patterns to counterfactual shocks, with four key results. First, endogenous adjustment of firm-to-firm relationships dynamically amplifies the effects of changes in variable trade costs on trade volumes and welfare by more than three times. Second, reductions in the cost of maintaining relationships have effects on trade and welfare that are over 50% larger than cost-equivalent reductions in variable trade costs. Third, stickiness in firm-level relationships imparts a high degree of inertia to the dynamics of aggregate trade and output, with typical responses to shocks exhibiting half-lives of around two years. Finally, the model suggests that taxing trade flows to subsidize the formation of firm-level trading relationships can be welfare improving.

**klim@princeton.edu*. I am extremely grateful to Gene Grossman, Steve Redding, and Esteban Rossi-Hansberg for their support and guidance in writing this paper. I also thank Davin Chor, Oleg Itskhoki, Eduardo Morales, Ezra Oberfield, as well as numerous seminar participants and visiting faculty at Princeton University for helpful discussions and comments. This research benefited from financial support from the International Economics Section at Princeton University.

1 Introduction

Many of the goods and services that are traded between firms lack centralized markets or intermediaries facilitating their exchange.¹ Such firm-to-firm trade is therefore contingent on firms' active management of direct relationships with their customers and suppliers. This can often be an integral yet costly aspect of operations. Market analysts estimate, for example, that firms in the United States spent more than \$10bn in 2014 on customer relationship management (CRM) and supply chain management (SCM) software systems alone.² Motivated by this observation, this paper studies the quantitative implications of frictions in the creation and destruction of firm-to-firm trading relationships (henceforth referred to as relationship stickiness) for aggregate output and trade across locations. When it is costly to form and adjust trading relationships, how do firms vary their selection of trade partners in response to changes in the economic environment? Consequently, how do these decisions translate into the responses of aggregate output and trade to macroeconomic shocks?

To answer these questions, I develop a structural model of trade between heterogeneous firms in which the network of firm-level input-output linkages is determined both dynamically and endogenously. In the model, monopolistically-competitive firms in different locations produce output using a technology exhibiting constant returns to scale and a constant elasticity of substitution across inputs. Access to additional customers therefore increases the variable profit of a firm, while access to additional suppliers lowers its marginal cost. These incentives to form trading relationships are counterbalanced by assuming that firms face a fixed cost per active relationship, and that the opportunity to activate or terminate each relationship arrives randomly over time.³ The static fixed cost creates a meaningful tradeoff for firms in their selection of relationships, while the dynamic opportunity cost makes these selection problems forward-looking. These assumptions therefore allow the model to generate rich predictions regarding the distributions of customers and suppliers across firms, the

¹This is a point dating back at least to Rauch (1999), who was one of the first to argue using empirical evidence for a view of trade as characterized by networks of buyers and sellers rather than by frictionless markets.

²See for instance the reports by Gartner, Inc. (2014a, 2014b). Software platforms marketed by industry leaders such as Salesforce and SAP offer solutions for a wide range of relationship management tasks, such as the organization of contact databases, monitoring of customer and supplier financial information, tender and contract management, supplier performance assessment, and so on. This highlights the potentially complex nature of the costs that firms face in managing business relationships, of which expenses on software are only one particular facet.

³The fixed relationship cost is analogous to the fixed cost of exporting in Melitz (2003), except that here it is paid at the firm-to-firm level. The random arrival of opportunities to reset the status of a relationship is analogous to the price reset shock in Calvo (1983), except that here firms are constrained in their ability to adjust relationships along the extensive rather than the intensive margin.

assortativity of matching between firms, the persistence of firm-to-firm relationships across time, as well as the differential responses of these patterns to aggregate shocks in the short-versus the long-run.

At the same time, the model remains computationally tractable. Cross-sectional firm-level variables are pinned down by sufficient statistics that are easily computed for any input-output architecture, and solving for the model's transition dynamics under rational firm expectations typically requires about one hour on a standard personal computer. Computational tractability in turn permits structural estimation of the model and the quantitative analysis of counterfactual exercises. Using both cross-sectional and panel data on firm-level trading relationships in the United States (obtained from Standard and Poor's Capital IQ and Compustat platforms), I estimate the model's parameters via a simulated method of moments technique. I show that the model is able to replicate the majority of empirical regularities that I document in the paper, with larger firms tending to: (1) have more suppliers and customers; (2) trade with larger and more connected firms; and (3) have trading relationships that are more persistent. I then study the quantitative responses of trade patterns and welfare to counterfactual changes in trade costs, changes in relationship costs, and idiosyncratic firm-level fluctuations.

The key findings of this paper are as follows. First, the endogenous adjustment of firm-to-firm trading relationships dynamically amplifies the effects of changes in variable trade costs on aggregate interfirm trade and welfare. Intuitively, when relationships are sticky, a fall in trade costs induces firms to not only buy more from existing trade partners but also to accumulate more trade partners over time. Quantitatively, the magnitude of this amplification effect is large: the elasticities of aggregate trade and welfare with respect to trade costs are estimated to be between three to four times higher in the long-run than in the short-run.⁴ This suggests that taking into account the timing of policies aimed at reducing trade costs can be important, and in particular provides a rationale for quick rather than gradual reduction of trade barriers.

Second, reductions in relationship fixed costs have stronger effects on aggregate trade and welfare than cost-equivalent reductions in variable trade costs. Consider a planner with an exogenous subsidy budget who can choose to either subsidize the intensive margin of trade (through export or import subsidies for example) or to subsidize the fixed cost of each active relationship (by mitigating communication or meeting costs for instance). The model's counterfactuals predict that the latter option would generate increases in aggregate trade and

⁴The magnitude of this dynamic amplification effect is similar to the size of the amplification effect that Alessandria, Choi, and Ruhl (2015) estimate, which in their model is generated by firm-level investments in lowering export costs that respond endogenously to changes in trade barriers.

welfare that are more than 50% larger in the long-run than the gains that would be realized under the former option, with similar rates of dynamic adjustment. This implies that policy measures which reduce the frictions that firms face in establishing trading relationships can be equally as if not more cost-effective than traditional trade policy instruments in terms of their ability to increase trade and welfare. This may be of particular interest for policymakers who find the direct promotion of firm-to-firm relationships to be less politically-objectionable than adjustments in tariff barriers.

Third, when firm relationships are sticky, both macroeconomic shocks as well as idiosyncratic fluctuations in firm-level characteristics can have effects on aggregate trade and output that are not only large but persistent as well. Following a decline in trade or relationship costs, the dynamic adjustments of trade volumes and welfare typically exhibit half-lives of around two years. Similarly, idiosyncratic shocks to firm-level characteristics generate declines in trade and welfare that dissipate gradually with a half-life of around two years, even when such shocks leave the aggregate distribution of firm characteristics unchanged (and therefore would have no aggregate effect in a frictionless model). The endogenous adjustment of firm-to-firm relationships due to relationship stickiness therefore imparts a high degree of inertia to the dynamics of aggregate outcomes, whether these dynamics are driven by macroeconomic shocks or by idiosyncratic firm-level fluctuations.

Finally, a simple policy exercise shows that subsidies to the cost of maintaining relationships with customers financed by a tax on imports can improve welfare. This suggests that firms in the market equilibrium are trading too much at the intensive margin and too little at the extensive margin relative to the social optimum. I show analytically that inefficiency of the market equilibrium stems from two sources. The first is the standard markup distortion arising from firm monopoly power. The second is a novel source of inefficiency generated by the network structure of production (often referred to as a network externality): firms select relationships based only on profit-maximizing criteria and do not internalize the value of each relationship to all other firms in the network.

The modeling of frictions in firm-level trading relationships in this paper is most closely related to the models of Oberfield (2015) and Chaney (2014, 2015). In both of these models, potential buyer-supplier pairs also receive trading opportunities at a finite rate, and the network of firm-level input-output linkages is an endogenous and dynamic outcome of this exogenous stochastic process. However, there are two key differences between these frameworks and the model that I develop. First, I introduce a fixed cost to relationship formation, whereas activating a trading relationship is costless for firms in both Oberfield (2015) and Chaney (2014, 2015).⁵ It is this costly nature of relationship formation that gen-

⁵In Oberfield (2015), firms always have the option of buying from suppliers that they could have traded

erates the dynamic amplification of shocks discussed above, which I find to be quantitatively large. Explicitly modeling the costs of relationship formation also allows me to study the effect of reductions in such costs on aggregate patterns of output and trade.⁶ Second, both Oberfield (2015) and Chaney (2014, 2015) only partially model variations in the extensive margin of firm-to-firm trading relationships.⁷ These models therefore lose identifying power that might otherwise be gained by exploiting richer heterogeneity in empirically observed networks of firm-to-firm trade. For this reason, I construct a model that simultaneously generates non-trivial predictions about the distributions of both customers and suppliers across firms.

The theory developed in this paper is also related to the broader theoretical literature on social and economic network formation, within which there are two qualitatively different approaches to modeling the formation of ties between atomistic agents.⁸ The first approach posits an exogenous stochastic algorithm for the formation of links, and then proceeds to study the resulting network properties.⁹ As these models of network formation are non-structural, however, they cannot be used to study how networks of trade between firms respond to changes in economic incentives. The second approach to modeling network formation assumes that the creation and destruction of links are the result of strategic interactions between agents.¹⁰ These game-theoretic approaches therefore explicitly take into account optimizing behavior by the agents constituting the network, but the complexity of solving these models beyond simple illustrative examples precludes quantitative analysis.

The modeling of network formation in this paper can thus be viewed as a combination of the two approaches discussed above, or in the terminology of Currarini, Jackson, and Pin (2010), a combination of “chance and choice”: firms receive the opportunity to adjust relationships according to an exogenous stochastic process, but the activation or termination

with in the past, while in Chaney (2014, 2015), trade occurs automatically once a potential seller acquires contact with a buyer. In both models, there are no fixed costs of trade between firms.

⁶Bernard, Moxnes, and Saito (2015) and Bernard, Moxnes, and Ulltveit-Moe (2015) explicitly model fixed relationship costs between firms in the same way that I do here. However, these papers model only the static formation of relationships between one group of buyers and one group of sellers - in essence capturing only one tier of the static network of trade between firms.

⁷In Oberfield (2015), the number of suppliers per firm is exogenously fixed, while in Chaney (2014, 2015), every firm has the same number of suppliers even though the number of suppliers per firm grows over time.

⁸See Jackson (2005, 2011) for more in-depth surveys of the network formation literature.

⁹Well-known examples from the graph theory literature are the Erdős-Rényi (1959) random network, the Watts-Strogatz (1998) small world model, and the Barabási-Albert (1999) preferential attachment model. In the economics literature, Atalay et al (2011) combine the random and preferential attachment algorithms to model the buyer-supplier network in the US economy.

¹⁰Aumann and Myerson (1988) and Myerson (1991) model network formation as extensive-form and simultaneous move games respectively. Jackson and Wolinsky (1996) adopt a cooperative game theoretic approach, while Kranton and Minehart (2001) study buyer-seller networks in which ascending-bid auctions are used to determine the formation of links.

of a trading relationship conditional on having the opportunity to do so is an endogenous outcome. This hybrid approach is similar in spirit to the dynamic network formation models in Bala and Goyal (2000), Watts (2001), and Jackson and Watts (2002), but within the context of a structural model of trade between heterogeneous producers that can be used for quantitative analysis.¹¹

Finally, this paper contributes to several other areas of research. In studying quantitatively how firm-level relationship stickiness affects the responses of aggregate trade to shocks across different time horizons, this paper adds to the already-vast literature on the dynamics of firm-level trade and the estimation of trade elasticities.¹² Although the concept of trade studied in this paper focuses on trade between firms and is not explicitly international in nature, the notion of relationship stickiness applies to firm-to-firm trade in general, whether goods cross national borders or not. Understanding the effects of these frictions on trade within a country therefore also adds to our understanding of their effects on trade between countries. This paper also contributes to the study of how microeconomic shocks translate into aggregate fluctuations. Gabaix (2011) and Acemoglu et al (2012) argue that the firm size distribution and the network structure of linkages between sectors matter for how idiosyncratic firm- and sector-level shocks translate into aggregate movements, but do not seek to explain what determines these characteristics of the economy in the first place. The model that I develop endogeneizes both the firm size distribution as well as the firm-level input-output architecture, and therefore can be used to study the two-way interaction between these characteristics and aggregate fluctuations.

The outline of this paper is as follows. In section 2, I describe the data and document empirical regularities in the US production network. In section 3, I develop a static version of the theoretical model, in which the set of buyer-supplier relationships is taken as given. I characterize how firm size, trade volumes, and household welfare depend on the existing production network, and show how to solve the market equilibrium of the model for any given network of relationships. In section 4, I then endogeneize the formation of linkages between firms in the economy by introducing a dynamic matching process between potential buyers and sellers. I examine in detail the steady-state of the model, and show how to construct theoretical counterparts to the empirical moments described in section 2. In section 5, I take the model to data and estimate its parameters via simulated method of moments. Section

¹¹Bala and Goyal (2000), Watts (2001), and Jackson and Watts (2002) also assume for tractability that agents are myopic in their decisions about which links to form, whereas firms in my model optimally select relationships given rational expectations about the future costs and benefits of each relationship.

¹²Recent work on firm-level trade dynamics includes papers by Costantini and Melitz (2007), Eaton, Eslava, Kugler, and Tybout (2007), Burstein and Melitz (2011), Impullitti et al (2013), and Alessandria, Choi, and Ruhl (2015). For examples of recent work on estimating trade elasticities, see Arkolakis et al (2012) and Simonovska and Waugh (2014a, 2014b).

6 uses these parameter estimates to quantitatively study the model’s predicted responses of trade volumes and welfare to counterfactual shocks. Finally, section 7 concludes.

2 Data and Empirical Regularities

2.1 Data

Before describing the theoretical model, I first present several stylized facts about production networks in the US economy. These empirical regularities are documented using two overlapping datasets. The first is obtained from Standard and Poor’s Capital IQ platform, which collects fundamental data on a large set of companies worldwide, covering over 99% of global market capitalization. For a subset of these firms, both public and private but located mostly in the US, the database also records supplier and customer relationships based on a variety of sources, such as publicly available financial forms, company reports, and press announcements. From this database, I select all firms in the continental US for which relationship data is available and average revenue from 2003-2007 is positive. This gives me a dataset comprising 8,592 firms with \$16.3 trillion in total revenue, comparable to the value of \$30.0 trillion in total non-farm US business revenue as reported in the Census Bureau’s 2007 survey of business owners. The Capital IQ platform also provides the headquarters address of the majority of firms in this sample, which I geocode to obtain estimates of a firm’s location. Using these estimated locations, I then compute estimated distances between every supplier-customer pair in the dataset. Figure 1 shows the Capital IQ network for illustration, where each circle (node) represents a firm and each line (edge) represents a trading relationship.




Figure 1: The network of firm-to-firm trade in the continental United States, Capital IQ dataset

The second dataset is based on information from the Compustat platform, which is also

operated by Standard and Poor. The Compustat database contains fundamental information for publicly-listed firms in the US, compiled solely from financial disclosure forms, and includes firms' own reports of who their major customers are. In accordance with Financial Accounting Standards No. 131, a major customer is defined as a firm that accounts for at least 10% of the reporting seller's revenue. The Compustat relationship data was processed and studied by Atalay et al (2011), from whom the dataset was obtained. It contains 103,379 firm-year observations from 1979 to 2007.

Both the Capital IQ and Compustat datasets have their advantages and disadvantages. The Capital IQ platform offers greater coverage of firms with relationship data, as the database includes both public and private firms and records relationships based on sources other than financial disclosure forms. However, the main drawback of the dataset is that it is not possible to tell whether a particular relationship reported in a given year is still active at a later date. The Compustat data, on the other hand, is in panel form and therefore allows one to track the creation and destruction of trading relationships across time. The main weakness of the Compustat data is the 10% truncation level, which implies that a firm cannot have more than 10 customers reported in a given year, although there is still substantial variation in the number of recorded suppliers a firm has. For these reasons, I treat the capital IQ data as cross-sectional and primarily use it to estimate the steady-state of the model. I use the Compustat data to measure dynamic moments that are also used in the estimation.

2.2 Empirical regularities

In what follows, I document several empirical regularities characterizing the production network between firms in the data sample. In section 5, a subset of these moments will be used to estimate the theoretical model by simulated method of moments, and it is therefore useful at this point to formalize notation. Denoting the set of firms by S , I first define N_{bin} evenly-spaced quantile bins $\{B_b\}_{b \in \{1, \dots, N_{bin}\}}$, where:

$$B_b = \begin{cases} [q_{b-1}, q_b) & , b \in \{1, \dots, N_{bin} - 1\} \\ [q_{b-1}, q_b] & , b = N_{bin} \end{cases} \quad (2.1)$$

with $q_b \equiv \frac{b}{N_{bin}}$, and define $\bar{q}_b \equiv \frac{q_{b-1} + q_b}{2}$ as the midpoint of bin b . I then compute for each variable of interest X the quantile of this variable for firm s , $q^X(s)$, and define $b^X(s) \equiv \{b | q^X(s) \in B_b\}$ as the quantile bin of variable X for firm s . Finally, I define $S_b^X \equiv \{s \in S | b^X(s) = b\}$ as the set of firms for which variable X falls in quantile bin b .

2.2.1 Firm-level distributions

I begin by documenting the high degree of firm heterogeneity along several dimensions. To do so, I first compute for each variable X the Kaplan-Meier estimate of the cumulative distribution function of the normalized variable:¹³

$$\tilde{X}(s) \equiv \frac{\tilde{X}(s) - \min_{s' \in S} X(s')}{\max_{s' \in S} X(s') - \min_{s' \in S} X(s')} \quad (2.2)$$

I then evaluate the inverse empirical CDF at the points $\{\bar{q}_b\}_{b \in \{1, \dots, N_{bin}\}}$ via linear interpolation, obtaining estimates of the quantile levels $\{\bar{X}_b\}_{b \in \{1, \dots, N_{bin}\}}$ for each quantile bin. Figure 2 shows these moments for the distributions of log revenue, log employment, in-degree (number of suppliers), and out-degree (number of customers) across all firms in the Capital IQ dataset.

To gain some sense about the parametric form of the distributions, I first compare the revenue and employment distributions to log-normal distributions with the same mean and variance by Monte Carlo simulation. As can be seen from the graphs, the distributions are relatively well-modeled by log-normal distributions, as is a common finding in the literature on firm size distributions.¹⁴ The lognormal approximation slightly overstates the fraction of firms with revenue below a given amount, however, and does the opposite for the firm employment distribution.

Next, to characterize the firm-level degree distributions, I compare these to two distributions that play central roles in network theory. It is well-known that in random graph models, where links form between nodes with a constant probability, the degree distribution is approximately Poisson. On the other hand, in preferential attachment graph models, where nodes with a greater number of links form new links with a greater probability, the degree distribution exhibits a power law. I therefore compare the degree distributions to Poisson and Pareto distributions.¹⁵ From this, we see that the Poisson distribution is a poor approximation to the empirical degree distributions, strongly suggesting that relationships between firms are far from random, as might be expected. The Pareto distribution is a somewhat better approximation, although the approximation is also far from perfect.

¹³This normalization is employed so as to make computed moments of the univariate firm-level distributions scale-invariant, and therefore directly comparable to corresponding moments in the theoretical model.

¹⁴See for example Cabral and Mata (2003) and Rossi-Hansberg and Wright (2007).

¹⁵The Poisson parameter is chosen to match the mean of the empirical distribution, while the tail index of the Pareto distribution is computed using the Hill estimation procedure and the lower bound is set to match the mean of the empirical distribution.

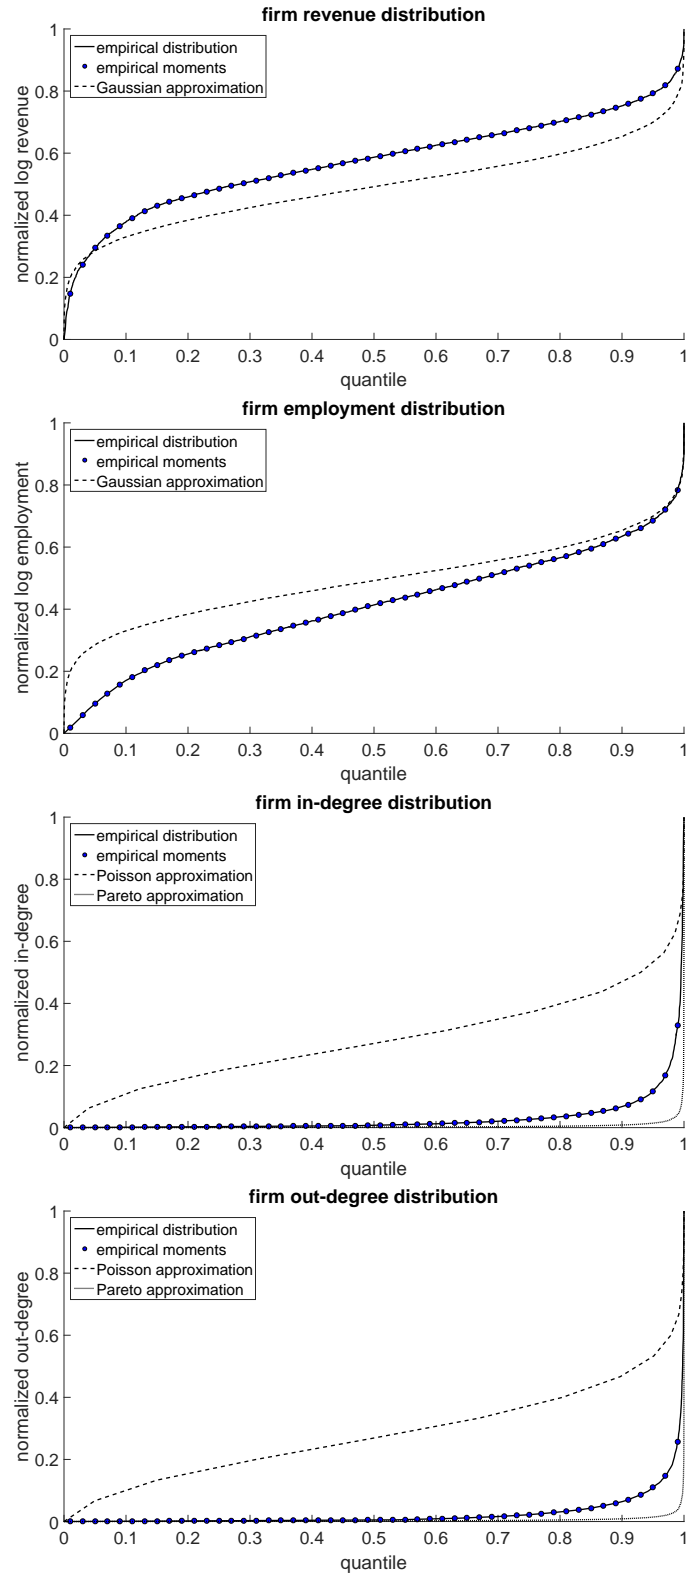


Figure 2: Firm-level distributions

2.2.2 Bivariate distributions

Next, I study how firm-level variables vary with firm size. Toward this end, I compute:

$${}^R\bar{Q}_b^X \equiv \frac{1}{|S_b^R|} \sum_{s \in S_b^R} q^X(s) \quad (2.3)$$

as the average quantile of variable X for all firms with revenue falling in quantile bin b . These moments are displayed in Figure 3 for employment, in-degree, and out-degree for all firms in the Capital IQ dataset. As expected, firm revenue and employment are highly correlated, but it is also clear from the graphs that larger firms tend to have larger numbers of customers and suppliers on average, with the rate of increase in degree also increasing in firm size. Firm-level variation in the numbers of suppliers and customers as well as the covariance of these measures with firm size will speak to the magnitude of the static aspect of relationship stickiness in the theoretical model.

