Supply Chain Dynamics with Search Frictions

PRELIMINARY AND INCOMPLETE

Lauri Esala[1]
UPF and BSE

June 26, 2023

Abstract

An economy's production network is, in essence, a collection of relationships between customers and suppliers. Micro-level evidence indicates that these buyer-supplier linkages between firms are persistent, time-consuming to form, and difficult to replace if lost. So far, the macroeconomic literature lacks dynamic models of endogenous production networks with realistic frictions to adjusting linkages. This paper proposes such a framework, where the key adjustment cost is a search friction. I build a continuous-time model of a multi-stage production process, where firms look for both buyers and sellers - for now, assuming that a production chain needs only one firm in each stage. Tractability is then ensured by assuming that all forms of uncertainty follow a jump process, so only one event is possible at each instant. In this paper, I explore the steady-state implications of between-firm matching in this environment. Firms seek to join chains with high flow profit, but also take into account forward-looking considerations such as the fragility and growth potential of a chain. I also provide an example of aggregate dynamics in the model, studying the impulse response of aggregate variables to destroying buyer-supplier connections.

1 Introduction

Production in modern economies is organized into increasingly complex networks of producers and intermediaries. Final goods sold to consumers and firms are now typically combined from inputs that themselves require multiple stages of elaboration. While some intermediate goods needed in production are readily bought on the marketplace, considerations such as the need for reliable supplies of highly customised inputs call for the formation of long-term buyer-supplier relationships, which can be a time-consuming process. Indeed, Miyauchi (2021) finds that firms take time to replace unexpectedly lost key suppliers, while Huneeus (2020) shows that firm-to-firm trade linkages react sluggishly to shocks. The widespread supply chain disturbances that plagued the global economy after the COVID-19 pandemic made it clear that the the destruction and re-creation of these relationships can also play a critical role in the business cycle.

While economists have formulated static theories of supply chain formation and fragility (e.g. Elliott et al., 2022), many macroeconomic questions – such as how the churn in relationships

[1]Email: lauri.esala@upf.edu. Telephone: +358 40 8484662. Address: Ramon Trias Fargas, 25-27, 08005 Barcelona, Spain. I thank my doctoral advisors Edouard Schaal and Isaac Baley for their invaluable guidance, and Mishel Ghassibe and Andrea Sy for useful comments and discussions.

reacts to shocks and affects the depth of a recession and the pace of the subsequent recovery – call for a dynamic perspective. However, tractable models of forward-looking network formation with realistic frictions have proved hard to formulate due to the massive number of possible network configurations that forward-looking firms have to consider in such models. Indeed, firms need to evaluate each configuration not only today, but given all possible paths of shocks in the future. Furthermore, adding aggregate shocks to such a model, which could be considered a prerequisite for business cycle analysis, would typically imply that the entire network becomes a state variable for each firm.

This paper takes a first step towards filling the gap. In particular, I propose a model of multi-stage production where firms' choice of trading partners is subject to search frictions. In the current baseline version of model, production is organised in linear supply chains, where each firm matches with a maximum of one buyer and one seller (as, for example, in Antràs and de Gortari, 2020). Each intermediate goods firm is endowed with a stage-specific Cobb-Douglas technology combining labor and a single input that can either be produced in-house, or bought from an upstream supplier. Buyer-supplier linkages are formed using random search, while allowing for chains of different lengths to emerge. Matched firms continuously search for better trade partners, similarly to "on-the-job search" in labor market models. Pricing is kept simple, assuming each firm matched with a customer earns a mark-up on their unit cost.

In steady state, the state variable of the firm becomes the vector of productivities of all chain members. The continuous-time environment is specified in such a way that only one event – a match, a separation, or a jump productivity shock – may alter this state at each instant, implying that the model remains quite tractable. However, the ultimate goal of my project is to provide a full framework in which to study the dynamics of production chains and their interaction with the business cycle. White the full-information model is easy to solve when chains are relatively short and firms' productivites can take on only a small set of values, the cost grows quickly when enlarging the state space. To get around this problem, I am working on a solution method to be applied in larger and more realistic models. Inspired by Krusell et al. (1998), I assume that firms do not use the entire chain's state when forecasting their environment, and instead base their prognostications on a limited set of "moments".

In the preliminary work documented here, I explore the steady state implications of firm-to-firm search across multiple production stages. The productivity ladder and "on-the-job" search imply that all but the most productive firms in the most productive chains may want to switch to more profitable customer and supplier relationships, if given the chance. However, this freedom to switch chains extends to all of the chain's members. Therefore, an intermediate goods firm not only faces the threat of being replaced by its own buyers and sellers, but more distant members of the chain may also get "poached" by other chains. This implies a need for firms consider the *fragility* of the production chain, in addition to its flow profitability, when making matching decisions. For example, a firm may sometimes optimally turn down a match with a more lucrative customer, if it deems the new chain more unstable than the current one.