2.2.3 Matching distributions

Having characterized both the distributions and correlations of revenue, employment, in-degree, and out-degree across firms, I now examine what kinds of firms match up with what kinds of firms in the network. In particular, I study how matching between firms varies with firm size by first computing $\bar{q}^{S,X}(s)$ and $\bar{q}^{C,X}(s)$ as the quantile of the mean level of variable X amongst suppliers and customers respectively of firm s (conditional on firm s having positive in- or out-degree). Next, as in section (2.2.2), I compute the averages of these firm-level measures within each revenue quantile bin:

$${}^R\bar{Q}_b^{S,X} \equiv \frac{1}{|S_b^R|} \sum_{s \in S_b^R} \bar{q}^{S,X}(s) \quad (2.4)$$

$${}^R\bar{Q}_b^{C,X} \equiv \frac{1}{|S_b^R|} \sum_{s \in S_b^R} \bar{q}^{C,X}(s) \quad (2.5)$$

Figure 5 shows these moments for supplier and customer revenue, employment, in-degree, and out-degree, for all firms in the Capital IQ dataset. From these graphs, we see that the assortativity of matching between firms is unambiguously positive, whether measured in terms of firm size or connectivity. On average, larger firms tend to buy and sell from firms that are also larger and better connected. This finding stands in contrast with the report of negative assortative matching in Bernard et al (2015) between exporting Norwegian firms and their trade partners, but agrees with the finding of Sugita et al (2014) that matching

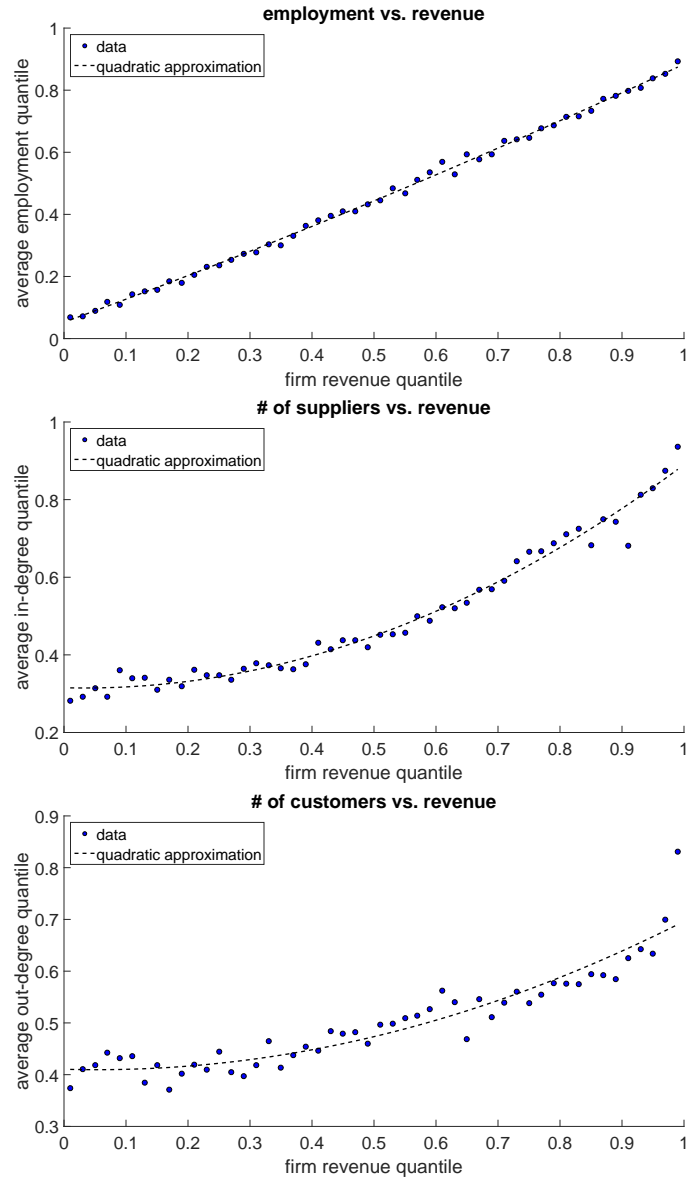


Figure 3: Bivariate distributions

assortativity is positive between textile firms in Mexico selling to firms in the US. These patterns of firm matching will be important in identifying the shape of the distribution of relationship fixed cost shocks in the theoretical model.

2.2.4 Relationship geography

In addition to characterizing the assortativity of firm matching, the geocoded locations of firms in the Capital IQ dataset allow me to examine the geographic distribution of a firm's suppliers and customers. To do so, I first compute $D_S(s)$ and $D_C(s)$ as the average distance between firm s and its suppliers and customers respectively, normalized by the maximum trading distance in the Capital IQ dataset.¹⁶ I then compute:

$$\bar{D}_b^S \equiv \frac{1}{|S_b^R|} \sum_{s \in S_b^R} D_S(s) \quad (2.6)$$

$$\bar{D}_b^C \equiv \frac{1}{|S_b^R|} \sum_{s \in S_b^R} D_C(s) \quad (2.7)$$

as the averages of the supplier and customer distance measures respectively for all firms with revenue falling in quantile bin b . These moments are shown in Figure 5. Perhaps somewhat surprisingly, larger firms tend to sell to customers that are located nearer by, while average supplier distance does not appear to vary much with firm size.¹⁷

2.2.5 Relationship dynamics

Finally, I make use of the panel nature of the Compustat data to study the dynamics of firm-to-firm relationships, which will be used to infer the magnitude of the dynamic aspect of relationship stickiness in the theoretical model. In particular, I examine how the rates at which firms retain existing suppliers and customers vary with firm size. To address this, I first compute for every firm s that exists in the dataset in both periods $t - 1$ and t the variables $\rho_t^{S,ret}(s)$ and $\rho_t^{C,ret}(s)$, which denote the fraction of that firm's suppliers and customers at date $t - 1$ respectively that are retained in period t . I then compute the following cross-

¹⁶The maximum distance is 4,415 kilometers, which is approximately equal to the horizontal width of the continental United States. Again, this normalization is employed so as to make empirical moments directly comparable to the simulated moments in the theoretical model.

¹⁷This finding is surprising in the context of trade models featuring fixed costs of exporting, for example, since these models predict that larger firms are more likely to sell to customers in more distant locations. On the other hand, it is perhaps less surprising in the context of models featuring agglomeration effects.

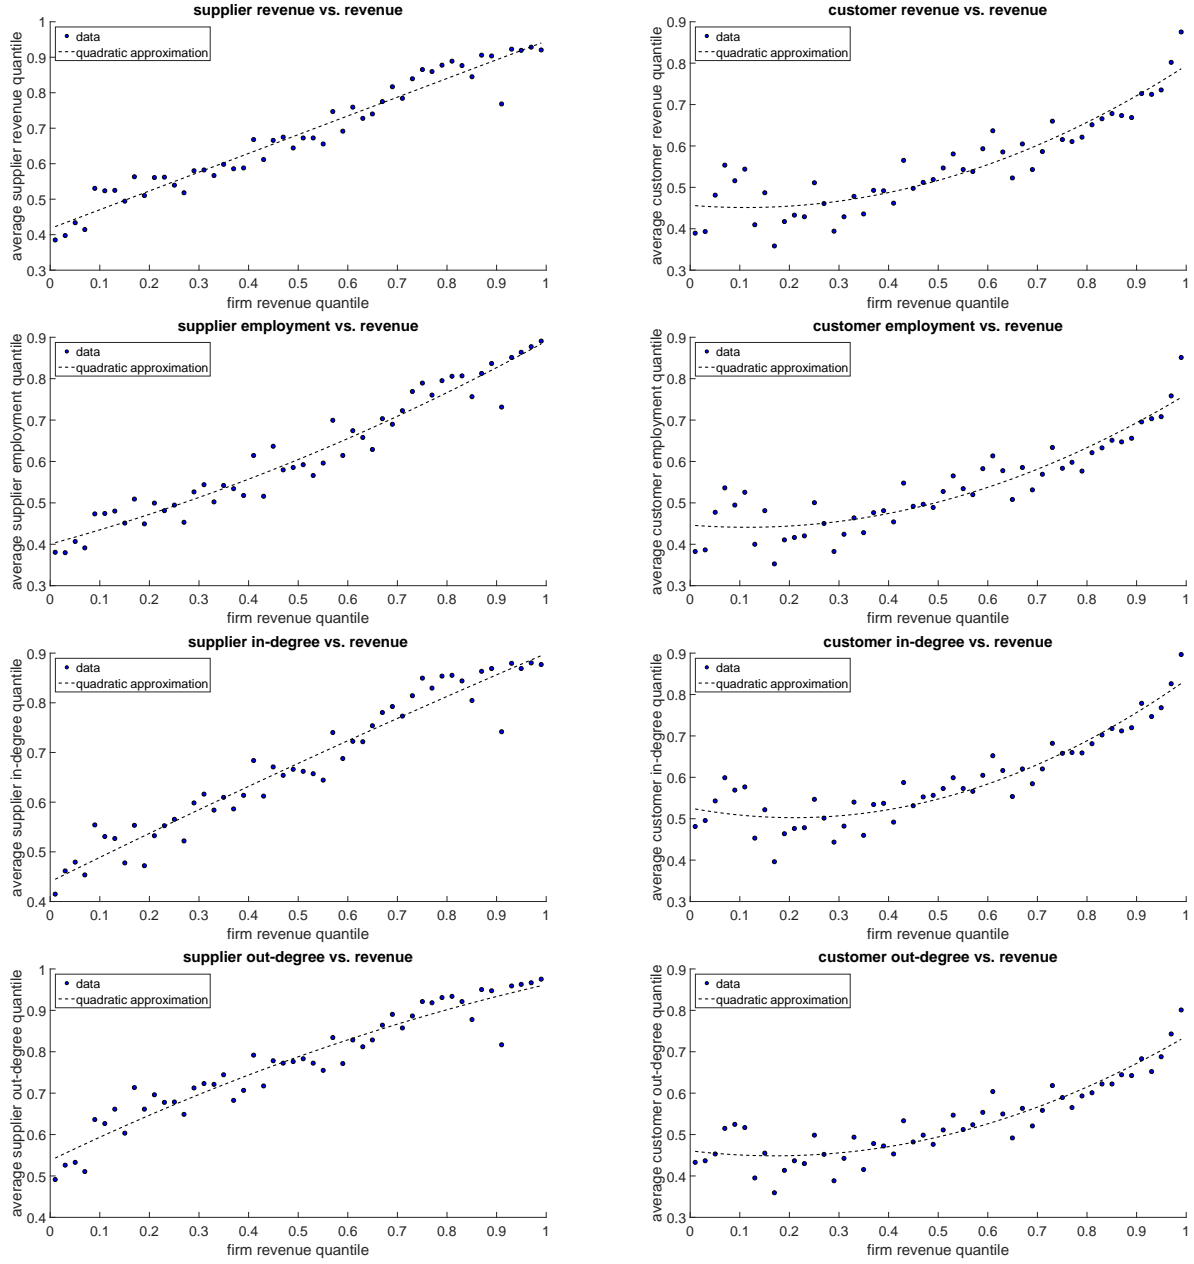


Figure 4: Matching distributions

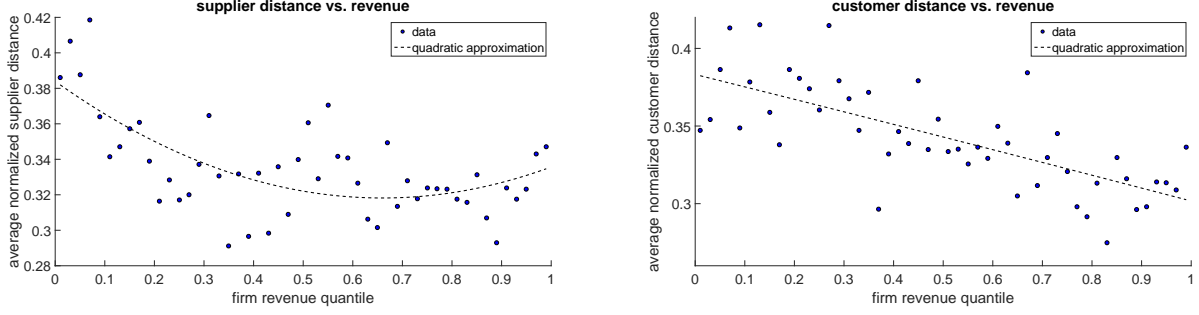


Figure 5: Geographic distribution of suppliers and customers

sectional averages:

$$\bar{\rho}_{b,t}^{S,ret} \equiv \frac{1}{|S_{b,t}^R|} \sum_{s \in S_{b,t}^R} \rho_t^{S,ret}(s) \quad (2.8)$$

$$\bar{\rho}_{b,t}^{C,ret} \equiv \frac{1}{|S_{b,t}^R|} \sum_{s \in S_{b,t}^R} \rho_t^{C,ret}(s) \quad (2.9)$$

where $S_{b,t}^R$ denotes the set of firms in revenue quantile bin b at date t (relative to the cross-sectional revenue distribution at that date). Finally, I compute the time-series averages of these moments across time:

$$\bar{\rho}_b^{S,ret} \equiv \frac{1}{T} \sum_{t=1}^T \bar{\rho}_{b,t}^{S,ret} \quad (2.10)$$

$$\bar{\rho}_b^{C,ret} \equiv \frac{1}{T} \sum_{t=1}^T \bar{\rho}_{b,t}^{C,ret} \quad (2.11)$$

where $T = 29$ is the number of years in the Compustat dataset.

These moments are shown in Figure 6. From these graphs, we see that larger firms tend to retain a larger fraction of both existing suppliers and customers, and by implication, the average duration of relationships is longer for relationships involving larger firms. The mean duration of trading relationships across all firms in the Compustat dataset is 1.74 years, and the average rate at which suppliers and customers are terminated year-to-year are 38.4% and 30.1% respectively.

2.2.6 Summary of stylized facts

In sum, the production network between firms in the data sample can be characterized by the following stylized facts:




Figure 6: Dynamic distributions

1. The firm size distribution is approximately log-normal, and the degree distributions deviate from both the Poisson and Pareto distributions predicted by statistical network formation models.
2. Larger firms tend to have more suppliers and customers.
3. The assortativity of matching between firms in terms of revenue, employment, and degree is unambiguously positive.
4. Larger firms tend to buy from and sell to suppliers and customers that are located nearer by.
5. Larger firms retain a larger fraction of suppliers and customers year-to-year.

Having documented these empirical regularities, I now turn to development of a simple model of trade between heterogeneous firms featuring sticky trading relationships, in which the firm-level degree distributions and matching between firms are endogenous outcomes. I return to the data in section 5 when I make use of the moments described above to estimate the model.

3 Static Model

I begin by describing a static version of the model in which the network of trading relationships between firms is fixed, and show how to characterize and solve for the static equilibrium conditional on the network. Having done so, I then focus attention on endogenizing dynamic formation of the production network in section 4.

3.1 Basic environment

The economy consists of a representative household and an exogenously-given unit continuum of heterogeneous firms that each produce a unique variety of a differentiated product. Firms are heterogeneous over states $\chi = (\phi, \delta)$, where ϕ and δ are what I refer to as the *fundamental productivity* of a firm's production process and the *fundamental quality* of a firm's product respectively, to be defined below. The exogenous cumulative distribution function over firm states is denoted by F_χ , with density f_χ and support S_χ a bounded subset of \mathbb{R}_+^2 . For brevity, I also refer to firms with state χ as χ -firms. I begin by studying a simplified version of the model in which all firms belong to a single location. In section 3.3, I show how it is straightforward to incorporate multiple locations into the model, and in particular I embed geography which will allow the model to speak to the geographic distribution of firm-to-firm trade discussed in section 2.2.4.

3.1.1 Households

The representative household supplies L units of labor inelastically and has CES preferences over all varieties of the differentiated product, given by:

$$U = \left[\int_{S_\chi} [\delta x_H(\chi)]^{\frac{\sigma-1}{\sigma}} dF_\chi(\chi) \right]^{\frac{\sigma}{\sigma-1}} \quad (3.1)$$

where σ is the elasticity of substitution across varieties, and $x_H(\chi)$ is the household's consumption of χ -firm varieties.¹⁸ Given the price $p_H(\chi)$ charged by χ -firms to the household, household demand is given by:

$$x_H(\chi) = \Delta_H \delta^{\sigma-1} [p_H(\chi)]^{-\sigma} \quad (3.2)$$

Note that conditional on prices, households demand a greater amount of varieties for which fundamental quality δ is higher. As opposed to buyer-seller specific components of quality, I assume here that δ is a characteristic of the firm that is common across all customers. The household's demand shifter can then be written as:

$$\Delta_H \equiv U P_H^\sigma \quad (3.3)$$

¹⁸Note that given the assumed unit mass of firms, integrals of all firm-level variables over the distribution F_χ are equal to both the average as well as the total value of that variable across firms.

and the consumer price index is equal to:

$$P_H = \left[\int_{S_\chi} \left[\frac{p_H(\chi)}{\delta} \right]^{1-\sigma} dF_\chi(\chi) \right]^{\frac{1}{1-\sigma}} \quad (3.4)$$

3.1.2 Firm production technology

Each firm produces its variety of the differentiated product using labor and the output of other firms. I assume, however, that firm-to-firm trade is characterized by relationship frictions, such that every χ -firm is only able to purchase inputs from a given χ' -firm with probability $m(\chi, \chi')$. Given that there exists a continuum of firms of every state, this implies that $m(\chi, \chi')$ is also equal to the fraction of χ' -firms that supply a given χ -firm, as well as the fraction of χ -firms that purchase from a given χ' -firm. I refer to m as the *matching function* of the economy, which completely specifies the extensive margin of firm-to-firm trading relationships in the economy. I take m as given in this section, and endogenize formation of firm-to-firm trading relationships once dynamics are introduced into the model in section 4.

Given the matching function, the output of a χ -firm is given by the following constant returns to scale CES production function:

$$X(\chi) = \left[[\phi l(\chi)]^{\frac{\sigma-1}{\sigma}} + \int_{S_\chi} m(\chi, \chi') [\alpha x(\chi, \chi')]^{\frac{\sigma-1}{\sigma}} dF_\chi(\chi') \right]^{\frac{\sigma}{\sigma-1}} \quad (3.5)$$

where $l(\chi)$ is the quantity of labor demanded and $x(\chi, \chi')$ is the quantity of each χ' -variety used as inputs.¹⁹ The parameter α is a measure of input-suitability, which I take as constant across firm pairs for now. Once I introduce geography into the model in section 3.3, α will be a natural means of incorporating trade costs across firms in different locations.²⁰ As is standard in the literature, I assume that the elasticity of substitution across inputs for intermediate demand is the same as that for final demand.

Taking the wage as the numeraire and given prices $\{p(\chi, \chi')\}_{\chi' \in S_\chi}$ charged by other firms, the marginal cost of each χ -firm is therefore given by:

$$\eta(\chi) = \left[\phi^{\sigma-1} + \alpha^{\sigma-1} \int_{S_\chi} m(\chi, \chi') [p(\chi, \chi')]^{1-\sigma} dF_\chi(\chi') \right]^{\frac{1}{1-\sigma}} \quad (3.6)$$

¹⁹In the appendix, I show how the model is isomorphic to one in which firms face convex input costs rather than a production function exhibiting “love of variety”.

²⁰In section C.1 of the appendix, I also discuss how α can be used to capture differences in input suitability across industries and to match industry-level input-output shares, although I do not pursue this extension in the numerical analysis.

while the quantities of labor and intermediate inputs demanded are given respectively by:

$$l(\chi) = X(\chi) \eta(\chi)^\sigma \phi^{\sigma-1} \quad (3.7)$$

$$x(\chi, \chi') = X(\chi) \eta(\chi)^\sigma \alpha^{\sigma-1} p(\chi, \chi')^{-\sigma} \quad (3.8)$$

Note that conditional on prices, firms with greater fundamental productivity ϕ have lower marginal costs.

3.1.3 Relationship costs

It is evident from equation (3.6) that as long as prices are finite, access to additional suppliers always lowers the marginal cost of a firm, which follows from the CES property of the production function. Furthermore, since the production function exhibits constant returns to scale, access to additional customers always increases a firm's variable profit. These forces generate incentives for firms to form as many upstream and downstream trading relationships as possible. To allow for the endogenous selection of relationships in the dynamic model studied in section 4, I therefore impose a cost of forming relationships by assuming that a link between any two firms requires f units of labor. This can be interpreted as the cost of resources needed to manage ongoing relationships, such as expenditures on customer and supplier management systems as alluded to in the introduction to this paper or as more general man-hour costs.

In what follows, I further assume that this fixed relationship cost is paid fully by the selling firm. As we will see, this assumption implies that firm pricing decisions which are optimal in the static market equilibrium remain optimal in the dynamic market equilibrium, and that decisions about which relationships to keep active need to be analyzed only from the perspective of selling firms. In section C.2 of the appendix, I discuss how this assumption might be relaxed to allow for the buying firm to pay a positive share of the fixed relationship cost.

3.1.4 Market clearing

The labor market clearing condition can be written as:

$$\int_{S_\chi} l(\chi) dF_\chi(\chi) = L - L_f \quad (3.9)$$

where L_f is the total amount of labor used to pay the fixed relationship costs in the economy:

$$L_f = f \int_{S_\chi} \int_{S_\chi} m(\chi, \chi') dF_\chi(\chi) dF_\chi(\chi') \quad (3.10)$$

If we define the total mass of a χ -firm's suppliers and customers respectively as:

$$M_S(\chi) \equiv \int_{S_\chi} m(\chi, \chi') dF_\chi(\chi') \quad (3.11)$$

$$M_C(\chi) \equiv \int_{S_\chi} m(\chi', \chi) dF_\chi(\chi') \quad (3.12)$$

then total fixed labor costs can be written equivalently as $L_f = \int_{S_\chi} M_S(\chi) dF_\chi(\chi) = \int_{S_\chi} M_C(\chi) dF_\chi(\chi)$.

Since variable labor $l(\chi)$ must be non-negative, we see that labor market clearing can be satisfied for any arbitrary matching function $m : S_\chi \times S_\chi \rightarrow [0, 1]$, including the matching function $m(\chi, \chi') = 1$ for all $\chi, \chi' \in S_\chi$ specifying the complete network, if and only if the following assumption holds.

Assumption 1. *The fixed relationship cost f is less than the total labor supply L .*

Finally, market clearing for the output of a χ -firm requires:

$$X(\chi) = x_H(\chi) + \int_{S_\chi} m(\chi', \chi) x(\chi', \chi) dF_\chi(\chi') \quad (3.13)$$

3.1.5 Firm pricing and market structure

The market structure for all firm sales is assumed to be monopolistic competition. Given that the household and all purchasing firms face a continuum of sellers of every state and have demand functions (3.2) and (3.8) exhibiting a constant price elasticity, the profit-maximizing price charged by each firm is equal to the standard CES markup over marginal cost:

$$p_H(\chi) = \mu \eta(\chi) \quad (3.14)$$

$$p(\chi, \chi') = \mu \eta(\chi') \quad (3.15)$$

$$\mu = \frac{\sigma}{\sigma - 1} \quad (3.16)$$

As I discuss in section 4.1.3, the assumption that selling firms pay the entire share of the fixed relationship cost implies that the costly nature of relationships has no effect on the optimal price charged by firms. In section C.2 of the appendix, I also discuss how the model

might be enriched by allowing for a form of bargaining between buyers and sellers, so that the markups charged by firms remain constant but are not completely determined by the elasticity of substitution σ .

3.2 Static market equilibrium

3.2.1 Firm network characteristics

As described above, the parameters ϕ and δ capture exogenous productivity and quality characteristics that are fundamental to the firm, in the sense that they are independent of the firm's connection to other firms. Conditional on prices, firms with greater values of ϕ and δ enjoy lower marginal costs and greater final demand respectively. Firm-level outcomes in equilibrium, however, such as the overall size and profit of a firm, depend not only on a firm's fundamental characteristics but also on the characteristics of other firms that it is connected to in the production network. For an arbitrary matching function, a given firm-level outcome may therefore in principle be a function of very complicated moments of the production network, which would render solution of the model intractable.

Fortunately, however, we can rely on the structure of the CES production function specified in (3.5) to derive sufficient statistics at the firm level that will allow us to compute all variables of interest with minimal computational difficulty. In contrast with firm fundamental characteristics ϕ and δ , it is therefore useful to characterize the static market equilibrium of the model in terms of what I call a χ -firm's *network* productivity and quality, defined respectively by:

$$\Phi(\chi) \equiv \eta(\chi)^{1-\sigma} \quad (3.17)$$

$$\Delta(\chi) \equiv \frac{1}{\Delta_H} X(\chi) \eta(\chi)^\sigma \quad (3.18)$$

Note that $\Phi(\chi)$ is an inverse measure of a χ -firm's marginal cost, while $\Delta(\chi)$ is the demand shifter of a χ -firm in the intermediate demand function (3.8) relative to the household's demand shifter Δ_H .

In what sense do $\Phi(\chi)$ and $\Delta(\chi)$ capture the characteristics of a χ -firm in the production network as a whole, and how are these quantities determined? Combining the demand equations (3.2) and (3.8), the firm marginal cost equation (3.6), the goods market clearing condition (3.13), and the pricing conditions (3.14) and (3.15), we obtain the following system

of equations:

$$\Phi(\chi) = \phi^{\sigma-1} + \mu^{1-\sigma} \alpha^{\sigma-1} \int_{S_\chi} m(\chi, \chi') \Phi(\chi') dF_\chi(\chi') \quad (3.19)$$

$$\Delta(\chi) = \mu^{-\sigma} \delta^{\sigma-1} + \mu^{-\sigma} \alpha^{\sigma-1} \int_{S_\chi} m(\chi', \chi) \Delta(\chi') dF_\chi(\chi') \quad (3.20)$$

Given the matching function, (3.19) and (3.20) specify a pair of decoupled linear functional equations in $\Phi(\cdot)$ and $\Delta(\cdot)$ respectively, and show how a firm's network characteristics depend on both its fundamental characteristics as well as on the network characteristics of its suppliers and customers. Conditional on ϕ and δ , firms that are connected to firms with larger network productivities and qualities also have higher network productivities and qualities themselves.²¹

The following proposition shows that as long as input-suitability α is not too large relative to the markup μ , there exist unique solutions to the equations (3.19) and (3.20) for any matching function, and that starting from any arbitrary (but bounded) guesses for $\Phi(\cdot)$ and $\Delta(\cdot)$, iterating on (3.19) and (3.20) converges to these unique solutions with a known rate.²² The proof of Proposition 1, relegated to the appendix, entails showing that the functional equations (3.19) and (3.20) constitute contraction mappings with Lipschitz constants $\left(\frac{\alpha}{\mu}\right)^{\sigma-1}$ and $\frac{\alpha^{\sigma-1}}{\mu^\sigma}$ respectively.

Proposition 1. *Under assumption 2, there exist unique network productivity and quality functions $\Phi : S_\chi \rightarrow \mathbb{R}_+$ and $\Delta : S_\chi \rightarrow \mathbb{R}_+$ for any matching function $m : S_\chi \times S_\chi \rightarrow [0, 1]$. Furthermore, starting from any arbitrary functions $\tilde{\Phi} : S_\chi \rightarrow \mathbb{R}_+$ and $\tilde{\Delta} : S_\chi \rightarrow \mathbb{R}_+$, iteration on equations (3.19) and (3.20) converges to Φ and Δ at rates $\left(\frac{\alpha}{\mu}\right)^{\sigma-1}$ and $\frac{\alpha^{\sigma-1}}{\mu^\sigma}$ respectively.*

Assumption 2. *Input suitability α is less than the markup μ .*

Under assumption 2, we can also rewrite equations (3.19) and (3.20) to express the

²¹Note that Φ and Δ are conceptually similar to the measure of weighted average productivity in Melitz (2003), but in my model, these are measures at the firm-level on both the buyer and seller sides, and depend on the network structure specified by the matching function.

²²When assumption 2 is violated, it becomes feasible for a pair of firms that are connected to each other both as buyer and seller to use only each other's output as inputs for production, thereby generating infinite output and profits.

network productivity and quality of a χ -firm respectively as:

$$\Phi(\chi) = \int_{S_\chi} \left[\sum_{d=0}^{\infty} \left(\frac{\alpha}{\mu} \right)^{d(\sigma-1)} m^{(d)}(\chi, \chi') \right] (\phi')^{\sigma-1} dF_\chi(\chi') \quad (3.21)$$

$$\Delta(\chi) = \mu^{-\sigma} \int_{S_\chi} \left[\sum_{d=0}^{\infty} \left(\frac{\alpha^{\sigma-1}}{\mu^\sigma} \right)^d m^{(d)}(\chi', \chi) \right] (\delta')^{\sigma-1} dF_\chi(\chi') \quad (3.22)$$

where $m^{(d)}$ is the d^{th} -degree matching function, defined recursively by:

$$m^{(0)}(\chi, \chi') = \begin{cases} \frac{1}{f_\chi(\chi)}, & \text{if } \chi = \chi' \\ 0, & \text{if } \chi \neq \chi' \end{cases} \quad (3.23)$$

$$m^{(1)}(\chi, \chi') = m(\chi, \chi') \quad (3.24)$$

$$m^{(d)}(\chi, \chi') = \int_{S_\chi} m_i^{(d-1)}(\chi, \chi'') m(\chi'', \chi') dF_\chi(\chi'') \quad (3.25)$$

Intuitively, one can think of $m^{(d)}(\chi, \chi')$ for $d \geq 1$ as the probability that a χ -firm buys indirectly from a χ' -firm through a supply chain that is of length d . With this interpretation, equations (3.21) and (3.22) show how the network productivity and quality of a firm depend on its connections to all other firms via supply chains of all lengths. Note that the rate at which the value of an indirect relationship decays with the length of the supply chain is decreasing in input suitability α and increasing in the markup μ .

3.2.2 Firm size and inter-firm trade

Once the fundamental and network characteristics of a firm are known, the total revenue, variable profit, and variable employment of a χ -firm are completely determined up to the scale factor Δ_H , and are given respectively by:

$$R(\chi) = \mu \Delta_H \Delta(\chi) \Phi(\chi) \quad (3.26)$$

$$\pi(\chi) = (\mu - 1) \Delta_H \Delta(\chi) \Phi(\chi) \quad (3.27)$$

$$l(\chi) = \Delta_H \Delta(\chi) \phi^{\sigma-1} \quad (3.28)$$

Intuitively, if a firm is twice as productive and produces a product that is twice as good from the perspective of the entire networked economy, its revenue and profit gross of fixed

relationship costs quadruples. Total firm profit and employment are given by:

$$\Pi(\chi) = \pi(\chi) - fM_C(\chi) \quad (3.29)$$

$$L(\chi) = l(\chi) + fM_C(\chi) \quad (3.30)$$

Total output of a χ -firm is also completely determined by firm fundamental and network characteristics up to a scale factor:

$$X(\chi) = \Delta_H \Delta(\chi) \Phi(\chi)^{\frac{\sigma}{\sigma-1}} \quad (3.31)$$

as are the value and quantity of output traded from χ' to χ -firms:

$$r(\chi, \chi') = \left(\frac{\alpha}{\mu}\right)^{\sigma-1} \Delta_H \Delta(\chi) \Phi(\chi') \quad (3.32)$$

$$x(\chi, \chi') = \frac{\alpha^{\sigma-1}}{\mu^\sigma} \Delta_H \Delta(\chi) \Phi(\chi')^{\frac{\sigma}{\sigma-1}} \quad (3.33)$$

3.2.3 Household welfare and demand

To complete characterization of the static market equilibrium, it remains to determine the scale factor Δ_H . From the labor market clearing condition (3.9) and the firm variable employment equation (3.28), this is given by:

$$\Delta_H = \frac{L - L_f}{\int_{S_\chi} \Delta(\chi) \phi^{\sigma-1} dF_\chi(\chi)} \quad (3.34)$$

Equations (3.3) and (3.4) then give the CPI and household welfare respectively as:

$$P_H = \mu \left[\int_{S_\chi} \Phi(\chi) \delta^{\sigma-1} dF_\chi(\chi) \right]^{\frac{1}{1-\sigma}} \quad (3.35)$$

$$U = \mu^{-\sigma} (L - L_f) \frac{\left[\int_{S_\chi} \Phi(\chi) \delta^{\sigma-1} dF_\chi(\chi) \right]^{\frac{\sigma}{\sigma-1}}}{\int_{S_\chi} \Delta(\chi) \phi^{\sigma-1} dF_\chi(\chi)} \quad (3.36)$$

while household demand is given by:

$$x_H(\chi) = \mu^{-\sigma} \Delta_H \delta^{\sigma-1} \Phi(\chi)^{\frac{\sigma}{\sigma-1}} \quad (3.37)$$

Using equations (3.21) and (3.22) to substitute for $\Phi(\chi)$ and $\Delta(\chi)$ respectively, we see that the numerator and denominator of (3.36) are identical except for the terms $\left(\frac{\alpha}{\mu}\right)^{d(\sigma-1)}$

and $\left(\frac{\alpha^{\sigma-1}}{\mu^\sigma}\right)^d$, with the difference going to zero exponentially as d increases. An intuitive approximation to the value of household welfare is therefore:

$$U \approx (L - L_f) \left[\int_{S_\chi} \int_{S_\chi} \left[\sum_{d=0}^{\infty} \left(\frac{\alpha}{\mu}\right)^{d(\sigma-1)} m^{(d)}(\chi, \chi') \right] (\delta \phi')^{\sigma-1} dF_\chi(\chi) dF_\chi(\chi') \right]^{\frac{1}{\sigma-1}} \quad (3.38)$$

which is exact in the limit as $\mu \rightarrow 1$ (perfect competition). Equation (3.38) suggests that household welfare is greater when buyers of greater fundamental quality δ are connected with sellers of greater fundamental productivity ϕ' , with the cost to welfare of additional relationships appearing in the term $L - L_f$. When $\mu > 1$, the same general intuition applies, although household utility is only given exactly by the slightly more complicated expression (3.36).

3.2.4 Static market equilibrium definition

Given the matching function m , the exogenous distribution over fundamental firm characteristics F_χ , and the model parameters $\{L, \sigma, \alpha, f\}$, we can now define a static market equilibrium of the economy as follows. In section A.1 of the appendix, I describe the computational algorithm used to solve for the static market equilibrium.

Definition 1. A static market equilibrium of the economy is a pair of firm network characteristic functions $\Phi : S_\chi \rightarrow \mathbb{R}_+$ and $\Delta : S_\chi \rightarrow \mathbb{R}_+$ satisfying equations (3.19) and (3.20), a scalar household demand shifter Δ_H satisfying (3.34), and allocation functions $\{l(\cdot), X(\cdot), x(\cdot, \cdot), x_H(\cdot)\}$ given respectively as side equations by (3.28), (3.31), (3.33), and (3.37).

3.2.5 Static market equilibrium efficiency

To characterize the efficiency of a static market equilibrium, we can compare the resulting allocation with the allocation that would be chosen by a social planner seeking to maximize household welfare subject to the same exogenous matching function, production technology, and resource constraints. The following proposition (proved in section B.1 of the appendix) summarizes the solution to the planner's problem.

Proposition 2. *Given a matching function $m : S_\chi \times S_\chi \rightarrow [0, 1]$, the network characteristic*

functions under the social planner's allocation satisfy:

$$\Phi^{SP}(\chi) = \phi^{\sigma-1} + \alpha^{\sigma-1} \int_{S_\chi} m(\chi, \chi') \Phi^{SP}(\chi') dF_\chi(\chi') \quad (3.39)$$

$$\Delta^{SP}(\chi) = \delta^{\sigma-1} + \alpha^{\sigma-1} \int_{S_\chi} m(\chi', \chi) \Delta^{SP}(\chi') dF_\chi(\chi') \quad (3.40)$$

and the allocations of output and labor are given by equations (3.28), (3.31), (3.33), and (3.37) with μ set equal to 1.

This result shows that any static market equilibrium allocation coincides with the corresponding planner's allocation if and only if all firms in the decentralized equilibrium are perfectly competitive. With monopolistically-competitive firms, the static market equilibrium allocation is therefore inefficient relative to the planner's allocation. This result can be interpreted as implying that the introduction of relationship frictions into the model through the exogenous matching function m imposes no additional inefficiency beyond the standard monopoly markup distortion. Once the matching function is endogeneized in section 4, this will no longer be true, as firm's decisions about which relationships to keep active generate an additional dynamic source of inefficiency.

3.3 Embedding geography

Before introducing dynamics and endogeneizing the formation of firm-to-firm trading relationships, it is useful to first describe how geography can be embedded into the model to study how relationship stickiness affects trade patterns across different locations, as this will be one area of focus of the numerical analysis and counterfactuals in sections 5 and 6. Toward this end, I assume that the unit mass of firms is evenly distributed along a unit circle, with each point on the circle indicating a different location. The distribution over firm states F_χ is assumed to be identical in all locations, and we can therefore focus on characterizing the market equilibrium in a single location.

To model trade costs, I assume that trade between two locations separated by a distance D along the unit circle is subject to iceberg trade costs equal to $\tau(D) \geq 1$, with $\tau(0) = 1$, $\tau'(D) > 0$, and τ log-subadditive.²³ Since all locations are identical, we can assume for notational simplicity and without loss of generality that firms in any one location can only sell to locations located clockwise of their own location. Given these assumptions, the static market equilibrium with geography embedded is simply characterized by analogous equations

²³That is, $\log \tau(D_1) + \log \tau(D_2) \geq \log \tau(D_1 + D_2)$ for any $D_1, D_2 \in [0, 1]$, which is equivalent to the assumption that trade costs satisfy the triangle inequality.

for the network productivity and quality functions:

$$\Phi(\chi) = \phi^{\sigma-1} + \left(\frac{\alpha}{\mu}\right)^{\sigma-1} \int_0^1 \int_{S_\chi} \tau(D)^{1-\sigma} m[\chi, \chi' | \tau(D)] \Phi(\chi') dF_\chi(\chi') dD \quad (3.41)$$

$$\Delta(\chi) = \mu^{-\sigma} \delta^{\sigma-1} + \mu^{-\sigma} \alpha^{\sigma-1} \int_0^1 \int_{S_\chi} \tau(D)^{-\sigma} m[\chi', \chi | \tau(D)] \Delta(\chi') dF_\chi(\chi') dD \quad (3.42)$$

where the matching function is now allowed to depend on distance through the trade cost $\tau(D)$.²⁴

As in the model without geography, there exist unique solutions to equations (3.41) and (3.42) for the functions Φ and Δ . Given these, the value of trade between a χ -buyer and a χ' -seller separated by a distance D is then given by:

$$R(\chi, \chi' | D) = \left(\frac{\alpha}{\mu}\right)^{\sigma-1} \tau(D)^{1-\sigma} \Delta_H \Delta(\chi) \Phi(\chi') \quad (3.43)$$

Notice that equation (3.43) resembles a gravity equation for trade volumes at the firm level, where $\Delta_H \Delta(\cdot)$ and $\Phi(\cdot)$ capture the economic size of the importer and exporter respectively. The total value of trade between locations a distance D apart, however, also depends on the mass of firms that match between the two locations, and is given by:

$$\bar{R}(D) = \left(\frac{\alpha}{\mu}\right)^{\sigma-1} \tau(D)^{1-\sigma} \Delta_H \int_{S_\chi} \int_{S_{\chi'}} m[\chi, \chi' | \tau(D)] \Delta(\chi) \Phi(\chi') dF_\chi(\chi) dF_{\chi'}(\chi') \quad (3.44)$$

Observe that if the matching function is held fixed, then as in models of trade with CES roundabout production such as Melitz (2003), the elasticity of trade volumes with respect to trade costs depends only on the elasticity of substitution σ . However, once the matching function is endogenously determined as the result of firms' decisions to trade or not to trade with other firms in various locations, the response of trade volumes to trade costs also depends on the extent of relationship frictions between firms.

4 Dynamics and Endogenous Network Formation

Analysis of the static version of the model shows that given any arbitrary matching function m , numerical solution of all firm-level variables of interest is straightforward and

²⁴Note that by writing equation (3.42) in this way, we are implicitly assuming that the representative household in each location purchases goods only from firms in its own location. Making the alternative assumption that households also purchase directly from firms in other locations subject to the same trade costs would simply require multiplying the first term on the right-hand side of (3.42) by the term $\bar{\tau} \equiv \int_0^1 \tau(D)^{-\sigma} dD$, and would add nothing of qualitative substance to the model.

tractable. It is the matching function m , however, that captures all the relevant information determining the empirical moments in which we are interested, as described in section 2. Endogeneizing formation of the network is therefore crucial to my analysis, and I accomplish this by introducing a dynamic process of firm matching, as described below.

4.1 Dynamics of firm matching

Time is discrete and the representative household has preferences at date t defined by:

$$V_t = \sum_{s=t}^{\infty} \beta^{s-t} U_s \quad (4.1)$$

where U_t is given by the date t equivalent of (3.1). Since the household's value function is linear in per-period utility, household decisions every period are characterized exactly as in the static model, and the discount factor β exists only to characterize how firms (which are owned by the household) discount the future. To economize on notation, I first describe the dynamic model without geography embedded, and reintroduce geography once I conduct the numerical analysis and study counterfactuals. The dynamics of firm matching are modeled based on three main assumptions.

4.1.1 Random fixed relationship costs

First, I assume that the fixed relationship cost f_t is a random variable given by $f_t = f\xi_t$, where ξ_t is independent and identically distributed across firm pairs and time with cumulative distribution function F_ξ and unit mean. As in the static model, I assume that regardless of the realization of ξ_t , the selling firm always pays the full share of the fixed cost. The stochastic nature of the fixed relationship cost is the mechanism that generates the creation of new linkages between firms and the destruction of existing relationships, even in the steady-state of the model.

The assumption that ξ_t exhibits no serial correlation is made primarily for tractability, and might jar with one's intuition that relationship costs should be persistent. Nonetheless, the model generates non-trivial predictions about the persistence of relationships via assumptions about how often firms can reset relationships, described next.

4.1.2 Sticky relationships

I assume that firm-to-firm trading relationships are temporally sticky in the following sense. At each date, a firm receives the opportunity to sell to each firm that it did not

sell to in the previous period with probability $1 - \nu$, and also receives the opportunity to terminate trading relationships with each of its existing customers with probability $1 - \nu$. I refer to this as the *reset shock*, and assume that it is independent across all firm pairs. Although the model can easily accommodate differences in the probabilities with which a firm can create and destroy relationships, I assume for parsimony that these probabilities are the same. Furthermore, I assume that regardless of whether a reset shock is received, selling firms can costlessly adjust prices every period, so that firm-to-firm relationships are sticky only along the extensive margin.

The assumption that firms can only sell to new customers with a finite probability may be interpreted as modeling the fact that potential trading partners take time to meet and learn about the suitability of their output for each other's production processes or to negotiate new trading arrangements. Similarly, the assumption that firms cannot costlessly terminate existing relationships may be interpreted as either legal barriers to reneging on pre-negotiated contractual obligations, or more simply as the notion that winding down trading relationships also takes time. Allowing firms to costlessly adjust the intensive but not the extensive margin of trade may be interpreted as assuming that contracts between firms mandate only the provision of a good by the seller and not the price at which that good is sold.

Note that since the selling firm always pays the full share of the fixed relationship cost, the buying firm is always agreeable to any trading relationship, and therefore the decision to terminate or activate relationships only needs to be analyzed from the perspective of the selling firm. Under these assumptions, the matching function evolves according to the following law of motion:

$$m_t(\chi, \chi') = \nu m_{t-1}(\chi, \chi') + (1 - \nu) a_t(\chi, \chi') \quad (4.2)$$

where $a_t(\chi, \chi')$ is the probability that a χ' -firm sells to a χ -firm in period t conditional on being given the opportunity to reset that relationship. I refer to a_t as the acceptance function and characterize this in the following section. In any steady-state of the model, the matching function is simply equal to the acceptance function:

$$m(\chi, \chi') = a(\chi, \chi') \quad (4.3)$$

Note that f and ν capture respectively the static and dynamic aspects of relationship stickiness alluded to in the introduction of this paper. There are several qualitatively different cases that one can consider. First, in the absence of the dynamic friction ($\nu = 0$), the matching function converges immediately to its steady-state value of $a(\cdot, \cdot)$, and the short- and long-run elasticities of trade volumes with respect to aggregate shocks are equal. Second,

when the dynamic friction is extreme ($\nu = 1$), the production network exhibits no dynamics along the extensive margin. Third, in the presence of extreme static relationship costs ($f \rightarrow \infty$ and $\nu \in [0, 1)$), any steady-state of the model features an empty network in which no inter-firm trade occurs. Fourth, in the absence of the static friction ($f = 0$ and $\nu \in [0, 1)$), any steady-state of the model features a complete network in which all firms trade with one another. Trade therefore does not respond to external shocks along the extensive margin. Finally, with moderate static and dynamic frictions ($f \in (0, \infty)$ and $\nu \in (0, 1)$), the model exhibits both non-trivial steady-state production networks as well as non-trivial transition dynamics between steady-states.

4.1.3 Dynamic relationship activation decisions

The third and final assumption regards how and when firms decide to reset trading relationships conditional on having the opportunity to do so. First, note that the assumption that buying firms pay none of the fixed cost implies that it is never optimal for the selling firm to deviate from the standard CES markup pricing. Therefore, the variable profit earned by a χ' -firm from selling to a χ -firm at date t is the same as in the static market equilibrium, given by equations (3.20) and (3.27) as:

$$\pi_t(\chi, \chi') = \mu^{-\sigma} (\mu - 1) \alpha^{\sigma-1} \Delta_{H,t} \Delta_t(\chi) \Phi_t(\chi') \quad (4.4)$$

where $\Phi_t(\cdot)$, $\Delta_t(\cdot)$, and $\Delta_{H,t}$ are defined by the date t equivalents of equations (3.19), (3.20), and (3.34).

Now, let $V_t^+(\chi, \chi' | \xi_t)$ denote the value to a χ' -firm of selling to a χ -firm in period t conditional on the realization of the relationship cost shock ξ_t , and let $V_t^-(\chi, \chi')$ denote the value to the firm of not selling.²⁵ These value functions are given by the following Bellman equations:

$$V_t^+(\chi, \chi' | \xi_t) = \pi_t(\chi, \chi') - f\xi_t + \beta(1 - \nu) \mathbb{E}_t[V_{t+1}^O(\chi, \chi' | \xi_{t+1})] + \beta\nu \mathbb{E}_t[V_{t+1}^+(\chi, \chi' | \xi_{t+1})] \quad (4.5)$$

$$V_t^-(\chi, \chi') = \beta(1 - \nu) \mathbb{E}_t[V_{t+1}^O(\chi, \chi' | \xi_{t+1})] + \beta\nu V_{t+1}^-(\chi, \chi') \quad (4.6)$$

where $V_t^O(\chi, \chi' | \xi_t)$ denotes the value to a χ' -firm of having the option to reset its relationship with a χ -firm customer given the relationship cost shock ξ_t :

$$V_t^O(\chi, \chi' | \xi_t) = \max \left\{ V_t^+(\chi, \chi' | \xi_t), V_t^-(\chi, \chi') \right\} \quad (4.7)$$

²⁵Note that since the relationship cost shocks are i.i.d. over time, the value of not selling at date t does not depend on ξ_t . Furthermore, since there is no aggregate uncertainty in the model, this implies that there is no uncertainty over the value of V_t^- at any date for any pair of firms.

Note that the assumption of sticky relationships makes the activation and termination decisions facing a given firm forward-looking. If a firm chooses not to terminate a relationship given the chance to do so, it may find itself wishing to terminate the relationship in the future but lacking the opportunity to do so. Similarly, if a firm chooses not to sell to a potential customer despite having the chance to do so, it may be forced to wait several periods before being able to activate the relationship. Observe that if relationships are not sticky ($\nu = 0$) or firms are completely myopic ($\beta = 0$), then $V_t^+(\chi, \chi'|\xi_t) \geq V_t^-(\chi, \chi')$ if and only if $\pi_t(\chi, \chi') \geq f\xi_t$. In these two special cases, relationships are activated as long as the static profits accruing to selling firms are enough to cover the fixed relationship costs. The probability that a χ' -firm sells to a χ -firm at date t once it has the chance to do so is then given by:

$$\tilde{a}_t(\chi, \chi') = F_\xi \left[\frac{\pi_t(\chi, \chi')}{f} \right] \quad (4.8)$$

From (4.4), this implies that firms with larger network productivities and qualities are more likely to form downstream and upstream trading relationships respectively. The assumption of myopic agents in models of network formation is in fact somewhat standard in the network literature, and might seem to be a reasonable first approximation to firms' decision making processes.²⁶ We can, however, go further in characterizing the dynamic activation decisions of firms in this model.