On the other hand, firms also consider a chain's *growth potential*. For example, a highly productive upstream firm may find it advantageous to join a production chain with an inefficient final goods producer, if its other suppliers are also highly productive. The logic is that the new chain will likely be able replace the final goods firm with a more efficient one, and start earning much higher profits. In both examples, firms consider a notion of mismatch in chains when making their matching decisions – in particular, they account for whether any one of the chain's members is likely to leave or be replaced.

The resulting equilibrium distribution of production chains has features of sorting along productivity, with the highest-productivity firms in particular tending to stick together in the same production chains. Other supply chain arrangements than the most profitable ones remain sensitive to endogenous separations, since their higher-productivity members can and do leave to join other chains. This affects the lowest-productivity firms in particular, who experience frequent spells of being unmatched in my illustrative main calibration. I also present results from a pair of alternative calibrations to shed light on how the strength of the steady-state mechanisms described above depends on model parameters. Finally, I also display some preliminary work on the effect son aggregate shocks. In particular, in response to an MIT shock that destroys a randomly chosen set of supply chain linkages, the economy recovers slowly as firms rebuild their supply chains.

Related literature. My paper contributes to a small but growing literature on dynamic macroeconomic models in which the production network evolves endogenously. Seminal works by Lim (2018) and Acemoglu and Azar (2020) propose highly tractable frameworks where linkages can be adjusted frictionlessly, implying that firms' decisions are static and their response to shocks is immediate. Other papers, such as Taschereau-Dumouchel (2020), take the set of firms and their input-output relationships as given, but endogenize the set of active firms[2]. An older literature, exemplified by Atalay et al. (2011) and Carvalho and Voigtländer (2014), uses tools from network science to build models where the input-output structure evolves in a frictional but backward-looking manner[3].

However, micro-level empirical evidence on firm-to-firm trade suggests that building customer and supplier relationships is not a frictionless process. For example, firms that unexpectedly lose a key supplier only find a replacement over time, and face worse economic outcomes in the meantime (Miyauchi, 2021). Multiple studies have shown that firm-to-firm trade linkages are highly persistent over time (e.g. Huneeus, 2020) and that firms invest in and protect these relationships (Ersahin et al., 2022; Heise, 2021).

To my knowledge, only Huneeus (2020) has constructed a dynamic theory where breaking and forming buyer-supplier relationships is costly[4]. In his model, which builds on Lim (2018), continuum sets of firms match with each other to create a roundabout input-output network.

[2] Acemoglu and Tahbaz-Salehi (2020) also considers the entry decisions of firms given a fixed technological compatibility matrix, but in a completely static framework.

[3] There are also several purely static theories of how production networks form and respond to shocks. Some of the key work in this vein includes Elliott et al. (2022) and

[4] Ghassibe (2021) extends the model of Acemoglu and Azar (2020) by introducing nominal rigidities, which makes firms' linkage choices endogenously persistent.

Suppliers choose their customers subject to random relationship-specific fixed costs. This becomes a complicated forward-looking problem due to a Calvo-type time-dependent adjustment friction. In my paper, I provide a complementary approach, where instead of a roundabout continuum network, individual firms search for downstream buyers and upstream suppliers to form production chains of varying lengths and characteristics.

While several papers provide evidence of search frictions in firm-to-firm trade, my paper is one of the first to study their implications in a business cycle model. On the empirical side, Lenoir et al. (2022) and Eaton et al. (2021) are two recent examples of a long literature documenting search costs in exporting and importing. Evidence in Fernandez-Villaverde et al. (2021) and Miyauchi (2021) suggests that relationship formation between domestic firms also requires search effort. However, the theoretical implications of search frictions in domestic trade remain largely unstudied, with the two aforementioned papers providing key exceptions[5]. Fernandez-Villaverde et al. (2021) provide a new microfoundation for the notion of coordination failures in macroeconomics in a model of two-stage production, where final goods and input producers engage in costly search for each other. More closely related to my work is the model of Miyauchi (2021). He embeds a dynamic Diamond-Mortensen-Pissarides search and matching setup into a static model of trade between regions to provide a novel mechanism for agglomeration spillovers in economic geography.

Outline. The paper is structured as follows. Section 2 outlines the model and solution strategy. Section 3 presents the calibration and discusses firms' behavior in steady state, and an example of aggregate dynamics. Section 4 concludes.