It is instructive to first consider a steady-state of the model in which the functions π_t , V_t^+ , V_t^- , and V_t^O are all constant. From equations (4.5) and (4.6), it is straightforward to verify that:

$$\mathbb{E} \left[V^O(\chi, \chi'|\xi) \right] = \begin{cases} \frac{\pi(\chi, \chi') - f}{1 - \beta}, & \forall (\chi, \chi') \in S_+ \\ 0, & \forall (\chi, \chi') \notin S_+ \end{cases} \quad (4.9)$$

where $S_+ \equiv \{(\chi, \chi') \in S_\chi^2 | \pi(\chi, \chi') - f \geq 0\}$. This tells us that the option value of a relationship is positive if and only if the profit from that relationship exceeds the relationship cost on average. Substituting (4.9) into (4.5) and (4.6), we then find:

$$V^+(\chi, \chi'|\xi) - V^-(\chi, \chi') = \frac{\pi(\chi, \chi') - \beta\nu f}{1 - \beta\nu} - f\xi \quad (4.10)$$

and therefore the probability that a χ' -firm sells to a χ -firm conditional on having the chance

²⁶See for example Bala and Goyal (2000) and Jackson (2005).

to do so is given by:

$$a(\chi, \chi') = F_\xi \left[\frac{\pi(\chi, \chi') - \beta\nu f}{(1 - \beta\nu) f} \right] \quad (4.11)$$

Comparing this expression with equation (4.8), we again see that firms with greater network productivities and qualities are more likely to form downstream and upstream trading relationships respectively, but once the option values of relationships are taken into account, this effect becomes more pronounced. In particular, for firm pairs such that $\pi(\chi, \chi') > f$, there is a positive probability, equal to $a(\chi, \chi') - \tilde{a}(\chi, \chi')$, that temporarily-unprofitable relationships will still be activated because the relationship is profitable enough on average. Similarly, for firm pairs such that $\pi(\chi, \chi') < f$, there is a positive probability, given by $\tilde{a}(\chi, \chi') - a(\chi, \chi')$, that temporarily-profitable relationships will not be activated because the relationship is not profitable enough on average. Furthermore, note that (4.11) implies that firm pairs with $\pi(\chi, \chi') < \beta\nu f$ will never form trading relationships in steady-state.

How do we characterize the activation and termination decisions of firms outside the steady-state? Iterating forward on equations (4.5), (4.6), and (4.7), we can write the difference in the values of selling and not selling as:

$$V_t^+(\chi, \chi' | \xi_t) - V_t^-(\chi, \chi') = \pi_t(\chi, \chi') - f\xi_t + \sum_{s=1}^{\infty} (\beta\nu)^s [\pi_{t+s}(\chi, \chi') - f] \quad (4.12)$$

which can be interpreted as the expected future stream of profits net of fixed costs until the relationship can be reset. The acceptance function at date t is therefore given by:

$$a_t(\chi, \chi') = F_\xi \left[\frac{\pi_t(\chi, \chi')}{f} + \sum_{s=1}^{\infty} (\beta\nu)^s \left[\frac{\pi_{t+s}(\chi, \chi')}{f} - 1 \right] \right] \quad (4.13)$$

From this, we see that solving for the acceptance function at date t outside of the steady-state requires solving for the profit functions π_{t+s} for all $s \geq 1$. In section A.2 of the appendix, I describe the computational algorithm that I employ to accomplish this, which essentially involves iterating on the path of profit functions $\{\pi_{t+s}\}_{s=1}^T$ for some value of T large enough such that m_{t+T} is close to the eventual steady-state matching function. This allows me to solve exactly for the model's transition dynamics between steady-states under rational firm expectations. In section 6.4, I show why this is important, as the assumption of myopic firms leads to model predictions that are both qualitatively and quantitatively different from the rational expectations case.

Note that even though ξ_t is assumed to have unit mean, firms in the dynamic market equilibrium select relationships based on the realized values of the relationship cost shocks.

Therefore, the average cost of active relationships is no longer equal to f as it was in the static model, and the total mass of labor used to pay relationship fixed costs is now given by:

$$L_{f,t} = f \int_{S_\chi} \int_{S_\chi} \left[\nu m_{t-1}(\chi, \chi') + (1 - \nu) \bar{\xi}_t(\chi, \chi') \right] dF_\chi(\chi) dF_\chi(\chi') \quad (4.14)$$

The first term in the integral reflects the cost of relationships that cannot be reset (and hence for which there is no selection on ξ_t), while the second term reflects the cost of relationships that are voluntarily selected by firms. The term $\bar{\xi}_t(\chi, \chi')$ denotes the average value of the idiosyncratic component of the cost shock amongst $\chi - \chi'$ firm pairs that receive the reset shock:

$$\bar{\xi}_t(\chi, \chi') = \int_0^{\xi_{max,t}(\chi, \chi')} \xi dF_\xi(\xi) \quad (4.15)$$

and $\xi_{max,t}(\chi, \chi')$ is the maximum value of the cost shock for which $\chi - \chi'$ relationships are voluntarily selected:

$$\xi_{max,t}(\chi, \chi') = \max \left\{ \frac{\pi_t(\chi, \chi')}{f} + \sum_{s=1}^{\infty} (\beta\nu)^s \left[\frac{\pi_{t+s}(\chi, \chi')}{f} - 1 \right], 0 \right\} \quad (4.16)$$

4.2 Dynamic market equilibrium

4.2.1 Dynamic market equilibrium definition

Having characterized the dynamics of firm matching, we can now define a dynamic market equilibrium as follows.

Definition 2. Given an initial matching function $m_0 : S_\chi \times S_\chi \rightarrow [0, 1]$, a dynamic market equilibrium of the model is a list of sequences of matching functions $\{m_t\}_{t=1}^{\infty}$, acceptance functions $\{a_t\}_{t=0}^{\infty}$, profit functions $\{\pi_t\}_{t=0}^{\infty}$, and network characteristic functions $\{\Phi_t, \Delta_t\}_{t=0}^{\infty}$, as well as a list of scalars $\{\Delta_{Ht}\}_{t=0}^{\infty}$, all of which satisfy equations (3.19), (3.20), (3.34), (4.2), (4.4), and (4.13). Given the matching function m_t , the allocation at date t in a dynamic equilibrium is as defined in the static model.

Similarly, we can define a steady-state of the dynamic model as a dynamic market equilibrium in which all variables in Definition 2 are constant.

Definition 3. A steady-state equilibrium of the dynamic model is a matching function m , an acceptance function a , a profit function π , network characteristic functions $\{\Phi, \Delta\}$, as well as a scalar Δ_H , all of which satisfy equations (3.19), (3.20), (3.34), (4.3), (4.4), and (4.11).

Given the steady-state matching function m , the allocation in a steady-state equilibrium is as defined in the static model.

In section A.2 of the appendix, I describe the computational algorithms used to solve for both the model's transition dynamics as well as its steady-state.

4.2.2 Dynamic market equilibrium efficiency

To what extent are the dynamic relationship selection decisions made by firms socially optimal? Recall that the results of Proposition 2 showed how the static market equilibrium is inefficient relative to the social planner's allocation because of the monopoly markups charged by firms. Similarly, we can characterize the dynamic efficiency of the model by comparing the market equilibrium allocation with the dynamic allocation that would be chosen by a social planner subject to the same static and dynamic frictions faced by firms. In particular, we can compare the cutoff value for the relationship cost shock chosen by firms, given by equation (4.16), to the cutoff value that would be chosen by the planner. In section B.2 of the appendix, I show that the planner's solution is characterized by the following proposition.

Proposition 3. *The cutoff value for the cost shock at date t chosen by the social planner is given by:*

$$\xi_{max,t}^{SP}(\chi, \chi') = \max \left\{ \frac{\pi_t^{SP}(\chi, \chi')}{f} + \sum_{s=1}^{\infty} (\beta\nu)^s \left(\frac{\mathcal{C}_{t+s}}{\mathcal{C}_t} \right) \left[\frac{\pi_{t+s}^{SP}(\chi, \chi')}{f} - 1 \right], 0 \right\} \quad (4.17)$$

where π_t^{SP} is the planner's analog of the profit function:

$$\pi_t^{SP}(\chi, \chi') \equiv \frac{\alpha^{\sigma-1}}{\sigma-1} \Delta_{H,t}^{SP} \Delta_t^{SP}(\chi^*) \Phi_t^{SP}(\chi^{*'}) \quad (4.18)$$

and \mathcal{C}_t is a measure of the total connectivity between firms in the economy:

$$\mathcal{C}_t \equiv \left[\int_{S_\chi} \int_{S_\chi} \left[\sum_{d=0}^{\infty} \alpha^{d(\sigma-1)} m_t^{SP,(d)}(\chi, \chi') \right] (\delta\phi')^{\sigma-1} dF_\chi(\chi) dF_\chi(\chi') \right]^{\frac{1}{\sigma-1}} \quad (4.19)$$

Comparing equations (4.16) and (4.17), we now see that the criterion by which firms select relationships in the market equilibrium differs from the socially-optimal criterion in two ways. First, because of the monopoly markup distortion discussed in section 3.2.5, the static social value of a given relationship relative to its cost (measured by $\frac{\pi_t^{SP}}{f}$) differs from the ratio of profits to fixed costs ($\frac{\pi_t}{f}$) that are faced by selling firms in the market equilibrium.

Note that holding fixed the network productivity of the selling firm and the network quality of the buying firm, the function π_t^{SP} differs from the profit function π_t only by a constant fraction $\mu^{-\sigma}$.

Second, the planner internalizes the effect of each relationship on all other firms in the production network (often referred to as network externalities) whereas firms in the market equilibrium do not. To better understand this effect, it is useful to consider the social value of a given relationship at date t , which can be characterized by the static marginal change in household utility resulting from a marginal increase in the mass of active relationships between firms of given states. In the proof of Proposition 3, I show that this is given by:

$$\frac{dU_t}{d\bar{m}_t(\chi, \chi')} = \mathcal{C}_t \left[\pi_t^{SP}(\chi, \chi') - f \right] \quad (4.20)$$

where $\bar{m}_t(\chi, \chi') \equiv m_t(\chi, \chi') f_\chi(\chi) f_\chi(\chi')$ denotes the total mass of connections between χ -firm buyers and χ' -firm sellers. From equation (4.20), we see that the social value of each relationship is equal to the difference $\pi_t^{SP} - f$ amplified by the aggregate connectivity measure \mathcal{C}_t . Intuitively, when firms are more connected to each other (\mathcal{C}_t is larger), the activation or termination of a single relationship has larger aggregate effects. Since the amplification term \mathcal{C}_t potentially varies across time, the planner values changes in the extensive margin of firm relationships accordingly. This effect appears through the term $\frac{C_{t+s}}{C_t}$ in equation (4.17) but is absent in firms' decision making processes about which relationships to activate and terminate at each date.

4.3 Properties of the steady-state

4.3.1 Firm-level distributions

In our analysis of the static market equilibrium, we saw how the revenue and employment of a firm are completely determined (up to a scale factor) by the fundamental and network characteristics of that firm. I now show that variation in firm in-degrees (measured by M_S) and out-degrees (measured by M_C) is also completely determined by variation in network characteristics. To see this, first observe from equations (4.3), (4.4), and (4.11) that variations across firm-pairs in the profit, activation, and matching functions depend only on variations in the product $\Delta(\chi) \Phi(\chi')$.²⁷ In particular, the matching function in steady-state can be written as:

²⁷Given that each firm has a continuum of both suppliers and customers of each state, these functions do not depend on idiosyncratic realizations of the fixed cost shock ξ_t .

$$m(\chi, \chi') = \tilde{m} \left[\Delta_H \Delta(\chi) \Phi(\chi') \right] \quad (4.21)$$

where $\tilde{m} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is an increasing scalar function defined by:

$$\tilde{m}(x) = F_\xi \left[\frac{x - \beta\nu\bar{f}}{(1 - \beta\nu)\bar{f}} \right] \quad (4.22)$$

with $\bar{f} \equiv \frac{\mu^\sigma}{\mu-1} \alpha^{1-\sigma} f$. As a result, the network quality and productivity of a χ -firm are sufficient statistics for its in- and out-degrees respectively:

$$M_S(\chi) = \tilde{M}_S[\Delta(\chi)] \equiv \int_{S_\chi} \tilde{m} \left[\Delta(\chi) \Phi(\chi') \Delta_H \right] dF_\chi(\chi') \quad (4.23)$$

$$M_C(\chi) = \tilde{M}_C[\Phi(\chi)] \equiv \int_{S_\chi} \tilde{m} \left[\Delta(\chi') \Phi(\chi) \Delta_H \right] dF_\chi(\chi') \quad (4.24)$$

Since firm revenue is proportional to the product of firm network productivity and quality, this implies that firms with larger masses of suppliers and customers also tend to have larger revenue.

Figure 7 shows an example of the network productivity and quality functions in a steady-state of the model obtained through numerical solution, as well as the supplier and customer functions $M_S(\cdot)$ and $M_C(\cdot)$ defined by equations (3.11) and (3.12). Note that even though fundamental firm productivities and qualities ϕ and δ may be uncorrelated, a firm's network productivity $\Phi(\chi)$ is still increasing in δ because a firm with higher fundamental quality offers greater profit opportunities to potential suppliers, and therefore is more likely to form upstream trading relationships. Similarly, a firm's network quality $\Delta(\chi)$ is increasing in both its fundamental productivity and quality.

4.3.2 Matching assortativity

What determines the assortativity of matching between firms in the model? The average supplier and customer revenue of a χ -firm are given respectively by:

$$\bar{R}_S(\chi) = \frac{\int_{S_\chi} m(\chi, \chi') R(\chi') dF_\chi(\chi')}{M_S(\chi)} \quad (4.25)$$

$$\bar{R}_C(\chi) = \frac{\int_{S_\chi} m(\chi', \chi) R(\chi') dF_\chi(\chi')}{M_C(\chi)} \quad (4.26)$$

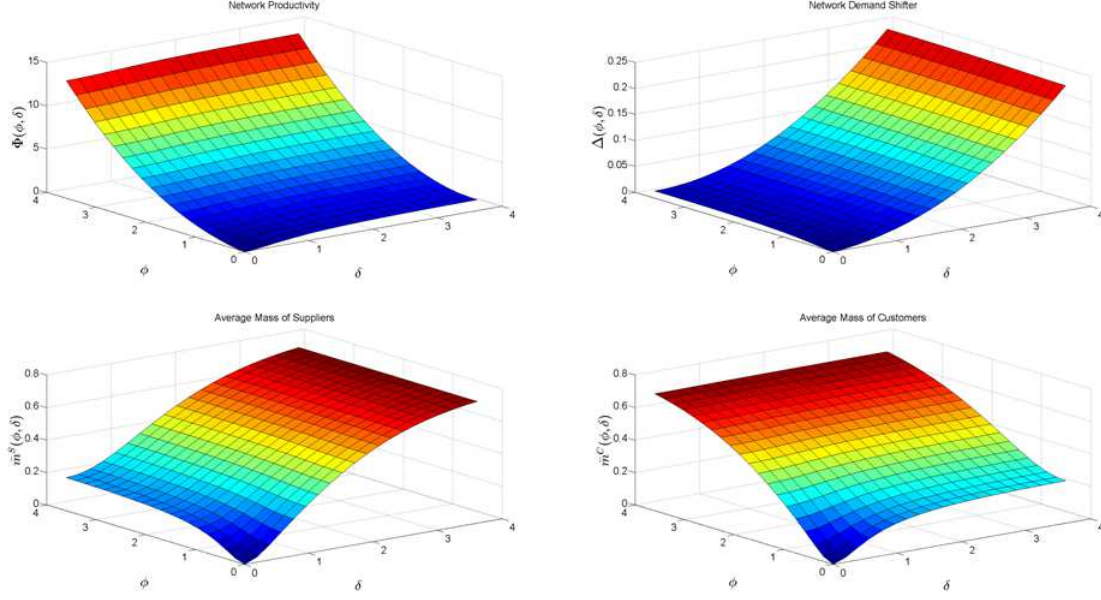


Figure 7: Firm network characteristics and matching in steady-state

Given the analysis in the previous section, the matching between a χ -firm and its suppliers and customers depends only on $\Delta(\chi)$ and $\Phi(\chi)$ respectively, and therefore we can alternatively consider the average supplier and customer revenue of firms with network quality Δ and productivity Φ respectively (which I henceforth refer to as Δ - and Φ -firms), given by:

$$\tilde{R}_S(\Delta) = \frac{\int_{S_\chi} \tilde{m}[\Delta\Phi(\chi')\Delta_H] R(\chi') dF_\chi(\chi')}{\tilde{M}_S(\Delta)} \quad (4.27)$$

$$\tilde{R}_C(\Phi) = \frac{\int_{S_\chi} \tilde{m}[\Delta(\chi')\Phi\Delta_H] R(\chi') dF_\chi(\chi')}{\tilde{M}_C(\Phi)} \quad (4.28)$$

Since firms with higher network productivity and quality also tend to have higher revenue, the assortativity of firm matching (in terms of revenue) can be characterized by the gradients of the functions \tilde{R}_S and \tilde{R}_C . Differentiating equation (4.27), for example, we obtain:

$$\tilde{R}'_S(\Delta) = \frac{\Delta}{\tilde{M}_S(\Delta)} \int_{S_\chi} [R(\chi') - \tilde{R}_S(\Delta)] \varepsilon_{\tilde{m}}[\Delta\Phi(\chi')\Delta_H] \tilde{m}[\Delta\Phi(\chi')\Delta_H] dF_\chi(\chi') \quad (4.29)$$

where $\varepsilon_{\tilde{m}}$ is the elasticity of the scalar matching function \tilde{m} . From equation (4.29) and the equivalent derivative of equation (4.28), we make the following observation: if the elasticity $\varepsilon_{\tilde{m}}$ is constant, then $\tilde{R}_S(\cdot)$ and $\tilde{R}_C(\cdot)$ are constant functions, and in this sense the assortativity of matching between firms is neutral, with average customer and supplier revenue independent of firm size. This suggests that the elasticity $\varepsilon_{\tilde{m}}$ plays a crucial role in shaping the assortativity of matching between firms in general.

We can characterize this even further by considering the average revenue of Δ' -firms that supply a Δ -firm, the derivative of which with respect to Δ is:

$$\tilde{R}'_S(\Delta|\Delta') = \frac{\Delta}{M_S(\Delta)} \int_0^1 [\mu\Delta_H\Delta'\Phi' - \tilde{R}_S(\Delta)] \varepsilon_{\tilde{m}}(\Delta\Phi'\Delta_H) \tilde{m}(\Delta\Phi'\Delta_H) dF_{\Phi}(\Phi') \quad (4.30)$$

Since \tilde{m} is an increasing function, then from this equation we can make an even stronger observation about the role of $\varepsilon_{\tilde{m}}$: the assortativity of matching between Δ -buyers and Δ' -sellers is positive if $\varepsilon_{\tilde{m}}$ is increasing, and is negative if $\varepsilon_{\tilde{m}}$ is decreasing. The same is also true regarding the assortativity of matching between Φ -buyers and Φ' -sellers.

This analysis then begs the question: what determines the elasticity of the matching function? From equation (4.22), the matching function elasticity is equal to:

$$\varepsilon_{\tilde{m}}(x) = \frac{x}{(1 - \beta\nu) \bar{f}} \left[\frac{F'_\xi \left[\frac{x - \beta\nu \bar{f}}{(1 - \beta\nu) \bar{f}} \right]}{F_\xi \left[\frac{x - \beta\nu \bar{f}}{(1 - \beta\nu) \bar{f}} \right]} \right] \quad (4.31)$$

In the special case when $\nu = 0$, so that the model is completely static, the elasticity of the matching function is completely determined by the elasticity of the distribution function F_ξ of the relationship cost shock. Consequently, this implies that the assumed parametric form for F_ξ will be crucial for determining the model's predictions regarding the assortativity of matching between firms, an issue that we will return to when we discuss numerical estimation of the model in section 5.

4.3.3 Geographic distribution of trade partners

Reintroducing geography into the model simply requires rewriting the matching function as:

$$m[\chi, \chi' | \tau(D)] = \tilde{m} \left[\frac{\Delta(\chi) \Phi(\chi') \Delta_H}{\tau(D)^{\sigma-1}} \right] \quad (4.32)$$

and using equations (3.41) and (3.42) to specify the network characteristic functions. We can then easily compute the average supplier and customer distance of a χ -firm as follows:

$$D_S(\chi) = \frac{\int_0^1 \int_{S_\chi} D m[\chi, \chi' | \tau(D)] dF_\chi(\chi') dD}{M_S(\chi)} \quad (4.33)$$

$$D_C(\chi) = \frac{\int_0^1 \int_{S_\chi} D m[\chi, \chi' | \tau(D)] dF_\chi(\chi') dD}{M_C(\chi)} \quad (4.34)$$

Exactly the same analysis as in section 4.3.2 can be used to show that the matching function elasticity plays a key role in determining whether larger firms tend to have suppliers and customers that are located further or nearer by. When the elasticity is increasing, larger firms tend to have closer trade partners than smaller firms.

4.3.4 Relationship dynamics

Even in the steady-state of the model, there is churning of firm relationships due to the stochastic nature of the fixed relationship cost. First, note that the unconditional probabilities that a χ -firm will retain any one of its suppliers or customers are given by:

$$\rho_S^{ret}(\chi) = \nu + (1 - \nu) \int_{S_\chi} a(\chi, \chi') dF_\chi(\chi') \quad (4.35)$$

$$\rho_C^{ret}(\chi) = \nu + (1 - \nu) \int_{S_\chi} a(\chi', \chi) dF_\chi(\chi') \quad (4.36)$$

Since these probabilities are constant in steady-state, the unconditional duration of relationships between a χ -firm and its suppliers and customers follows a geometric distribution, with means $\frac{1}{1-\rho_S^{ret}(\chi)}$ and $\frac{1}{1-\rho_C^{ret}(\chi)}$ respectively. Furthermore, since the matching function is equal to the acceptance function in the steady-state of the model, then equations (4.35) and (4.36) deliver sharp predictions about the relation between the retention probabilities and the masses of a firm's suppliers and customers:

$$\rho_S^{ret}(\chi) = \nu + (1 - \nu) M_S(\chi) \quad (4.37)$$

$$\rho_C^{ret}(\chi) = \nu + (1 - \nu) M_C(\chi) \quad (4.38)$$

Firms with more suppliers and customers are therefore more likely to retain existing trading relationships.

Note that firms in the model are also more likely to trade with existing partners than new ones because of the sticky nature of relationships. If a $\chi - \chi'$ relationship was active in the previous period, the probability that it will be maintained in the current period is equal to $\nu + (1 - \nu) a(\chi, \chi')$, whereas the probability that it will be newly-formed is equal to $(1 - \nu) a(\chi, \chi')$. The fractions of suppliers and customers that are new for a χ -firm every

period are therefore given respectively by:

$$\rho_S^{new}(\chi) = \frac{\int_{S_\chi} (1 - \nu) a(\chi, \chi') [1 - m(\chi, \chi')] dF_\chi(\chi')}{M_S(\chi)} \quad (4.39)$$

$$\rho_C^{new}(\chi) = \frac{\int_{S_\chi} (1 - \nu) a(\chi', \chi) [1 - m(\chi', \chi)] dF_\chi(\chi')}{M_C(\chi)} \quad (4.40)$$

Finally, it is useful to point out that the parameter ν controls the rate of convergence between steady-states. As an illustrative example, consider an economy that is in steady-state at $t = 0$ with both the relationship fixed cost f and the reset friction ν being finite, and denote the matching function in this economy by m_{ss} . Suppose then that the fixed relationship cost f becomes either infinite or zero, and denote the new steady-state matching function by m'_{ss} (identically zero or one respectively). From equations (4.2) and (4.11), the matching function evolves according to:

$$\hat{m}_t(\chi, \chi') = \nu^t \hat{m}_0(\chi, \chi') \quad (4.41)$$

where $\hat{m}_t(\chi, \chi') \equiv m_t(\chi, \chi') - m'_{ss}(\chi, \chi')$ is the deviation of the matching function from the new steady-state. When relationships are stickier (larger ν), convergence between steady-states is slower.

5 Numerical Analysis

Having characterized the theoretical counterparts of the empirical moments described in section 2.2, I now take the model to data by estimating the steady-state of the model via simulated method of moments. I begin by specifying the remaining parametric assumptions in the model.

5.1 Parametric assumptions

First, given that the firm size distribution appears to be approximately log-normal (Figure 2.2.1), I assume that the log of fundamental firm productivities and qualities, ϕ and δ , are jointly Gaussian with zero mean and covariance matrix given by:

$$\Sigma = \begin{bmatrix} v_\phi^2 & \rho v_\phi v_\delta \\ \rho v_\phi v_\delta & v_\delta^2 \end{bmatrix} \quad (5.1)$$

Note that in the empty network with $m(\chi, \chi') = 0$ for all $\chi, \chi' \in S_\chi$, this assumption would imply that firm revenue and employment are exactly log-normally distributed.

Parameterization of the distribution function F_ξ of the relationship cost shock requires slightly more careful consideration. As discussed in section 4.3.2, the elasticity of F_ξ plays a key role in determining qualitative properties of the model, and in particular the gradient of the elasticity of F_ξ is directly related to the assortativity of matching between firms. As it turns out, almost all of the standard continuous distributions with support on $[0, \infty)$ feature a monotonically decreasing elasticity.²⁸ One notable exception is the Gompertz or log-Weibull distribution, which is used extensively in survival analysis and has the following distribution function:

$$F_\xi(x) = 1 - e^{-b_\xi(e^{s_\xi x} - 1)} \quad (5.2)$$

where b_ξ is a scale parameter and s_ξ characterizes the shape of the distribution. From a mathematical point of view, assuming that the relationship cost shock follows a Gompertz distribution is desirable because the sign of the elasticity gradient of the distribution is variable when $s_\xi \in (0, 1)$, which therefore allows for flexibility in the model's predictions regarding the assortativity of firm matching.

From an economic standpoint, a Gompertz-distributed relationship cost shock can be interpreted as follows. Suppose that upon meeting, a pair of firms takes a random amount of time (within the period) to negotiate the potential arrangements of the trading relationship, and that the fixed cost of the relationship is proportional to the amount of time that it takes for negotiations to be completed. Suppose also that the probability with which negotiations continue to drag on conditional on no agreement having been reached at a given point in time declines with time. If this process is characterized by the negotiation time having an exponential hazard rate, then the fixed cost of the relationship has a Gompertz distribution. Based on these considerations, I parameterize the relationship cost shock according to (5.2). With the mean of ξ_t fixed at 1, this pins down the scale parameter b_ξ given a choice of the shape parameter s_ξ .

Finally, trade costs are parameterized according to:

$$\tau(D) = (1 + \kappa D)^\epsilon \quad (5.3)$$

where κ measures the overall level of trade costs and ϵ measures the elasticity of trade costs with respect to distance.²⁹ Since the maximum possible trading distance in the model is

²⁸These include (at least) the Fréchet, Weibull, log-normal, Gamma, generalized Pareto, and log-logistic distribution.

²⁹Note that with this parameterization, τ is log-subadditive for any $\kappa, \epsilon \geq 0$, and therefore trade costs satisfy the triangle inequality.

normalized to 1, κ can also be interpreted as the cost of trading with the most distant firms relative to trading with firms that are right next door. Note that trade costs are non-existent when either $\kappa = 0$ or $\epsilon = 0$.

5.2 Parameter estimation

The above parameterization of the model gives us a total of 12 parameters: the elasticity of substitution σ ; input suitability α ; mean f and shape s_ξ of the relationship fixed cost; reset friction ν ; parameters of the χ distribution, v_ϕ , v_δ , and ρ ; parameters of the trade cost function k and ϵ ; labor supply L ; and the household discount factor β .

Since the Compustat data is of annual frequency, I set $\beta = .96$. Also, note that the total labor supply L only enters the set of equilibrium conditions through equation (3.34). If we write the magnitude of the fixed relationship cost f as a fraction \hat{f} of the total labor supply, then from equations (4.4) and (4.13), we see that the activation function a is independent of L . Equation (4.3) then implies that the matching function is also independent of L , and therefore so are the network characteristic functions defined by (3.19) and (3.20). In other words, the parameter L affects equilibrium variables only by scaling firm size one-to-one. I therefore fix $L = 1$ and compare normalized moments of the model to the corresponding normalized moments of the data, as described in section 2.2.1.

The remaining 10 parameters of the model are estimated using simulated method of moments. Recall that the five sets of empirical moments discussed in sections 2.2.1-2.2.5 were respectively:

1. \bar{X}_b , the normalized quantile level of variable X evaluated at the midpoint of quantile bin b ;
2. ${}^R\bar{Q}_b^X$, the average quantile of variable X for all firms with revenue falling in quantile bin b , given by equation (2.3);
3. ${}^R\bar{Q}_b^{S,X}$ and ${}^R\bar{Q}_b^{C,X}$, the average quantile of variable X amongst all suppliers and customers respectively of all firms with revenue falling in quantile bin b , given by equations (2.4) and (2.5);
4. \bar{D}_b^S and \bar{D}_b^C , the average normalized supplier and customer distances respectively amongst all firms with revenue falling in quantile bin b , given by equations (2.6) and (2.7);
5. $\bar{\rho}_b^{S,ret}$ and $\bar{\rho}_b^{C,ret}$, the dynamic moments capturing the rates at which firms retain old trading partners, given by equations (2.10) and (2.11).

One option for the estimation procedure is to target all of the moments described above. Since employment is highly correlated with revenue in the data, however, I choose to omit targeting the firm employment distribution (\bar{L}_b), as well as the correlation between revenue and employment (${}^R\bar{Q}_b^L$). Furthermore, instead of targeting all of the moments that characterize firm-to-firm matching, I target only the revenue quantiles of suppliers and customers across firms (${}^R\bar{Q}_b^{S,R}$ and ${}^R\bar{Q}_b^{C,R}$), and use the remaining matching moments as overidentification tests of model fit. This leaves $13 \times N_{bin}$ sets of moments for estimating 10 parameters.

The estimation procedure is as follows. First, to reduce simulation error, I generate N_{sim} random seeds ($\tilde{\varepsilon}_\phi, \tilde{\varepsilon}_\delta$) from a two-dimensional standard multivariate normal distribution.³⁰ Then, for every candidate set of parameter values, I compute the theoretical moments corresponding to the targeted moments described above for a set of N_{sim} simulated firms. To do so, I first solve for the values of the steady-state network characteristic and matching functions at a set of $N_{grid} \times N_{grid}$ points using the algorithm described in the appendix. I then solve for the functions $R(\cdot)$, $M_S(\cdot)$, $M_C(\cdot)$, $D_S(\cdot)$, $D_C(\cdot)$, $\rho_S^{ret}(\cdot)$, and $\rho_C^{ret}(\cdot)$ at these same grid points using equations (3.26), (3.11), (3.12), (4.33), (4.34), (4.35), and (4.36). Given the current values of v_ϕ , v_δ , and ρ , I then compute:

$$\begin{bmatrix} \log \phi \\ \log \delta \end{bmatrix} = \begin{bmatrix} v_\phi \sqrt{1 - \rho^2} & \rho v_\phi \\ 0 & v_\delta \end{bmatrix} \begin{bmatrix} \tilde{\varepsilon}_\phi \\ \tilde{\varepsilon}_\delta \end{bmatrix} \quad (5.4)$$

for each simulated firm (thereby maintaining consistency with the desired covariance matrix (5.1)), and then use bilinear interpolation to obtain the theoretical values of R , M_S , M_C , D_S , D_C , ρ_S^{ret} , and ρ_C^{ret} for each firm.

Having computed the theoretical counterparts of the target moments, I then compute the distance between these and the empirical moments according to:

$$\mathcal{D} = (|\mathcal{M}_{data} - \mathcal{M}_{model}|)^T \mathcal{W} (|\mathcal{M}_{data} - \mathcal{M}_{model}|) \quad (5.5)$$

where \mathcal{M}_{data} and \mathcal{M}_{model} are vectors containing the stacked empirical and model moments respectively, and \mathcal{W} is the pseudo-inverse of the covariance matrix of the empirical moment vector, estimated by bootstrapping techniques.³¹ Starting from an arbitrary initial choice of parameter values, I then execute a simulated annealing algorithm to minimize \mathcal{D} . Standard errors are computed using a bootstrap procedure, in which I repeat the estimation procedure

³⁰In order to obtain bounded support for the joint distribution of ϕ and δ , which is necessary for numerical solution of the model, I truncate the distributions of both $\tilde{\varepsilon}_\phi$ and $\tilde{\varepsilon}_\delta$ at the 95th percentiles.

³¹I resample with replacement 2000 times from the set of firms for both the Capital IQ and Compustat datasets, and compute the covariance matrix resampled data. I do not perform resampling along the time dimension for the Compustat data, although in principle this is possible using block bootstrapping techniques.

described above after replacing \mathcal{M}_{data} by the corresponding moments from a bootstrap re-sampling of the original data. To account for simulation error, I also regenerate the random seeds $(\tilde{\varepsilon}_\phi, \tilde{\varepsilon}_\delta)$ each time the estimation is performed.

5.3 Results

5.3.1 Parameter estimates

The parameter values obtained using the estimation procedure described above are shown in Table 1. From this, we make several observations.

First, the estimated value of the mean static relationship cost f appears to be small, but recall that total labor supply is normalized to 1 in the estimation, and therefore the estimate implies that around 7% of total production labor is used for managing relationships. At the firm-level, the model predicts that labor costs associated with managing existing trade relationships within a firm account for around 1.3% of total labor costs on average.

Second, the reset friction parameter ν affects the rate at which firms form new trading relationships and destroy existing ones. At these parameter estimates, the model predicts that the mean duration of a firm's relationships with its suppliers and customers is around 1.9 years, which is very close to the empirically-measured mean relationship duration of 1.74 years. The model also predicts that the average relationship termination rate across firms is around 34%, which again is very close to the empirical supplier and customer termination rates of 38.4% and 30.1% respectively.

Third, although the substitution elasticity σ is not very precisely estimated, the point estimate plus or minus one standard error falls well within the range of values typically estimated in the literature.³² This is reassuring given that the estimation is based on data in which the intensive margin of trade (transaction values) is unobserved.

Finally, the parameters governing the distribution of fundamental firm characteristics appear to be well identified, with relatively small standard errors, but the trade cost parameters are less precisely estimated. As discussed below, this is perhaps related to the inability of the model to match the qualitative relationship between firm size and trading partner distance.

5.3.2 Model fit

To examine the model's fit with data, Figures 8-12 reproduce the graphs characterizing the empirical moments described in section 2.2, but with the model's simulated moments

³²See for example Broda and Weinstein (2006).

Parameter		Value	Standard Error
mean of relationship cost	f	.070	.02
meeting friction	ν	.647	.03
shape of relationship cost shock	s_ξ	.585	.11
elasticity of substitution	σ	3.02	.27
input suitability	α	.347	.09
variance of fundamental productivity	v_ϕ	.364	.06
variance of fundamental quality	v_δ	.544	.06
correlation between ϕ and δ	ρ	-.241	.07
trade cost level	κ	.688	.18
elasticity of trade cost with distance	ϵ	.348	.12

Table 1: Estimated parameter values

overlaid. With regard to the firm-level distributions shown in Figure 8, we see that the theoretical firm revenue distribution closely approximates the empirical distribution, and takes on the same log-normal shape. The firm in-degree and out-degree distributions, on the other hand, are harder for the model to match exactly, although the theoretical and empirical distributions share the same convex shape. Comparing the theoretical degree distributions to the Poisson (random matching) and Pareto (preferential attachment) approximations described in section 2.2.1, we see that the model’s predicted distributions lie somewhere between the distributions of the two parametric forms. This is perhaps not surprising, given that the structural model features both elements of random reset shocks as well as preferential activation (and non-termination) with larger suppliers and customers. The firm employment distribution predicted by the model (which is untargeted in the estimation) resembles the empirically-observed employment distribution in terms of the log-normal shape, but the fit is poorer compared to the revenue distribution.

Figure 9 shows the model’s fit with regard to the correlation of firm revenue with employment, in-degree, and out-degree. As in the data, the model predicts that firms with larger revenue also tend to have larger employment, more suppliers, and more customers. Furthermore, the model closely matches the specific quantiles of these variables for firms in each revenue quantile bin, even for the untargeted employment distribution.

Next, we examine the model’s fit with regard to the assortativity of matching between firms, shown in Figure 10. From these graphs, we see that the model is able to reproduce the positive assortative matching between firms documented in the data, whether with regard to revenue (targeted), or employment, in-degree, and out-degree (untargeted). However, in each case, the model fit is better for firms at the upper-end of the revenue distribution. The fit with regard to matching between firms and their suppliers in terms of revenue, for example,

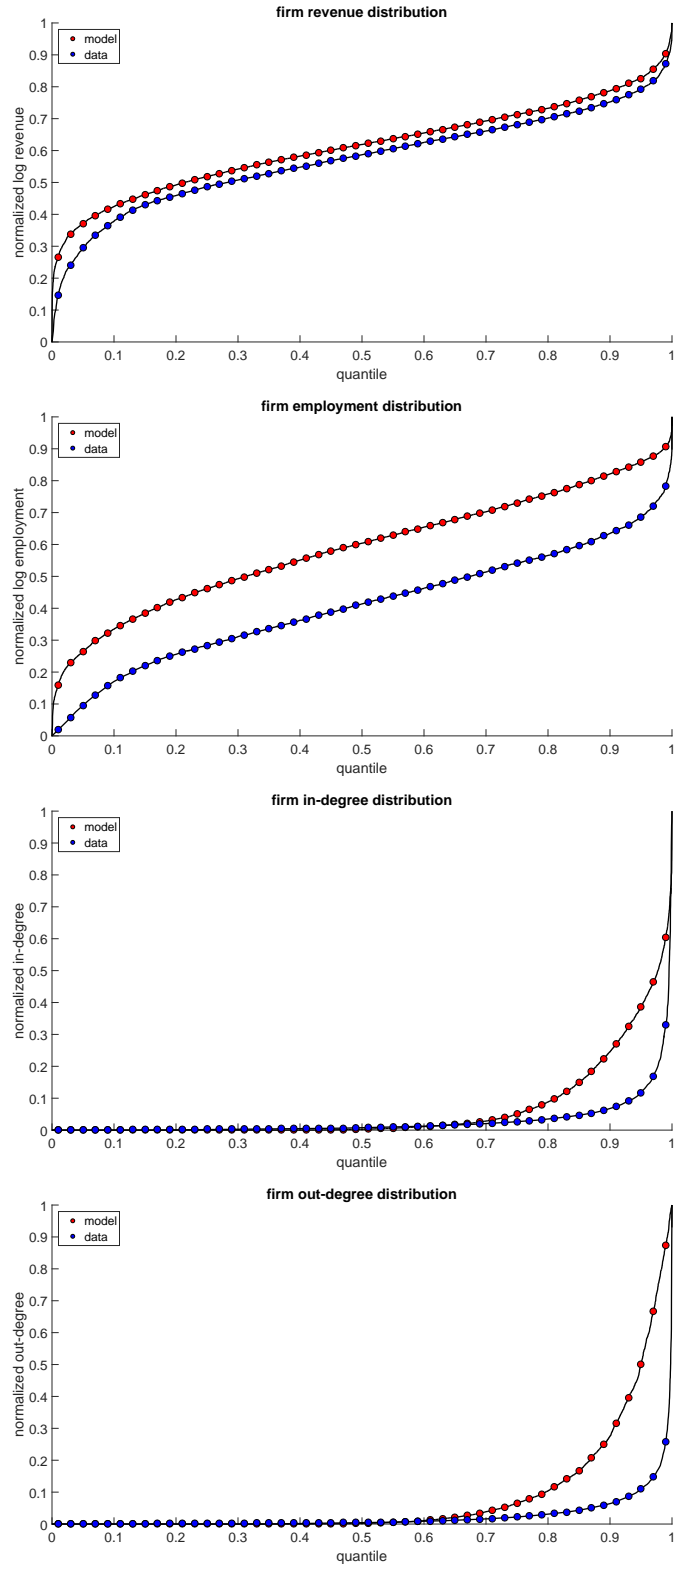


Figure 8: Model fit: firm-level distributions

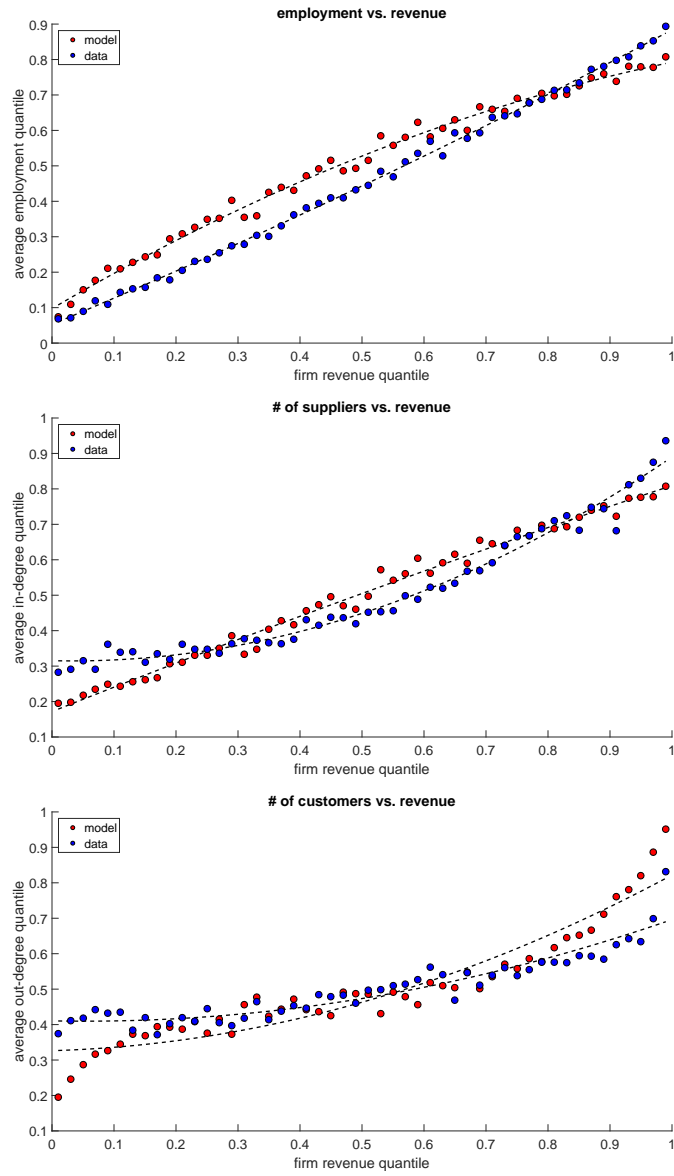


Figure 9: Model fit: Bivariate distributions

is almost perfect for firms with revenue above the median, but is poorer for firms with revenue below the median. This suggests that the economic tradeoffs involved in forming and terminating trading relationships may be significantly different for small versus large firms. In particular, the empirical moments of the matching distributions imply that small firms are likely to match with suppliers and customers that are larger than the theoretical mechanism in the model suggests.

With regard to the geographic distribution of a firm’s suppliers and customers, Figure 11 shows that the model is unable to replicate the qualitative feature of the data that larger firms tend to match with trade partners that are located closer to themselves, although in terms of levels the average normalized distances to suppliers and customers predicted by the model for larger firms are not too far off from the corresponding empirical moments. This discrepancy between model and data suggests that additional theoretical mechanisms beyond the relationship frictions studied in this paper are needed to generate both positive assortative matching between firms as well as average trade partner distances that decline with firm size. The pattern observed in Figure 11 might be generated by a trade model featuring an endogenous geographic distribution of firms with positive externalities in each location, for example, so that larger firms tend to be located closer to larger firms. Embedding endogenous geography, however, is beyond the scope of this paper.

Finally, Figure 12 shows the model’s fit with respect to the moments characterizing firm relationship dynamics. Here, we see that the model replicates the empirically-observed positive relation between firm size and the rate at which firms retain existing suppliers and customers, although the exact moments do not line up perfectly. Nonetheless, as discussed above, the predicted relationship durations and relationship termination rates are very close to their empirical counterparts on average.