2 Model

2.1 Model overview

Time is continuous and extends to infinity. The economy produces a continuum of differentiated final goods in production chains with varying levels of specialization. In particular, the multi-stage production process can feature up to with $H+1$ stages, indexed by $h = 0, \ldots, H$, from the most downstream to the most upstream.

The economy is populated by a continuum of firms, each one of which is endowed with the Cobb-Douglas technology for participating in a fixed stage of the production process. For simplicity, I assume that each stage has a fixed unit mass of firms, with no entry or exit. In each stage, firms must combine labor and a single input into goods to be shipped to the next stage downstream. If firms do not have an input supplier, they may also produce it in-house using only labor. Firms that do not belong to a production chain that sells to the final goods market have no source of demand, and thus do not produce. I call such supply chains *inactive*.

For now, I consider an economy whose production sector consists entirely of *linear supply*

[5] Here, I am excluding from consideration models of customer accumulation such as Gourio and Rudanko (2014) or Roldan-Blanco and Gilbukh (2021), which do not explicitly account for firm-to-firm trade.

chains, where each firm has a maximum of one buyer and one seller. This is similar to the models of sequential production in the global value chain literature, such as Antràs and de Gortari (2020) or Antràs (2023). Since my intention is to model *long-term relationships*, and later enlarge the model to include goods bought and sold directly to the market, I motivate my assumption of one upstream link with an observation by Duprez and Magerman (2018), who show that firms typically have one dominant supplier, whose input share is 25% on average[6].

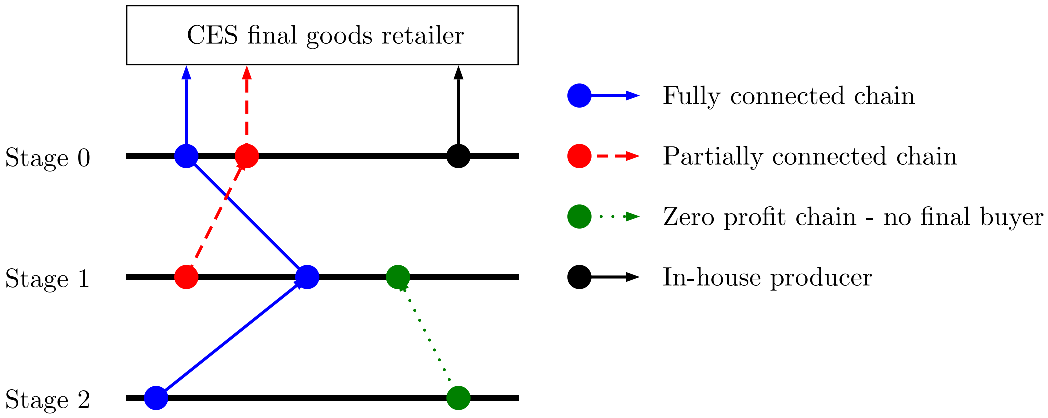

Figure 1: Possible production chains

Figure 1 visualizes the possible structures of production chains in an economy with three stages. For example, one may think of producing bread via a stylized process that involves turning wheat into flour, and flour into bread. The black vertical lines at each stage represent the continuum of firms in that particular stage. The solid line titled "fully connected chain" exemplifies a supply chain with the maximum level of specialization, where each stage of production is outsourced to a supplier, and no in-house production of multiple input goods takes place. The red dashed line titled "partially connected chain" features only two layers of specialization. Finally, the dotted line entitled "zero profit chain – no final buyer" shows just that – an inactive chain that does not produce and thus makes zero profit, since it lacks a firm doing the last stage of the production process.

As will be shown below, the model is specified in a way such that from the point of view of cost-effectiveness and flow profitability, it is always better to specialize as much as possible. However, since firms engage in random search for trade partners, including "on-the-job" search that may result in them leaving their production chain for a better one, a trade-off emerges: despite being more profitable, longer and more specialized production chains may nevertheless be unstable enough to be unattractive for firms.