6 Counterfactuals

Having estimated the parameters of the model, I now return to addressing the key question initially posed in the introduction to this paper: what are the quantitative implications of stickiness in firm-to-firm relationships for the responses of aggregate trade patterns and welfare to shocks? To answer these questions, I study the model’s transition dynamics in response to three kinds of counterfactual changes: declines in trade costs (section 6.1), declines in relationship costs (section 6.2), and idiosyncratic fluctuations in firm-level characteristics (section 6.3). I also examine the importance of accounting for rational firm expectations in computing these counterfactual dynamics (section 6.4), and revisit the efficiency of the dynamic market equilibrium by studying a simple policy exercise in which the fixed rela-

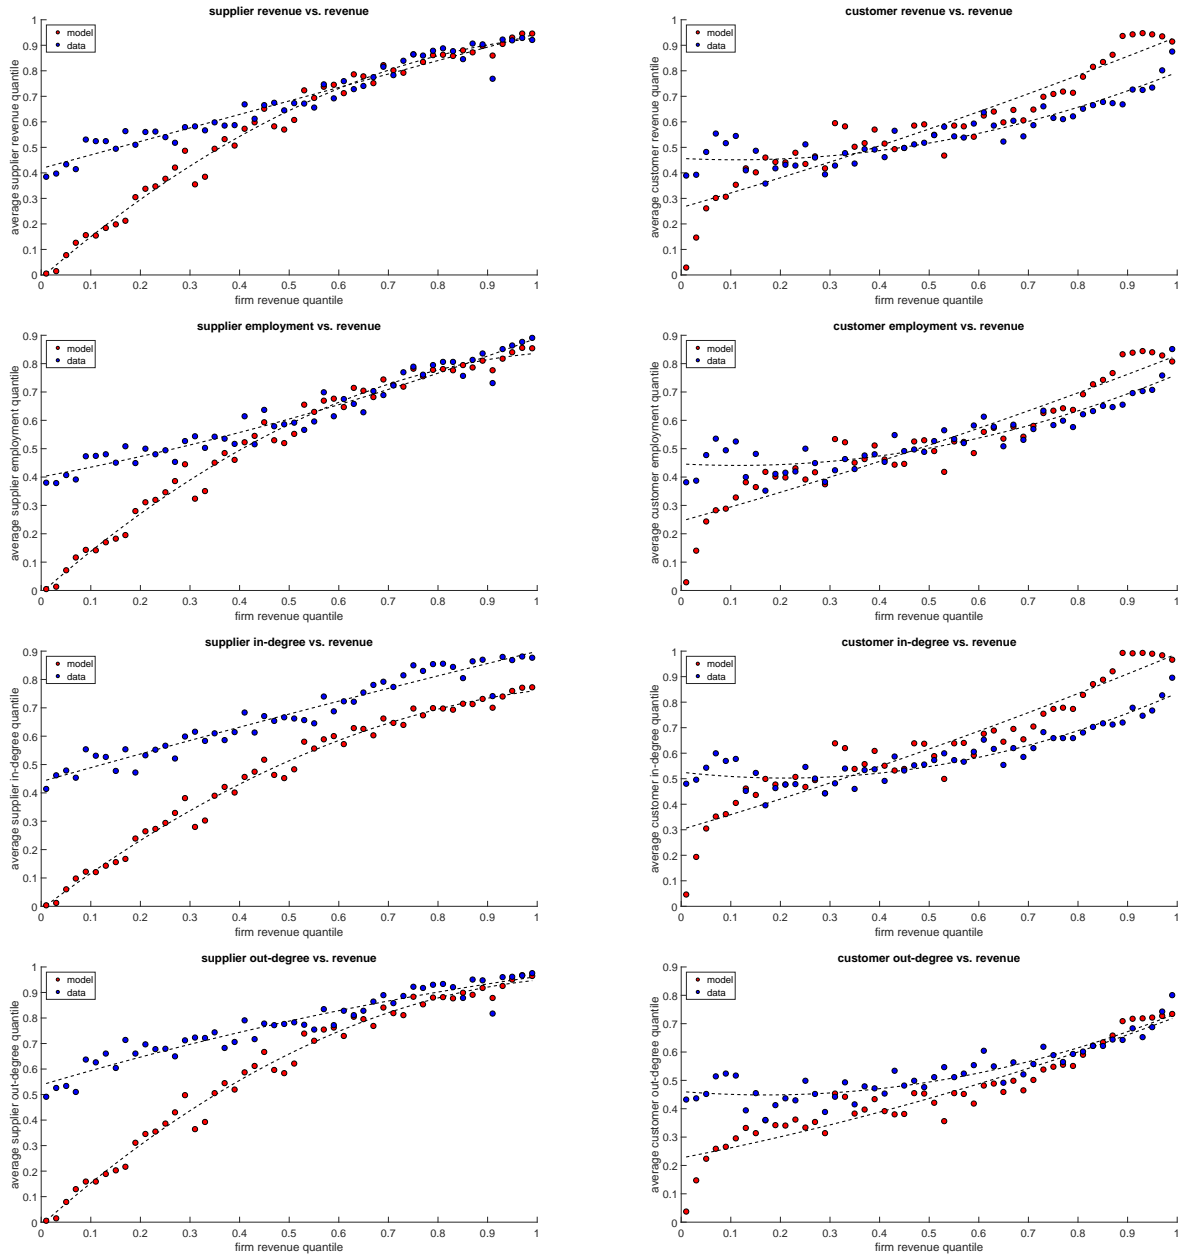


Figure 10: Model fit: matching distributions

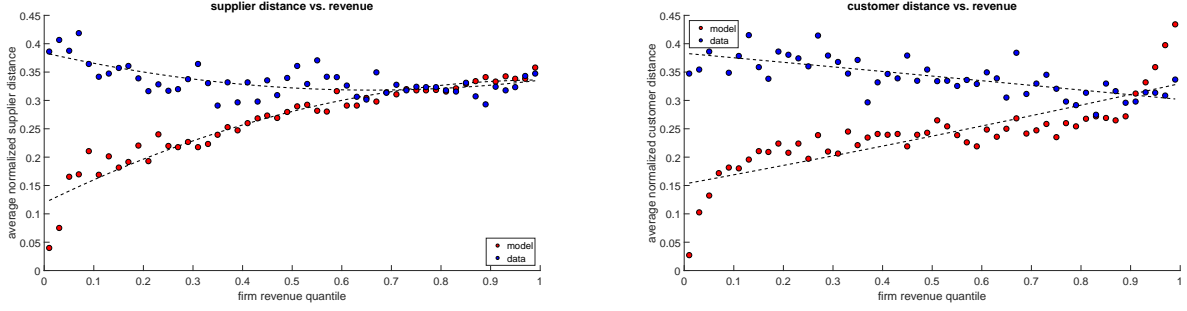


Figure 11: Model fit: geographic distribution of suppliers and customers

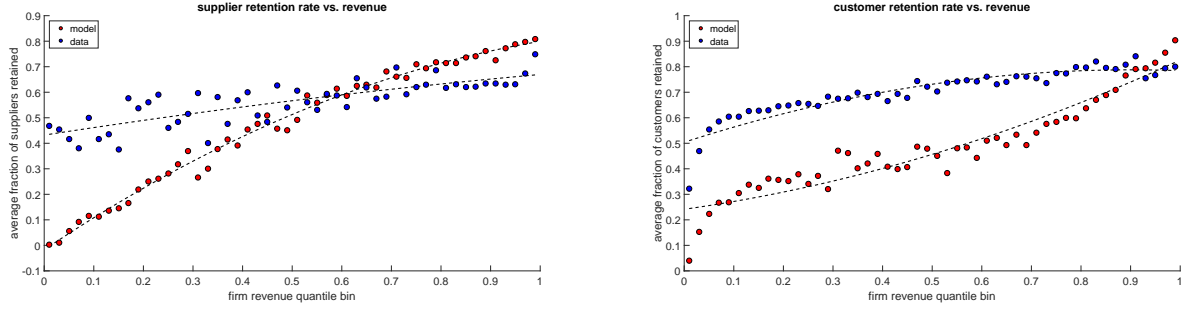


Figure 12: Model fit: dynamic distributions

tionship cost is subsidized by a planner who obtains revenue from an ad valorem import tax (section 6.5).

6.1 Trade cost shocks

To examine how sticky relationships affect the dynamic responses of aggregate trade volumes and welfare to trade cost shocks, I study the model's transition dynamics following a change in the overall trade cost level κ to some counterfactual level, starting from the steady-state of the model with parameters set at the SMM estimates. I assume that the shock hits the economy at $t = 0$ *after* all relationship cost shocks have been realized and all activation and termination decisions have been made, so that firms can readjust the intensive margin of trade in the initial period post-shock but not the extensive margin. In other words, the initial response of the economy to the trade cost shock takes the network of firm trade as fixed. From $t = 1$ onwards, firms adjust both the intensive and extensive margins of trade in response to the shock.

Recall that the aggregate value of imports at date t from a location a distance D away

is given by:

$$\bar{R}_t(D) = \left(\frac{\alpha}{\mu}\right)^{\sigma-1} \tau(D)^{1-\sigma} \Delta_{H,t} \int_{S_x} \int_{S_x} m_t[\chi, \chi' | \tau(D)] \Delta_t(\chi) \Phi_t(\chi') dF_\chi(\chi) dF_\chi(\chi') \quad (6.1)$$

A decline in the cost of trade $\tau(D)$ therefore affects trade volumes statically through a direct reduction in the cost of inputs purchased (via the term $\tau(D)^{1-\sigma}$), as well as dynamically through changes in the incentives that firms face in forming and terminating relationships (via the matching function m_t). In the initial period of the shock, the matching function is assumed to be fixed, and the short-run change in trade therefore occurs only through the static channel. In the long-run, the total effect of the trade cost shock on trade volumes incorporates adjustments of firm-to-firm trade along both the intensive and extensive margins.

Figure 13 shows the dynamic responses of trade and welfare following a uniform 5% decline in gross trade costs across all locations.³³ The first graph shows the transition paths of exports from a given location (measured as the percentage change relative to the pre-shock steady-state) to locations integrated over each quadrant of the unit circle.³⁴ The second and third graphs decompose these changes in trade volumes into changes along the extensive and intensive margins respectively, while the fourth graph shows changes in welfare. From these graphs, we observe the following. First, in the initial period of the shock, exports to all locations increase, with the total value of exports rising by around 8%. Since the set of active trading relationships is assumed to be fixed, all of these gains are generated by firms selling more to existing customers. Notice also that the initial increase in exports is larger for locations that are further away, so that the geographic distribution of trade immediately becomes more dispersed following the shock.

After the initial period, the decline in trade costs induces firms to accumulate more trading partners. Over time, the value of exports to all locations therefore continues to grow. Observe that along the transition path, the growth in the mass of active relationships is accompanied by a decline in the amount of trade per active relationship. The dynamic gains in aggregate trade are therefore driven solely by increases in the extensive margin of firm-to-firm trade. Once firms have fully adjusted their trading relationships in response to the shock, total exports to all locations are almost 30% higher relative to the pre-shock steady-state. The endogenous adjustment of firm-level relationships therefore amplifies the elasticity of aggregate trade with respect to trade costs by more than three times. Similarly,

³³Specifically, a change in κ corresponding to a 5% decline in the average trade cost measure $\int_0^1 (1 + \kappa D)^\epsilon dD$.

³⁴Since all locations are symmetric, the values of exports and imports between any pair of locations are identical.

the welfare gains from the reduction in trade costs are close to four times higher in the post-shock steady-state than in the initial period of the shock (although the absolute welfare gains are small). Note that the dynamic amplification effect is larger for exports to more distant locations, so that the geographic dispersion of trade also increases over time.

In addition to studying a uniform decline in the cost of trade across all locations, we can also use the model to study the effects of a bilateral reduction in the costs of trade between a given pair of locations. Since the set of locations is continuous, a change in trade costs between a single pair of locations leaves aggregate variables in each location unchanged.³⁵ The response of trade is therefore given by equation (6.1) with $\Delta_{H,t}$, $\Delta_t(\cdot)$ and $\Phi_t(\cdot)$ held fixed at their respective pre-shock steady-states. Nonetheless, the economic mechanisms remain the same: the bilateral decline in trade costs affects trade volumes both statically and dynamically.

Figure 14 shows the responses of trade following a 5% decline in gross bilateral trade costs for different distances between importing and exporting locations.³⁶ Again, we see that the initial increase in trade is dynamically amplified by the accumulation of additional trading partners by firms in response to the trade cost shock, and that the magnitude of the amplification is around a factor of three for all locations but is larger for more distant locations. Note that the response of trade in the initial period of the shock (the x-intercept in the first graph) is determined solely by the elasticity of substitution σ , as it would be in the frictionless model.

6.2 Relationship cost shocks

Lower variable trade costs reduce the cost of firm-to-firm trade along the intensive margin. How do trade patterns and welfare respond to changes in the cost of firm-to-firm trade along the extensive margin when firm relationships are sticky? To study this, I examine the model's transition dynamics following a change in the average value f of the relationship cost shock. Again, I assume that the shock hits the economy at $t = 0$ after all relationships have been set, and only allow firms to create and terminate relationships from $t = 1$ onwards. Furthermore, to enable consistent quantitative comparison with the results of the previous section, I compute the magnitude of the change in f in the following way.

Consider a decline in variable trade costs across all locations corresponding to a change in κ to some counterfactual level κ' . The cost of this change across steady-states if it were

³⁵One can think of this as a small open economy assumption but applied to a pair of locations.

³⁶Specifically, a change in κ corresponding to a 5% decline in $(1 + \kappa D)^\epsilon$ for each value of D .

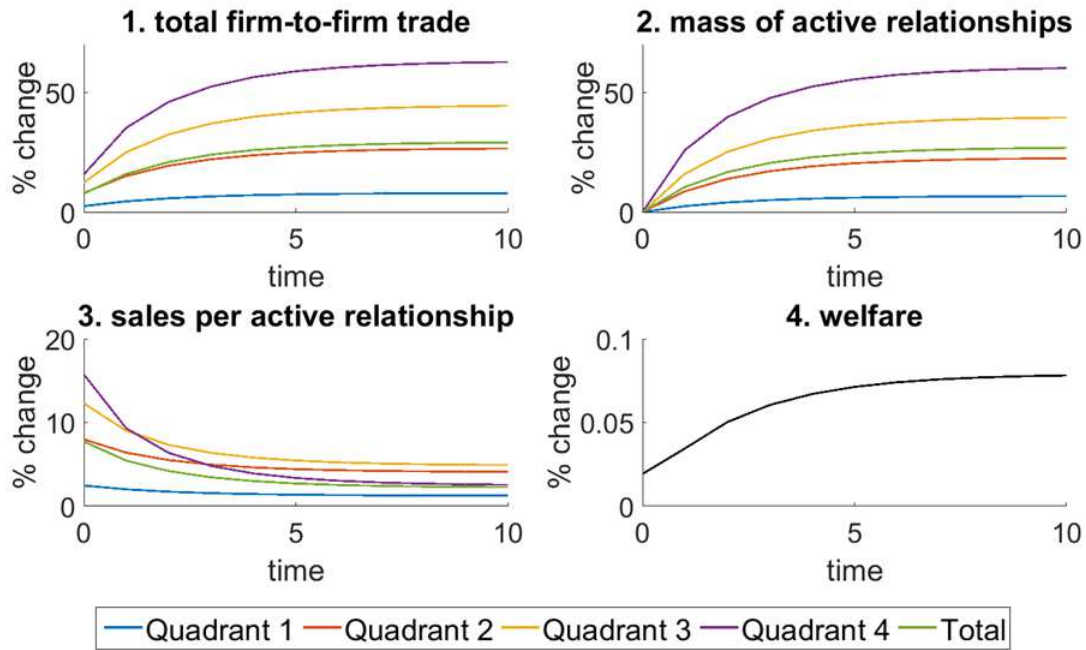


Figure 13: Responses of trade and welfare to 5% decline in global trade costs

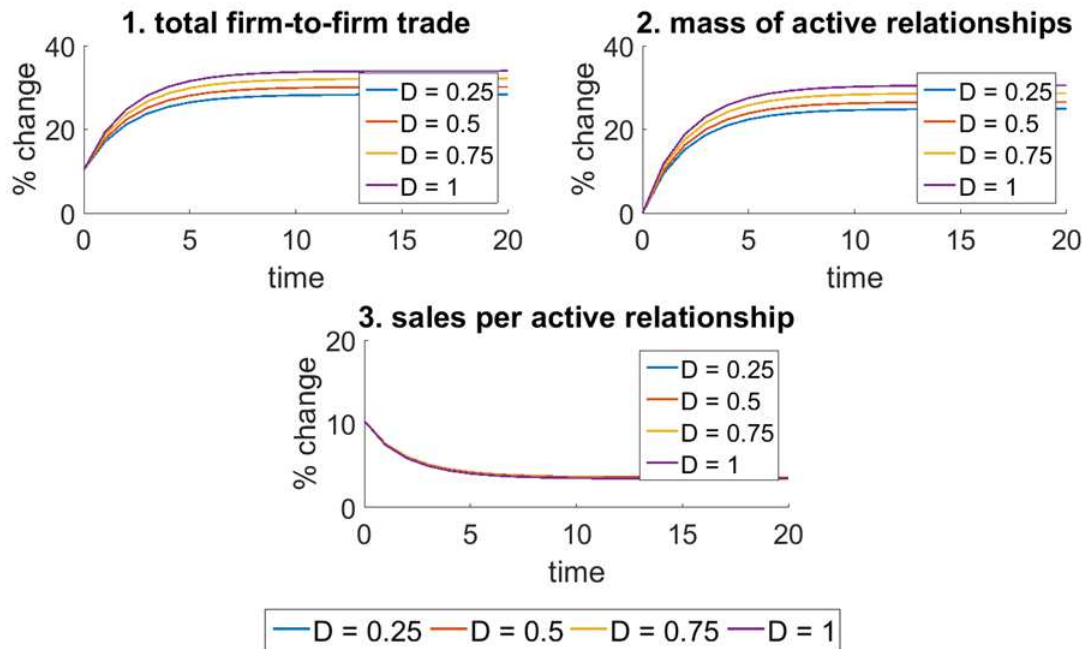


Figure 14: Responses of trade and welfare to 5% decline in bilateral trade costs

to be implemented by an ad valorem subsidy to exports would be given by:

$$T_{\kappa}(\kappa, \kappa') = \int_0^1 \left[(1 + \kappa D)^{\epsilon} - (1 + \kappa' D)^{\epsilon} \right] \bar{R}(D|\kappa') dD \quad (6.2)$$

where $\bar{R}(\cdot|\kappa')$ is the aggregate value of trade in the steady-state corresponding to κ' . Similarly, the cost of a decline in f to some counterfactual value f' if it were to be implemented by a subsidy to the cost of maintaining relationships would be equal to:

$$T_f(f, f') = (f - f') L_f(f') \quad (6.3)$$

where here $L_f(f')$ is the total mass of labor used to pay relationship fixed costs in the steady-state corresponding to f' . With κ and f set at the SMM parameter values, I therefore compute the value of f' such that $T_f(f, f') = T_{\kappa}(\kappa, \kappa')$ for a given value of κ' .

Figure 15 shows the responses of aggregate trade and welfare in response to a decline in f corresponding to the 5% decline in global variable trade costs studied in section 6.1.³⁷ From these graphs, we see that the effects of lower relationship costs are qualitatively similar to the effects of lower variable trade costs: exports to all locations increase over time, driven by growth in the mass of active relationships and accompanied by a decline in the intensive margin of trade. Quantitatively, however, the effects of a decrease in f on aggregate trade and welfare are much larger than the corresponding effects following a decrease in κ . The increase in total exports in the post-shock steady-state relative to the pre-shock steady-state is around 50% higher than the corresponding increase resulting from the decline in variable trade costs. Similarly, the long-run welfare gains are around 75% higher. Since the rates of adjustment in response to the shocks are similar in the two cases, these results suggest that policy measures targeting the frictions that firms face in establishing trading relationships can be equally as if not more cost-effective than ad valorem trade subsidies.

As in section 6.1, we can also study the effects of a decline in the bilateral cost of relationships between firms in a given pair of locations. The results (not shown) are similar, with a decline in f generating larger gains in trade and welfare than a cost-equivalent decline in κ .

6.3 Idiosyncratic fluctuations and aggregate dynamics

To study how shocks to firm-level fundamental characteristics translate into aggregate dynamics, I next consider the following counterfactual exercise. Suppose that at $t = 0$, the

³⁷In terms of parameter values, the comparison is between a 50% decline in κ versus an 18% decline in f .

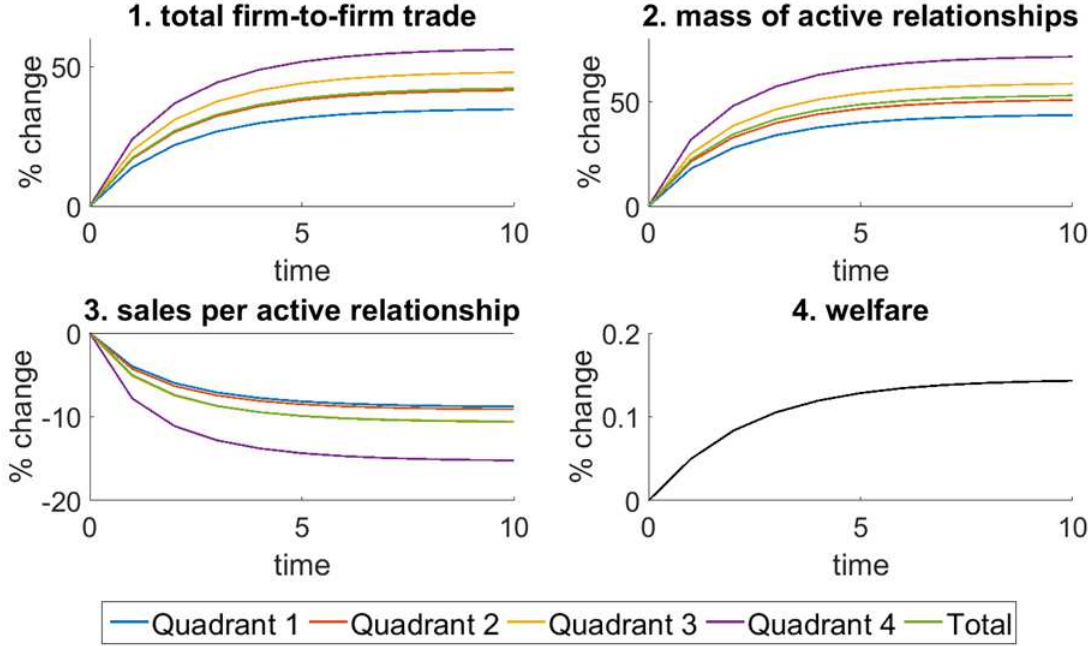


Figure 15: Responses of trade and welfare to global decline in relationship costs equivalent to 5% decline in trade costs

economy is initially in steady-state. Next, suppose that all firms receive an unexpected but permanent shock to their fundamental characteristics that leaves the distribution of states across firms unchanged. In particular, suppose that the post-shock fundamental productivities and qualities of a firm are given respectively by:

$$\log \hat{\phi} = \sqrt{1-s} \log \phi + \sqrt{s} \hat{\omega}_{\phi} \quad (6.4)$$

$$\log \hat{\delta} = \sqrt{1-s} \log \delta + \sqrt{s} \hat{\omega}_{\delta} \quad (6.5)$$

where the idiosyncratic shocks $\hat{\omega}_{\phi}$ and $\hat{\omega}_{\delta}$ are jointly normal with the same covariance matrix as $\log \phi$ and $\log \delta$, and where the parameter s captures the ratio of the shock variance to the variance of pre-shock firm states. Under this specification, it is straightforward to verify that the distribution of $\hat{\phi}$ and $\hat{\delta}$ across firms is identical to the pre-shock distribution of ϕ and δ . It is immediately obvious from this that in a model without costly relationships ($f = 0$), this shock would have no effect on the aggregate economy at all. In a world with sticky relationships, however, even such idiosyncratic fluctuations have aggregate effects.

As before, I assume that the shock hits the economy at $t = 0$ after all relationships have been set. Even though individual firm pairs cannot activate new relationships or terminate existing ones, however, the matching function still responds instantaneously to the fluctuation shock, not because firms adjust the identity of their trading partners, but because the

states of individual firms change. In particular, the matching function at date 0 adjusts instantaneously to:

$$\hat{m}_0(\hat{\chi}, \hat{\chi}') = \frac{\int_{S_\chi} \int_{S_\chi} m_{ss}(\chi, \chi') q(\hat{\chi}|\chi) q(\hat{\chi}'|\chi') dF_\chi(\chi) dF_\chi(\chi')}{\int_0^\infty \int_0^\infty q(\hat{\chi}|\chi) q(\hat{\chi}'|\chi') dF_\chi(\chi) dF_\chi(\chi')} \quad (6.6)$$

where q is the transition function between pre- and post-shock states implied by (6.4) and (6.5). Since the structural parameters of the model remain unchanged, the steady-state of the economy is the same as before the shock. However, firm relationships are “scrambled” by the idiosyncratic fluctuation in firm fundamental characteristics, and it takes time for the economy to return to its steady-state as firms readjust their relationships.

Figure 16 shows the responses of trade and welfare to the fluctuation shock for different values of the relative shock variance s . We observe that when s is very small, the fluctuation in firm states has little effect on aggregate quantities. However, as s starts to increase, the responses of trade and welfare grow quickly. With relative shock variances of 10% and 20%, aggregate trade falls immediately by about 10% and 30% respectively. Welfare also falls as firm states are scrambled, although again the magnitude of the effect is small. Furthermore, the economy only gradually returns to the steady-state, with the half-life of the trade and welfare responses being approximately two years.

This effect of idiosyncratic fluctuations on aggregate dynamics in the model can be considered complementary to the effects studied in Acemoglu et al (2012), where the authors examine the role of sector-level input-output structures in translating idiosyncratic shocks into aggregate fluctuations. In the model studied here, idiosyncratic shocks generate aggregate dynamics because the input-output structure of the economy at the firm level is endogenous, and responds to shocks that would have no aggregate effects in a model without relationship frictions.

6.4 The importance of rational expectations

Being able to solve for the model’s exact transition dynamics under rational expectations allows us to compare the model’s predictions to what would be obtained under the assumption that firms are myopic. As previously discussed, a common approach to modeling strategic network formation between atomistic agents is to assume that agents receive the chance to create or destroy links with finite probability, but that given the chance to change a relationship, the decision is made myopically based only on the static changes to the agent’s payoff.