2.2 Intermediate firms' behavior

Production technology and unit cost. Consider an individual firm, indexed by $i \in [0, 1]$, located in stage $h \in \{0, 1, 2, .., H\}$. As stated above, the firm produces using labor and, if

[6]Extending to model to accommodate multiple buyers per firm, as well as multiple sellers inasmuch it is possible, is a key task for future work.

matched to a supplier, an input purchased from the next stage. Since the firm's technology is Cobb-Douglas, its unit cost in logs is given, up to a constant, by

$$c_{hi,t} = -z_{hi,t} + (1-\omega)w + \omega(p_{i,h+1,t} - q). \tag{1}$$

Here, $z_{hi,t}$ is the firm's idiosyncratic productivity, w is the (constant and homogeneous) wage, ω is the (homogeneous across stages) input share in technology, $p_{i,h+1,t}$ is the price at which the firm obtains inputs from stage $h+1$ at time t, and $q > 1$ is a cost-reducing quality term. It captures benefits from forming long-term relationships, such as product customization, or other forms of "relationship capital" as studied by e.g. Heise (2021)[7]. The firm's idiosyncratic productivity is drawn from a finite $z_{hi,t} \in \{z_1, \dots, z_{N_z}\}$, and it envolves according to a continuous-time Markov chain with jump matrix $Q_z(z' \mid z)$. As stated above, if a firm has no supplier, it can produce using only labor, with log unit cost

$$c_{hi,t} = -z_{hi,t} + w. \tag{2}$$

In the current version, pricing is kept exogenous for simplicity. Following a large swathe of the production networks literature, such as Acemoglu and Azar (2020), the price is set to an exogenous mark-up over unit cost,

$$P_{hi,t} = (1+\mu)C_{hi,t}. \tag{3}$$

For now, the markup μ is fixed and homogeneous across stages. The implication is that for firms matched with a supplier, the unit cost may be written

$$c_h = -z_h + (1-\omega)w + \omega(m + c_{h+1} - q), \tag{4}$$

where $m := \log(1+\mu)$, and I have dropped the firm and time indices.

Link creation and destruction. Firms in stage h engage in random search for suppliers in the upstream stage $h+1$, and for buyers in the downstream stage $h-1$. Moreover, matched firms may search for better buyers and suppliers, similarly to "on-the-job" search in labor market models. Note that a match between two firms implies a match between their two upstream or downstream sub-chains.

I assume that search is undirected and costless, so that each firm searches continuously at rate $\lambda \in (0,1)$. The number of meetings is determined according to a Leontief matching function. This assumption, combined with the assumption that each stage has a unit mass of firms, means the arrival rate of meetings with buyer or suppliers in each stage is simply λ. I assume that matches are exogenously destroyed at a rate $\delta \in (0,1)$, after which both firms go back to searching for a replacement partner.

[7]However, in the current version of the model, the term is largely inconsequential. I plan for it to play a larger role in future version of the model where unmatched firms have the possibility of producing and purchasing uncustomized "generic" goods.

Flow profit. I assume that firms in the most downstream stage face a CES final demand, derived from the preferences of consumers, given by $Y_0 = Dp_0^{-\sigma}$, or in logs, $y_0 = d - \sigma(m_0 + c_0)$, where $m_0 := \log(\frac{\sigma}{\sigma-1})$ is the final goods markup. From an individual firm's point of view, D is a constant. In general equilibrium, it depends on preference parameters, household endowments, and the economy's price index, as will be made clear below. The price elasticity of final demand, σ, is an important determinant of the profitability gains from cost-effectiveness.

The environment specified above implies that a firm's profit in a production chain is determined by the $(H+1)$-vector of the chain members' idiosyncratic productivities, which I denote by $\mathbf{x}$. In terms of notation, I use the empty set symbol $\emptyset$ to mark unmatched stages in the chain. For example, $\mathbf{x} = (z_0, z_1, z_2, \emptyset)$ means that the firm is matched with a chain where the firms in stages $h = 0, 1, 2$ have productivities z_0, z_1, z_2, respectively, and the most upstream stage remains unmatched.[8]

Consider then the flow profit made by a stage-h firm in chain $\mathbf{x}$. By the Cobb-Douglas assumption, each firm's profit equals a fraction of the total cost of their buyer:

$$\Pi_h(\mathbf{x}) = \underbrace{\frac{\mu}{1+\mu}}_{\substack{\text{profit}\\ \text{margin}}} \times \underbrace{\omega}_{\substack{\text{intermediate}\\ \text{share}}} \times \underbrace{C_{h-1}Y_{h-1}}_{\substack{\text{buyer's}\\ \text{expenditure}}} \tag{5}$$

Furthermore, each firm's expenditure is a share of their buyer's expenditure:

$$C_h Y_h = \frac{\omega}{1+\mu} C_{h-1} Y_{h-1} = \ldots = \left(\frac{\omega}{1+\mu}\right)^h C_0 Y_0, \tag{6}$$

Again, set profit is zero if the chain is not linked all the way the the most downstream stage 0. Combining these observations, we get that a firm's flow profit in logs is given by

$$\pi_h(\mathbf{x}) = \log \hat{\mu}_h + d - \sigma m_0 + h(\log \omega - m) - (\sigma - 1)c_0(\mathbf{x}), \tag{7}$$

where $m := \log(1+\mu)$, $\hat{\mu}_0 = \frac{\sigma}{\sigma-1}$ and $\hat{\mu}_h = \mu$ otherwise, and $c_0(\mathbf{x})$ is the (log) unit cost of the most downstream firm, which is a measure of the cost-effectiveness of the entire chain. This expression of flow profit implies that when searching for suppliers, the firm wants to minimize the unit cost *of the entire chain* c_0, but weighs the stability of the chain (the distribution of future costs) as well[9].