To study the implications of myopia and therefore the importance of taking rational firm

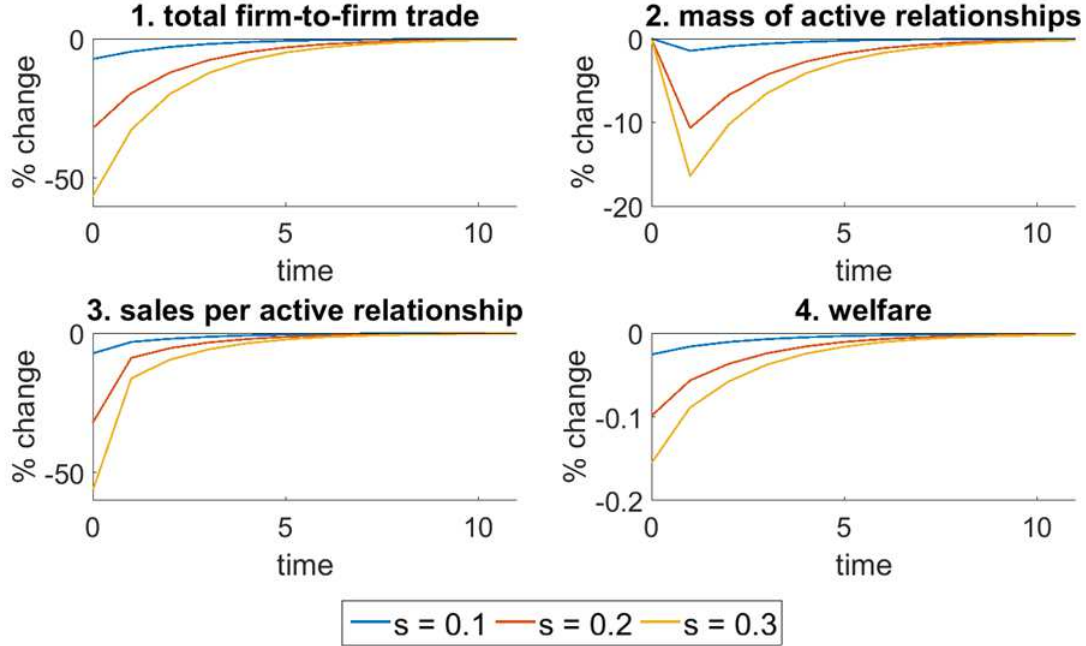


Figure 16: Responses of trade and welfare to idiosyncratic fluctuations in firm states

expectations into account, I study the model's predictions under the alternative assumption that the relationship acceptance function is given by (4.8) instead of (4.13), and compute the transition dynamics in response to the same global decline in variable trade costs discussed in section 6.1. Figure 17 shows the transition paths of trade and welfare (analogous to Figure 13), from which we observe the following. First, the short-run change in trade and welfare under both myopia and rational expectations is the same, because the matching function is held fixed. However, once firms are allowed to adjust the extensive margin of trade, the transition dynamics and the eventual steady-state of the model differ substantially under myopia relative to the rational expectations equilibrium. In particular, myopic firms form too many relationships relative to the rational expectations equilibrium, and welfare initially declines following the trade cost shock before increasing to a steady-state level that is about 25% lower than the rational expectations equilibrium steady-state. This divergence in both the qualitative as well as quantitative properties of the model under myopia clearly shows that taking agents' rational expectations into account can have a crucial impact on theoretical predictions.

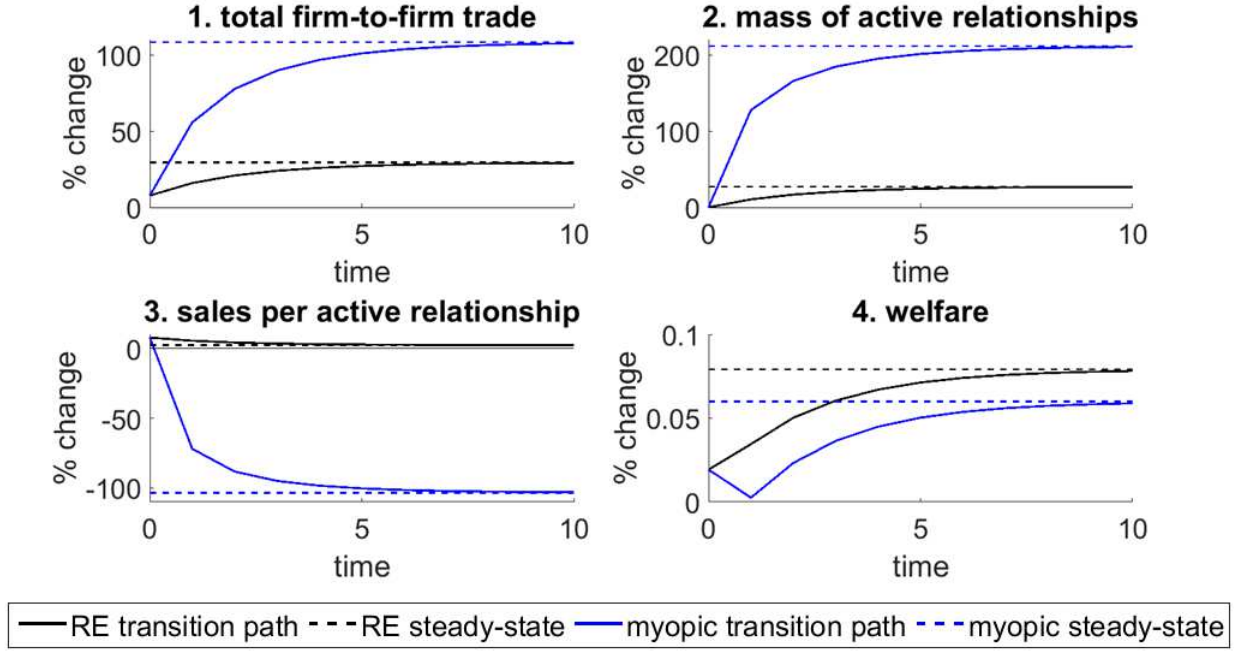


Figure 17: Responses of trade and welfare to 5% decline in global trade costs with myopic firms

6.5 Trade policy and sticky relationships

Given the central role of relationship stickiness in this paper, a natural policy question to ask is: can household welfare be improved by subsidies to the cost of forming relationships? To provide a first look into the effects of trade policy under sticky firm relationships, I consider the following stylized counterfactual. Suppose that for every relationship formed by a seller in each location, the policymaker in that location pays a fraction S_f of the fixed relationship cost, financed fully by an ad valorem import tax T_M . In other words, policymakers tax the intensive margin of trade to subsidize the extensive margin. Without transport costs ($\kappa = 0$), for example, the steady-state matching function under such a combination of policies would be:

$$m(\chi, \chi') = \tilde{m} \left[\frac{\Delta(\chi) \Phi(\chi') \Delta_H}{(1 - S_f)(1 + T_M)^{\sigma-1}} \right] \quad (6.7)$$

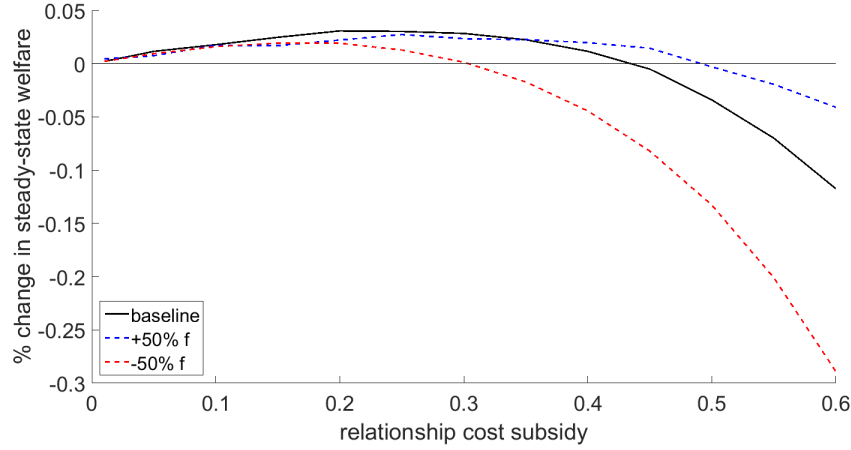


Figure 18: Effect of relationship cost subsidies on household welfare

where \tilde{m} is as defined by (4.22), and where the firm network characteristic functions are given by:

$$\Phi(\chi) = \phi^{\sigma-1} + \left[\frac{\alpha}{\mu(1+T_M)} \right]^{\sigma-1} \int_{S_\chi} m(\chi, \chi') \Phi(\chi') dF_\chi(\chi') \quad (6.8)$$

$$\Delta(\chi) = \mu^{-\sigma} \delta^{\sigma-1} + [\mu(1+T_M)]^{-\sigma} \alpha^{\sigma-1} \int_{S_\chi} m(\chi', \chi) \Delta(\chi') dF_\chi(\chi') \quad (6.9)$$

Balanced budgets in each location then require:

$$S_f L_f = T_M \bar{R} \quad (6.10)$$

where L_f is given by equation (4.14) and \bar{R} is total import expenditure:

$$\bar{R} = \left[\frac{\alpha}{\mu(1+T_M)} \right]^{\sigma-1} \Delta_H \int_{S_\chi} m(\chi, \chi') \Delta(\chi) \Phi(\chi') dF_\chi(\chi) dF_\chi(\chi') \quad (6.11)$$

Figure 18 shows the percentage change in household welfare across steady-states relative to the no-policy equilibrium for different values of S_f . Evidently, the model implies that firm relationship cost subsidies can be welfare improving even when financed by import taxes that distort the intensive margin of trade. This is a result of the fact that the market equilibrium is inefficient relative to the social planner's allocation, as characterized by Propositions 2 and 3.

7 Conclusion

This paper set out to study and quantify the effects of stickiness in firm-to-firm trading relationships on aggregate patterns of trade. The theoretical model developed to address these questions is able to adeptly match the majority of empirical moments relating to the distributions of relationships across firms, the correlation between firm connectivity and firm size, the assortativity of matching between firms, and the persistence of firm-to-firm relationships. Numerical estimation and counterfactual simulation of the model then suggest that firm-level relationship frictions matter for understanding patterns of aggregate trade in several key ways. First, endogenous adjustment of sticky firm relationships dynamically amplifies the response of trade and welfare to macroeconomic shocks. Second, subsidies to the cost of firm-level trade along the extensive margin can be a more cost-effective means of increasing aggregate trade and welfare than subsidies along the intensive margin. Third, idiosyncratic fluctuations at the firm-level can generate large and persistent aggregate trade dynamics when firm relationships are sticky. Finally, selection of trading relationships by profit-maximizing firms in the presence of relationship stickiness can be socially sub-optimal, with scope for welfare-improving subsidies to the formation of firm-level linkages.

The issues confronted in this paper also provide scope for future research. In particular, the model's inability to fit the matching distributions of firms at the lower-end of the revenue distribution suggest that more nuanced theory regarding the matching process may be needed to resolve this discrepancy. Extensions of the model, for instance, may consider the role of information in firm network formation, how such information propagates across firms, and how informational frictions may affect smaller versus large firms differentially. Furthermore, the empirical finding that larger firms tend to trade with partners that are closer by on average goes against not only predictions of the model developed in this paper, but also the standard intuition arising from heterogeneous-firm models of international trade that larger firms are more likely to export to more costly locations. This hints at a role for economic geography models in exploring the potentially-rich interaction between sticky firm relationships and the endogenous geographic locations of firms.

References

- Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdagalar, and Alireza Tahbaz-Salehi. 2012. “The Network Origins of Aggregate Fluctuations.” *Econometrica*, 80(5): 1977-2016.
- Alessandria, George, Horag Choi, and Kim Ruhl. 2015. “Trade Adjustment Dynamics and the Welfare Gains from Trade.” Mimeo: University of Rochester.
- Atalay, Enghin, Ali Hortaçsu, James Roberts, and Chad Syverson. 2011. “Network Structure of Production.” *Proceedings of the National Academy of Sciences*, 108(13): 5199-5202.
- Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. “New Trade Models, Same Old Gains?” *American Economic Review*, 102(1), 94-130.
- Aumann, Robert J, and Roger Myerson. 1988. “Endogenous Formation of Links Between Players and Coalitions: An Application of the Shapley Value,” in Roth, Al (ed.), *The Shapley Value*, Cambridge University Press.
- Bala, Venkatesh, and Sanjeev Goyal. 2000. “A Noncooperative Model of Network Formation.” *Econometrica*, 68(5): 1181-1229.
- Barabási, Albert-László, and Réka Albert. 1999. “Emergence of Scaling in Random Networks.” *Science*, 286 (5439): 509–512.
- Bernard, Andrew, Andreas Moxnes, and Yukiko Saito. 2015. “Production Networks, Geography, and Firm Performance.” Mimeo: Dartmouth University.
- Bernard, Andrew, Andreas Moxnes, and Karen Helene Ulltveit-Moe. 2015. “Two-sided Heterogeneity and Trade.” Mimeo: Dartmouth University.
- Burstein, Ariel, and Marc J Melitz. 2011. “Trade Liberalization and Firm Dynamics.” NBER Working Paper No. 16960.
- Cabral, Luís M B, and José Mata. 2003. “On the Evolution of the Firm Size Distribution: Facts and Theory.” *American Economic Review*, 93(4), 1075-1090.
- Calvo, Guillermo A. 1983. “Staggered Prices in a Utility-maximizing Framework.” *Journal of Monetary Economics*, 12, 383-398.
- Chaney, Thomas. 2015. “The Gravity Equation in International Trade: An Explanation.” Mimeo: Toulouse School of Economics.

- Chaney, Thomas. 2014. "The Network Structure of International Trade." *American Economic Review*, 104(11): 3600-3634.
- Constantini, James A, and Marc J Melitz. 2007. "The Dynamics of Firm-level Adjustment to Trade Liberalization," in Helpman, Elhanan, Dalia Marin, and Thierry Verdier (eds.), *The Organization of Firms in a Global Economy*, Cambridge, MA: Harvard University Press.
- Currarini, Sergio, Matthew O Jackson, and Paolo Pin. 2010. "Identifying the Roles of Race-based Choice and Chance in High School Friendship Network Formation" *Proceedings of the National Academy of Science*, 107: 4857-4861.
- Eaton, Jonathan, Marcela Eslava, Maurice Kugler, and James Tybout. 2007. "The Margins of Entry into Export Markets: Evidence from Colombia," in Helpman, Elhanan, Dalia Marin, and Thierry Verdier (eds.), *The Organization of Firms in a Global Economy*, Cambridge, MA: Harvard University Press.
- Eaton, Jonathan, Samuel Kortum, and Francis Kramarz. 2015. "Firm-to-firm Trade: Imports, Exports, and the Labor Market." Mimeo: Yale University.
- Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography and
- Erdős, Paul, and Alfréd Rényi. 1959. "On Random Graphs." *Publicationes Mathematicae*, 6: 290–297.
- Gabaix, Xavier. 2011. "The Granular Origins of Aggregate Fluctuations." *Econometrica*, 79(3): 733-772.
- Gartner, Inc. 2014a. "Market Share Analysis: Customer Relationship Management Software, Worldwide, 2014." <http://www.gartner.com/document/3045522>.
- . 2014b. "Market Share Analysis: Supply Chain Management and Procurement Software, Worldwide, 2014." <http://www.gartner.com/document/3044217>.
- Impullitti, Gianmario, Alfonso A Irarrazabal, and Luca David Oromolla. 2013. "A Theory of Entry into and Exit from Export Markets." *Journal of International Economics*, 90(1), 75-90.
- Jackson, Matthew O. 2011. "An Overview of Social Networks and Economic Applications," in Benhabib, Jess, Alberto Bisin, and Matthew O Jackson (eds.), *The Handbook of Social Economics*, North Holland Press: Amsterdam.

- Jackson, Matthew O. 2005. "A Survey of Models of Network Formation: Stability and Efficiency," in Demange, Gabrielle, and Myrna Wooders (eds.), *Group Formation in Economics: Networks, Clubs, and Coalitions*, Cambridge University Press: Cambridge, Chapter 1.
- Jackson, Matthew O, and Alison Watts. 2002. "The Evolution of Social and Economic Networks." *Journal of Economic Theory*, 106(2): 265-295.
- Jackson, Matthew O, and Asher Wolinsky. 1996. "A Strategic Model of Social and Economic Networks." *Journal of Economic Theory*, 71(1): 44-74.
- Kranton, Rachel E, and Deborah F Minehart. 2001. "A Theory of Buyer-Seller Networks." *American Economic Review*, 91(3): 485-508.
- Myerson, Roger. 1991. *Game Theory: Analysis of Conflict*, Harvard University Press: Cambridge, MA.
- Oberfield, Ezra. 2015. "Business Networks, Production Chains, and Productivity: A Theory of Input-Output Architecture." Mimeo: Princeton University.
- Rauch, James E. 1999. "Networks Versus Markets in International Trade." *Journal of International Economics*, 48: 7-35.
- Rossi-Hansberg, Esteban, and Mark L J Wright. 2007. "Establishment Size Dynamics in the Aggregate Economy." *American Economic Review*, 97(5), 1639-1666.
- Simonovska, Ina, and Michael Waugh. 2014a. "Trade Models, Trade Elasticities, and the Gains from Trade." NBER Working Paper No. 20495.
- . 2014b. "The Elasticity of Trade: Estimates and Evidence." *Journal of International Economics*, 92(1), 34-50.
- Stole, Lars A, and Jeffrey Zwiebel. 1996. "Intra-firm Bargaining Under Non-binding Contracts." *Review of Economic Studies*, 63: 375-410.
- Watts, Alison. 2001. "A Dynamic Model of Network Formation." *Games and Economic Behavior*, 34: 331-341.
- Watts, Duncan J, and Steven H Strogatz. 1998. "Collective Dynamics of 'Small-World' Networks." *Nature*, 393(6684): 440-442.

APPENDIX

A Computational Algorithms

A.1 Static algorithm

Given the matching function m , the static market equilibrium specified in Definition 1 can be solved for easily using the following algorithm.

1. Make initial guesses $\hat{\Phi}$ and $\hat{\Delta}$ for the network productivity and quality functions, and iterate on equations (3.19) and (3.20) until convergence.
2. Solve for Δ_H using equations (3.10) and (3.34).
3. Compute the allocation $\{l(\chi), X(\chi), x(\chi, \chi'), x_H(\chi)\}_{\chi \in S_\chi}$ using (3.28), (3.31), (3.33), and (3.37) respectively.

Since the functional equations (3.19) and (3.20) constitute contraction mappings with Lipschitz constants $\left(\frac{\alpha}{\mu}\right)^{\sigma-1}$ and $\frac{\alpha^{\sigma-1}}{\mu^\sigma}$ respectively, the iteration procedure in step 1 of the algorithm is guaranteed to converge at those rates. In practice, numerical solution of the model requires discretization of the state space S_χ into a mesh grid, of say $N_{grid} \times N_{grid}$ points. One can then solve for the functions $\Phi(\cdot)$ and $\Delta(\cdot)$ in step 1 at each point in the mesh grid, and then use bilinear interpolation to obtain numerical approximations of these functions as well as of the allocations $\{L(\chi), X(\chi), x(\chi, \chi'), x_H(\chi)\}$ for any desired value of $\chi \in S_\chi$.

A.2 Dynamic algorithm

I first describe the computational algorithm used to solve for the steady-state equilibrium specified in Definition 3, which is as follows.

1. Make initial guesses $\hat{\Phi}$ and $\hat{\Delta}_H \Delta$ for the network productivity function and the network quality function scaled by the household demand shifter.
2. Compute the implied profit function $\tilde{\pi}$ from equation (4.4).
3. Compute the implied matching and acceptance functions, \tilde{m} and \tilde{a} , from equations (4.3) and (4.11).
4. Compute the implied network productivity and quality functions, $\tilde{\Phi}$ and $\tilde{\Delta}$, from equations (3.19) and (3.20).

5. Compute the implied household demand shifter $\tilde{\Delta}_H$ from equations (3.34), (4.14), (4.15), and (4.16), and obtain the implied guess for the scaled network quality function, $\Delta_H \Delta = \tilde{\Delta}_H \tilde{\Delta}$.
6. Compute the residual $\mathcal{R} \equiv \max \{\mathcal{R}_\Phi, \mathcal{R}_\Delta\}$ where $\mathcal{R}_\Phi \equiv \max_{\chi \in S_\chi} |\hat{\Phi}(\chi) - \tilde{\Phi}(\chi)|$ and $\mathcal{R}_\Delta \equiv \max_{\chi \in S_\chi} |\Delta_H \Delta(\chi) - \tilde{\Delta}_H \tilde{\Delta}(\chi)|$; if $\mathcal{R} > \epsilon$ for some tolerance level ϵ , update the guesses for the network productivity and scaled quality functions according to $\hat{\Phi}' = \frac{\hat{\Phi} + \tilde{\Phi}}{2}$ and $\Delta_H \Delta'(\chi) = \frac{\Delta_H \Delta + \tilde{\Delta}_H \tilde{\Delta}}{2}$, and repeat from step 1 until $\mathcal{R} \leq \epsilon$.

I now discuss the computational algorithm used to solve for the model's transition dynamics as specified in Definition 2. Suppose that the matching and profit functions at date 0 are given by m_0 and π_0 respectively, and that the economy is not in steady-state. The goal is to solve for the model's transition path to the eventual steady-state characterized by the matching function denoted by m_{ss} . Note that given the matching function m_t , it is straightforward to solve for the static market equilibrium at date t using the algorithm discussed in section A.1. The challenge in solving the model's transition dynamics therefore lies in computing the matching function at date t given the matching function at date $t - 1$. As we see from equation (4.13), doing so while fully taking into account firm rational expectations requires solving for the profit functions $\{\pi_{t+s}\}_{s \geq 0}$. To accomplish this, I employ an algorithm that iterates on the path of profit functions $\{\pi_t\}_{t=1}^T$ for some value of T large enough such that the matching function at date T is close enough to the eventual steady-state matching function m_{ss} . Formally, the algorithm is as follows.

1. Make a guess \hat{T} for the number of periods that it takes for convergence to the steady-state.
2. Make an initial guess for the profit functions $\{\hat{\pi}_t\}_{t=2}^{\hat{T}}$ (e.g. $\hat{\pi}_t = \frac{1}{2}(\pi_0 + \pi_{ss})$ for all $t \in \{2, \dots, \hat{T}\}$).
3. At each date $t \in \{1, \dots, \hat{T}\}$, given \hat{m}_{t-1} (with $\hat{m}_0 = m_0$):
 - (a) Make initial guesses $\hat{\Phi}_t$ and $\Delta_H \Delta_t$ for the network productivity function and the network quality function scaled by the household demand shifter.
 - (b) Compute the implied profit function $\tilde{\pi}_t$ from equation (4.4).
 - (c) Compute the implied acceptance function \tilde{a}_t (4.11), setting $\pi_{t+s} = \tilde{\pi}_{t+s}$ for $s \in \{1, \dots, \hat{T} - t\}$ and $\pi_{t+s} = \pi_{ss}$ for $s > \hat{T} - t$.
 - (d) Compute the implied matching function \tilde{m}_t from equation (4.2).

- (e) Compute the implied network productivity and quality functions, $\tilde{\Phi}_t$ and $\tilde{\Delta}_t$, from equations (3.19) and (3.20).
 - (f) Compute the implied household demand shifter $\tilde{\Delta}_{H,t}$ from equations (3.34), (4.14), (4.15), and (4.16), and obtain the implied guess for the scaled network quality function, $\Delta_H \Delta_t = \tilde{\Delta}_{H,t} \tilde{\Delta}_t$.
 - (g) Compute the residual $\mathcal{R} \equiv \max \{ \mathcal{R}_\Phi, \mathcal{R}_\Delta \}$ where $\mathcal{R}_\Phi \equiv \max_{\chi \in S_\chi} \left| \hat{\Phi}_t(\chi) - \tilde{\Phi}_t(\chi) \right|$ and $\mathcal{R}_\Delta \equiv \max_{\chi \in S_\chi} \left| \Delta_H \hat{\Delta}_t(\chi) - \tilde{\Delta}_H \tilde{\Delta}_t(\chi) \right|$; if $\mathcal{R} > \epsilon$ for some tolerance level ϵ , update the guesses for the network productivity and scaled quality functions according to $\hat{\Phi}'_t(\chi) = \frac{1}{2} \left[\hat{\Phi}_t(\chi) + \tilde{\Phi}_t(\chi) \right]$ and $\Delta_H \hat{\Delta}'_t(\chi) = \frac{1}{2} \left[\Delta_H \hat{\Delta}_t(\chi) + \tilde{\Delta}_H \tilde{\Delta}_t(\chi) \right]$, and repeat from step (a) until $\mathcal{R} \leq \epsilon$, then set $\hat{m}_t = \tilde{m}_t$.
4. Compute the residual $\mathcal{R}_\pi \equiv \max_{t \in \{2, \dots, \hat{T}\}} \max_{(\chi, \chi') \in S_\chi^2} \left| \hat{\pi}_t(\chi, \chi') - \tilde{\pi}_t(\chi, \chi') \right|$; if $\mathcal{R}_\pi > \epsilon_\pi$ for some tolerance level ϵ_π , update the guesses for the profit functions according to $\hat{\pi}'_t = \frac{\hat{\pi}_t + \tilde{\pi}_t}{2}$ for all $t \in \{2, \dots, \hat{T}\}$, and repeat from step 2 until $\mathcal{R}_\pi \leq \epsilon$.
 5. Compute the residual $\mathcal{R}_m \equiv \max_{(\chi, \chi') \in S_\chi^2} \left| \hat{m}_{\hat{T}}(\chi, \chi') - m_{ss}(\chi, \chi') \right|$; if $\mathcal{R}_m > \epsilon_m$ for some tolerance level ϵ_m , increment \hat{T} and repeat from step 1.

As in solving for the static market equilibrium, numerical solution of the dynamic market equilibrium requires discretization of the state space S_χ into a mesh grid of $N_{grid} \times N_{grid}$ points, and bilinear interpolation can then be used to obtain numerical approximations of firm-level equilibrium variables off the grid points. Note that given the guess of future profit functions, step 3 of the algorithm has the same computational complexity as solving for the model's steady-state, and this part of the computation can be sped up by using the terminal guesses at the previous date when initializing the guesses for the network characteristic functions in step 3(a). Furthermore, upon increasing the guess for \hat{T} to $\hat{T} + 1$ in step 5, the new guess for the profit functions up to date \hat{T} used in step 2 can be set at the previous terminal guesses for the profit functions up to that date, which also speeds up the computation.

With a grid size of $N_{grid} = 20$ and tolerance levels $\epsilon = \epsilon_\pi = \epsilon_m = 10^{-4}$, executing the steady-state algorithm typically takes around 30 seconds, while solving for a transition path such as those discussed in the main text typically takes about one hour on a standard computer. Since estimation of the model's parameters only requires solving for steady-state equilibria, the complexity of executing the dynamic algorithm does not factor into the tractability of estimating the model.

B Static and Dynamic Efficiency

B.1 Static efficiency

To characterize the efficiency of the static market equilibrium, I compare the resulting allocation with the allocation that would be chosen by a social planner whose goal is to maximize household welfare subject to the production technology and market clearing constraints. Given the matching function m , the social planner chooses the allocation $\mathcal{A} \equiv \left\{ l(\chi), X(\chi), \{x(\chi, \chi')\}_{\chi' \in S_\chi}, x_H(\chi) \right\}_{\chi \in S_\chi}$ according to :

$$U = \max_{\mathcal{A}} \left[\int_{S_\chi} [\delta x_H(\chi)]^{\frac{\sigma-1}{\sigma}} dF_\chi(\chi) \right]^{\frac{\sigma}{\sigma-1}}$$

subject to the following constraints:

$$X(\chi) = \left[[\phi l(\chi)]^{\frac{\sigma-1}{\sigma}} + \int_{S_\chi} m(\chi, \chi') [\alpha x(\chi, \chi')]^{\frac{\sigma-1}{\sigma}} dF_\chi(\chi') \right]^{\frac{\sigma}{\sigma-1}} \quad (\text{B.1})$$

$$X(\chi) = x_H(\chi) + \int_{S_\chi} m(\chi', \chi) x(\chi', \chi) dF_\chi(\chi') \quad (\text{B.2})$$

$$\int_{S_\chi} l(\chi) dF_\chi(\chi) = L - L_f \quad (\text{B.3})$$

where $L_f = f \int_{S_\chi} \int_{S_\chi} m(\chi, \chi') dF_\chi(\chi) dF_\chi(\chi')$ is taken as given.

Denoting the Lagrange multipliers on constraints (B.2) and (B.3) by $\left(\frac{U}{\Delta_H}\right)^{\frac{1}{\sigma}} \eta(\chi) f_\chi(\chi)$ and $\left(\frac{U}{\Delta_H}\right)^{\frac{1}{\sigma}}$ respectively, the first-order conditions for the planner's problem can be expressed as:

$$x_H(\chi) = \Delta_H \delta^{\sigma-1} \eta(\chi)^{-\sigma} \quad (\text{B.4})$$

$$l(\chi) = X(\chi) \eta(\chi)^\sigma \phi^{\sigma-1} \quad (\text{B.5})$$

$$x(\chi, \chi') = X(\chi) \eta(\chi)^\sigma \alpha^{\sigma-1} \eta(\chi')^{-\sigma} \quad (\text{B.6})$$

Substituting these equations into (B.1) and (B.2), we get:

$$\Phi(\chi) = \phi^{\sigma-1} + \alpha^{\sigma-1} \int_{S_\chi} m(\chi, \chi') \Phi(\chi') dF_\chi(\chi') \quad (\text{B.7})$$

$$\Delta(\chi) = \delta^{\sigma-1} + \alpha^{\sigma-1} \int_{S_\chi} m(\chi', \chi) \Delta(\chi') dF_\chi(\chi') \quad (\text{B.8})$$

where $\Phi(\chi) \equiv \eta(\chi)^{1-\sigma}$ and $\Delta(\chi) \equiv \frac{1}{\Delta_H} X(\chi) \eta(\chi)^\sigma$.

Note that equations (B.4)-(B.8) are identical to equations (3.2), (3.7), (3.8), (3.19), and (3.20) respectively only when $\mu = 1$. This tells us that the static market equilibrium allocation is identical to the planner's allocation if and only if the markups charged by all firms are equal to one. With a finite elasticity of substitution σ , the static market equilibrium is therefore inefficient relative to the planner's allocation because of the monopoly markup distortion.

B.2 Dynamic efficiency

To study the efficiency properties of the dynamic market equilibrium, we consider the problem of a social planner that chooses the set of relationships to activate and terminate at each date so as to maximize the present discounted value of household welfare, subject to the same dynamic frictions faced by firms in the market equilibrium. From the results in section B.1, we know that given the matching function m_t and the total mass of labor used to pay relationship costs $L_{f,t}$, household utility at date t under the planner's optimal allocation can be written as:

$$U_t = (L - L_{f,t}) \mathcal{C}_t \quad (\text{B.9})$$

where \mathcal{C}_t measures the total connectivity of the static production network:

$$\mathcal{C}_t \equiv \left[\int_{S_\chi} \int_{S_\chi} \left[\sum_{d=0}^{\infty} \alpha^{d(\sigma-1)} m_t^{(d)}(\chi, \chi') \right] (\delta \phi')^{\sigma-1} dF_\chi(\chi) dF_\chi(\chi') \right]^{\frac{1}{\sigma-1}} \quad (\text{B.10})$$

$$= \left[\int_{S_\chi} \Phi_t(\chi) \delta^{\sigma-1} dF_\chi(\chi) \right]^{\frac{1}{\sigma-1}} \quad (\text{B.11})$$

$$= \left[\int_{S_\chi} \Delta_t(\chi) \phi^{\sigma-1} dF_\chi(\chi) \right]^{\frac{1}{\sigma-1}} \quad (\text{B.12})$$

and Φ_t and Δ_t are given by the date t equivalents of equations (B.7) and (B.8) respectively.

To study the planner's dynamic optimization problem, let $V_t(m_{t-1})$ denote the present value of discounted household utility at date t under the planner's optimal dynamic allocation when the matching function in the previous period is given by m_{t-1} . At each date t , the planner's choice about which relationships to activate and terminate is equivalent to a choice over the values $\{\xi_{max,t}(\chi, \chi')\}_{(\chi, \chi') \in S_\chi^2}$, where $\xi_{max,t}(\chi, \chi')$ specifies the maximum value of the idiosyncratic relationship cost shock component for which $\chi - \chi'$ firm pair relationships

are accepted. The Bellman equation for the planner's problem can therefore be written as:

$$V_t(m_{t-1}) = \max_{\{\xi_{max,t}(\chi, \chi')\}_{(\chi, \chi') \in S_\chi^2}} [U_t + \beta V_{t+1}(m_t)] \quad (\text{B.13})$$

where the maximization is subject to $\xi_{max,t}(\chi, \chi') \geq 0$ for all t and $(\chi, \chi') \in S_\chi^2$, as well as the following constraints:

$$U_t = (L - L_{f,t}) C_t \quad (\text{B.14})$$

$$C_t = \left[\int_{S_\chi} \Phi_t(\chi) \delta^{\sigma-1} dF_\chi(\chi) \right]^{\frac{1}{\sigma-1}} \quad (\text{B.15})$$

$$\Phi_t(\chi) = \phi^{\sigma-1} + \alpha^{\sigma-1} \int_{S_\chi} m_t(\chi, \chi') \Phi_t(\chi') dF_\chi(\chi') \quad (\text{B.16})$$

$$L_{f,t} = f \int_{S_\chi} \int_{S_\chi} \left[\nu m_{t-1}(\chi, \chi') + (1-\nu) \int_0^{\xi_{max,t}(\chi, \chi')} \xi dF_\xi(\xi) \right] dF_\chi(\chi) dF_\chi(\chi') \quad (\text{B.17})$$

$$m_t(\chi, \chi') = \nu m_{t-1}(\chi, \chi') + (1-\nu) F_\xi[\xi_{max,t}(\chi, \chi')] \quad (\text{B.18})$$

For brevity, denote $\xi_{max,t}^* \equiv \xi_{max,t}(\chi^*, \chi^{*'})$ and $m_t^* \equiv m_t(\chi^*, \chi^{*'})$ for a given firm pair $(\chi^*, \chi^{*'})$. The first step in solving the dynamic planner's problem is to find an expression for the derivative of U_t with respect to $\xi_{max,t}^*$. First, we differentiate (B.17) with respect to $\xi_{max,t}^*$ to get:

$$\frac{dL_{f,t}}{d\xi_{max,t}^*} = (1-\nu) H(\chi^*, \chi^{*'}, \xi_{max,t}^*) f \xi_{max,t}^* \quad (\text{B.19})$$

where $H(\chi, \chi', \xi) \equiv f_\chi(\chi) f_\chi(\chi') f_\xi(\xi)$ is the product of three probability densities. Next, differentiating (B.18) for $(\chi, \chi') = (\chi^*, \chi^{*'})$ with respect to $\xi_{max,t}^*$ gives:

$$\frac{dm_t^*}{d\xi_{max,t}^*} = (1-\nu) f_\xi(\xi_{max,t}^*) \quad (\text{B.20})$$

Differentiating the functional equation (B.8) with respect to $\xi_{max,t}^*$, we then obtain:

$$\frac{d\Phi_t(\chi)}{d\xi_{max,t}^*} = \frac{d\Phi_t(\chi)}{dm_t^*} \frac{dm_t^*}{d\xi_{max,t}^*} \quad (\text{B.21})$$

$$= (1-\nu) f_\xi(\xi_{max,t}^*) \left[\alpha^{\sigma-1} \Phi_t(\chi^{*'}) \mathbf{1}_{\chi^*}(\chi) + \alpha^{\sigma-1} \int_{S_\chi} m_t(\chi, \chi') \frac{d\Phi_t(\chi')}{d\xi_{max,t}^*} dF_\chi(\chi') \right] \quad (\text{B.22})$$

$$= (1-\nu) H(\chi^*, \chi^{*'}, \xi_{max,t}^*) \left[\sum_{d=0}^{\infty} \alpha^{d(\sigma-1)} m_t^{(d)}(\chi, \chi^*) \right] \alpha^{\sigma-1} \Phi(\chi^{*'}) \quad (\text{B.23})$$

where $\mathbf{1}_{\chi^*}(\chi)$ is the indicator function that equals 1 if $\chi = \chi^*$ and 0 otherwise. (Note that equation (B.23) summarizes the effect of a change in the mass of connections between

$\chi^* - \chi^{*'}$ firm pairs on the network productivities of all firms that are downstream of χ^* firms.) Differentiating equation (B.14) with respect to $\xi_{max,t}^*$ and using (B.19) and (B.23), we then get:

$$\frac{dU_t}{d\xi_{max,t}^*} = (1 - \nu) H\left(\chi^*, \chi^{*'}, \xi_{max,t}^*\right) \mathcal{C}_t \left[\tilde{\pi}_t\left(\chi^*, \chi^{*'}\right) - f \xi_{max,t}^* \right] \quad (\text{B.24})$$

where we have defined:

$$\tilde{\pi}_t\left(\chi^*, \chi^{*'}\right) \equiv \frac{\alpha^{\sigma-1}}{\sigma-1} \Delta_{H,t} \Delta_t\left(\chi^*\right) \Phi_t\left(\chi^{*'}\right) \quad (\text{B.25})$$

Note that conditional on the network characteristic functions, $\tilde{\pi}_t$ differs from the profit function π_t in the dynamic market equilibrium (given by equation (4.4)) only by a constant fraction $\mu^{-\sigma}$.

The next step in solving the planner's problem is to derive an expression for the derivative of the continuation value $V_{t+1}(m_t)$ with respect to $\xi_{max,t}^*$. First, we note that:

$$\frac{dV_{t+1}}{d\xi_{max,t}^*} = (1 - \nu) f_\xi\left(\xi_{max,t}^*\right) \frac{dV_{t+1}}{dm_t^*} \quad (\text{B.26})$$

The envelope condition then gives us:

$$\frac{dV_{t+1}}{dm_t^*} = \frac{dU_{t+1}}{dm_t^*} + \beta \nu \frac{dV_{t+2}}{dm_{t+1}^*} \quad (\text{B.27})$$

Using the same approach as in solving for $\frac{dU_t}{d\xi_{max,t}^*}$, it is straightforward to show that:

$$\frac{dU_{t+1}}{dm_t^*} = \nu f_\chi\left(\chi^*\right) f_\chi\left(\chi^{*'}\right) \mathcal{C}_{t+1} \left[\tilde{\pi}_{t+1}\left(\chi^*, \chi^{*'}\right) - f \right] \quad (\text{B.28})$$

Combining (B.26), (B.27) and (B.28), we then obtain:

$$\frac{dV_{t+1}}{d\xi_{max,t}^*} = \nu(1 - \nu) H\left(\chi^*, \chi^{*'}, \xi_{max,t}^*\right) \sum_{s=0}^{\infty} (\beta \nu)^s \mathcal{C}_{t+1+s} \left[\tilde{\pi}_{t+1+s}\left(\chi^*, \chi^{*'}\right) - f \right] \quad (\text{B.29})$$

Piecing together equations (B.24) and (B.29), we can finally write the first-order condition with respect to $\xi_{max,t}(\chi, \chi')$ in the planner's problem as:

$$\xi_{max,t}(\chi, \chi') = \max \left\{ \frac{\tilde{\pi}_t(\chi, \chi')}{f} + \sum_{s=1}^{\infty} (\beta \nu)^s \left(\frac{\mathcal{C}_{t+s}}{\mathcal{C}_t} \right) \left[\frac{\tilde{\pi}_{t+s}(\chi, \chi')}{f} - 1 \right], 0 \right\} \quad (\text{B.30})$$

C Model Extensions

C.1 Multiple industries

To introduce multiple industries into the model, we can partition the set of firms Ω into N subsets of equal mass and allow the input suitability parameter α to vary across industry pairs. This variation in input suitability captures how “upstream” or “downstream” one industry is relative to another, and allows the model to match industry-level input-output tables. Assuming that the distribution of fundamental firm characteristics is the same in all industries and denoting by α_{uv} the suitability of inputs from industry v for use in producing goods in industry u , the analogs of equations (3.19) and (3.20) in steady-state are then:

$$\Phi_u(\chi) = \phi^{\sigma-1} + \frac{1}{N} \sum_{v=1}^N \left(\frac{\alpha_{uv}}{\mu} \right)^{\sigma-1} \int_{S_\chi} m_{uv}(\chi, \chi') \Phi_v(\chi') dF_\chi(\chi') \quad (\text{C.1})$$

$$\Delta_u(\chi) = \mu^{-\sigma} \delta^{\sigma-1} + \frac{1}{N} \sum_{v=1}^N \mu^{-\sigma} \alpha_{vu}^{\sigma-1} \int_{S_\chi} m_{vu}(\chi', \chi) \Delta_v(\chi') dF_\chi(\chi') \quad (\text{C.2})$$

where now the network productivity and quality functions Φ_u and Δ_u are industry-specific, and the matching function m_{uv} is industry-pair-specific. The matching function for each industry pair can in turn be computed using the corresponding version of equation (4.11).

Given the network characteristic functions for each industry and the matching function for each industry pair, we can then use equation (3.32) to calculate input-output shares. The share of industry u 's inputs that are sourced from industry v , for example, is given by:

$$S_{uv}^I = \frac{\alpha_{uv}^{\sigma-1} \int_{S_\chi} \int_{S_\chi} m_{uv}(\chi, \chi') \Delta_u(\chi) \Phi_v(\chi') dF_\chi(\chi) dF_\chi(\chi')}{\sum_{w=1}^N \alpha_{uw}^{\sigma-1} \int_{S_\chi} \int_{S_\chi} m_{uw}(\chi, \chi') \Delta_u(\chi) \Phi_w(\chi') dF_\chi(\chi) dF_\chi(\chi')} \quad (\text{C.3})$$

while the share of industry u 's intermediate sales that accounted for by customers in industry v is:

$$S_{uv}^O = \frac{\alpha_{vu}^{\sigma-1} \int_{S_\chi} \int_{S_\chi} m_{vu}(\chi, \chi') \Delta_v(\chi) \Phi_u(\chi') dF_\chi(\chi) dF_\chi(\chi')}{\sum_{w=1}^N \alpha_{wu}^{\sigma-1} \int_{S_\chi} \int_{S_\chi} m_{wu}(\chi, \chi') \Delta_w(\chi) \Phi_u(\chi') dF_\chi(\chi) dF_\chi(\chi')} \quad (\text{C.4})$$

C.2 Customer-supplier Bargaining and Cost-sharing

In this section, I discuss how the model's assumptions can be modified to allow for a more general split of both the relationship surplus and the relationship fixed cost between the buying and selling firm.

First, note that without loss of generality, we can write the prices charged by a χ -firm to the household and to a potential χ' -buyer as markups $\mu_H(\chi)$ and $\mu(\chi, \chi')$ respectively over

the seller's marginal cost $\eta(\chi)$. The system of equations defining the network productivity and quality functions in the static market equilibrium can then be written as:

$$\Phi(\chi) = \phi^{\sigma-1} + \alpha^{\sigma-1} \int_{S_\chi} \mu(\chi, \chi')^{1-\sigma} m(\chi, \chi') \Phi(\chi') dF_\chi(\chi') \quad (\text{C.5})$$

$$\Delta(\chi) = \mu_H(\chi)^{-\sigma} \delta^{\sigma-1} + \alpha^{\sigma-1} \int_{S_{\chi|chi}} \mu(\chi', \chi)^{-\sigma} m(\chi', \chi) \Delta(\chi') dF_\chi(\chi') \quad (\text{C.6})$$

while the profit that a χ -firm makes from its sales to a χ' -firm is given by:

$$\pi(\chi, \chi') = \mu(\chi, \chi')^{-\sigma} [\mu(\chi, \chi') - 1] \alpha^{\sigma-1} \Delta_H \Delta(\chi) \Phi(\chi') \quad (\text{C.7})$$

Note also that the total profit of a χ -firm can be written as:

$$\pi(\chi) = \Delta_H \hat{\Delta}_i \Phi_i \quad (\text{C.8})$$

where

$$\hat{\Delta}_i \equiv \left[\mu_H(\chi)^{-\sigma} [\mu_H(\chi) - 1] \delta^{\sigma-1} + \alpha^{\sigma-1} \int_{S_\chi} \mu(\chi', \chi)^{-\sigma} [\mu(\chi', \chi) - 1] \Delta(\chi') dF_\chi(\chi') \right] \quad (\text{C.9})$$

depends only on variables relating to firm i 's customers.

Now suppose that instead of assuming a market structure characterized by monopolistic competition, we assume that firms take the markups charged by all other firms as given, and that the markup $\mu(\chi, \chi')$ is chosen to maximize the product $[v^C(\chi, \chi')]^\theta [v^S(\chi, \chi')]^{1-\theta}$. In other words, buyers and sellers engage in bilateral Nash bargaining (which we will soon see is equivalent to multilateral Nash bargaining in the static model), with $v^C(\chi, \chi')$ and $v^S(\chi, \chi')$ denoting the surplus to the customer and supplier respectively of the relationship between a χ -buyer and a χ' -seller. The parameter $\theta \in [0, 1]$ measures the bargaining power of the customer relative to the supplier.

From (C.5), (C.7), and (C.8), the surplus values can be written as:

$$v^C(\chi, \chi') = \mu(\chi, \chi')^{1-\sigma} \alpha^{\sigma-1} \hat{\Delta}(\chi) \Phi(\chi') \quad (\text{C.10})$$

$$v^S(\chi, \chi') = \mu(\chi, \chi')^{-\sigma} [\mu(\chi, \chi') - 1] \alpha^{\sigma-1} \Delta(\chi) \Phi(\chi') \quad (\text{C.11})$$

Note that $\hat{\Delta}(\chi)$ and $\Delta(\chi)$ depend only interactions between the χ -buyer and its own customers, while $\Phi(\chi')$ depends only on interactions between the χ' -seller and its own suppliers. In other words, because of the CES structure of the production function, the surplus of the relationship between a χ -buyer and a χ' -seller is independent of the interactions between

the buying firm and its other suppliers, and is also independent of the interactions between the selling firm and its other customers. As a result, bilateral Nash bargaining is equivalent in the model to the multilateral generalization of Nash bargaining proposed in Stole and Zwiebel (1996).

From equations (C.10) and (C.11), it is then straightforward to verify that firms again charge a constant markup over marginal cost, but that this markup is now given by:

$$\mu = \frac{\sigma - \theta}{\sigma - 1} \quad (\text{C.12})$$

Note that when all bargaining power resides with the supplier ($\theta = 0$), the markup charged is the same as that under monopolistic competition, whereas when all bargaining power resides with the buyer ($\theta = 1$), the markup is the same as that under perfect competition. In general, we have $\mu \in [1, \frac{\sigma}{\sigma-1}]$. Furthermore, if we assume that firms sell to households indirectly via a unit continuum of retailers that produce differentiated varieties of a retail good, and that sales between producers and retailers are characterized by the same bargaining process, then the same analysis as above can be used to rationalize markups for final sales that are also constant and given by (C.12).

We can also allow for a more general split of relationship costs between buyers and sellers by assuming that the buying firm pays a constant fraction b of the fixed cost in each relationship. In this case, whether a potential relationship is mutually desired by both buyer and seller depends on how the respective cost shares compare to the surplus values (C.10) and (C.11). Supposing that firms' pricing decisions remain characterized by constant markups equal to μ , it is straightforward to verify that a relationship is mutually desirable if and only if profits from that relationship are at least greater than an effective fixed cost given by:

$$f_{eff} \equiv f \max \{b\mu, 1 - b\} \quad (\text{C.13})$$

Note that the effective fixed cost is minimized when $b = \frac{1}{\mu+1}$. This implies that relationships are more likely to form if selling firms pay a larger share of relationship costs whenever the markups that they charge are also higher.

Through these additional assumptions, the model therefore allows for richer variation in inter-firm markups and effective relationship costs. It is important to point out, however, that these assumptions about bargaining and cost-sharing become much more restrictive once embedded in the dynamic model with endogenous network formation. For example, the characterization of the dynamic model discussed in the main text remains valid with buyer-supplier Nash bargaining only if we rule out repeated bargaining between potential

buyer-supplier pairs. The possibility of transfers between buyers and sellers also needs to be ruled out once the fixed relationship cost is taken into account. Furthermore, as discussed in the main text, once the buying firm pays a positive share of the relationship cost, constant markup pricing is not necessarily optimal for all firms in the dynamic model. For these reasons, I retain monopolistic competition as the assumed market structure and set $b = 0$ in the main model, and leave development of richer models of bargaining and cost-sharing under the setting of sticky relationships for future work.