[8] Of course, this also implies that the maximum number of stages in a production process is four.

[9] One feature of the current setup, to be fixed in future work, is that with homogeneous markups and input shares across stages, more downstream firms generally make less profit due to the term $h(\log \omega - m) \approx h(\log \omega - \mu) < 0$.

The unit cost of the chain $c_0(\mathbf{x})$ can furthermore be expressed as

$$\begin{aligned} c_0(\mathbf{x}) &= -\sum_{l=0}^{h-1} \omega^l \left[z_l - (1-\omega)w\right] - \sum_{l=0}^{h-1} \omega^{l+1}(q-m) + \omega^h c_h \\ &= \underbrace{-\sum_{l=0}^{\bar{h}} \omega^l z_l}_{\text{chain productivity}} \underbrace{-\sum_{l=0}^{\bar{h}-1} \omega^{l+1}(q-m)}_{\text{gain from specialization}} + w, \end{aligned}$$

where the cost in the last linked stage $\bar{h}$ is $c_{\bar{h}} = -z_{\bar{h}} + w$, as it is unmatched.

Matching decisions and the Hamilton-Jacobi-Bellman equation Given that optimization of input demand is a simple static decision, the firm's main problem is whether to match or not match with the buyers and suppliers that it meets[10]. These policies can be summarized by an *acceptance matrix*,

$$\mathcal{A}_h(\mathbf{x}, \mathbf{x}'),$$

which specifies whether the stage-h firm accepts matches that lead from state $\mathbf{x}$ to state $\mathbf{x}'$.

Of course, for a link to be formed after a meeting, both sides have to accept the match. Therefore, I define a *matching indicator* between stage h and $h+1$ firms as:

$$\mathcal{A}^*_{h:h+1}(\mathbf{x}^h, \mathbf{x}^{h+1}) := \mathcal{A}_h(\mathbf{x}^h, \mathbf{x}')\mathcal{A}_{h+1}(\mathbf{x}^{h+1}, \mathbf{x}').$$

where $\mathbf{x}^h$ and $\mathbf{x}^{h+1}$ are the states of the firms, and $\mathbf{x}' = (\mathbf{x}^h_{0:h}, \mathbf{x}^{h+1}_{h+1:H})$ is the new chain that results from the match.

The firms' problem is tremendously simplified by the use of continuous time and the particular specification of uncertainty. In accordance with labor market search models in the Diamond-Mortensen-Pissarides tradition, I've assumed that meetings with other firms and exogenous link destruction follow Poisson processes. In addition, productivity shocks are drawn from a jump process. This implies that all the stochastic features of the firms' environment follow jump processes. In other words, from the perspective of each firm, only one event can happen in each chain at each instant[11]. This removes complicated decision problems that could arise when multiple events take place the same time. For example, firms never have decide on a match while being uncertain about whether there has been a productivity shock or separation in another part of the chain.

The firms' Hamilton-Jacobi-Bellman equation can be expressed as

$$\rho V_h(\mathbf{x}) = \max_{\mathcal{A}_h(\cdot)} \left\{ \Pi_h(\mathbf{x}) + \sum_{\mathbf{x}' \in \mathcal{S}_h} \mathbb{P}_h(\mathbf{x}, \mathbf{x}') \Big[V_h(\mathbf{x}') - V_h(\mathbf{x}) \Big] \right\},$$

[10]Each firm always has the option to unilaterally separate from any buyer/seller relationship. However, in the current setup, they would never want to do so, as each match has a positive surplus. Also, decision rules are not functions of equilibrium distributions in steady state – but they would be in a fully dynamic model

[11]Technically, this is only true almost surely / with probability one.

where $\mathcal{S}_h$ is the state space (collection of all chain vectors $\mathbf{x}$), and the *transition matrix* $\mathbb{P}_h(\mathbf{x}, \mathbf{x}')$ is a conceptually uncomplicated but tedious-to-compute function of the distributions $\mathbf{m}_h(\mathbf{x})$ of firms across chains in each stage, and of each firms' acceptance matrices $\mathcal{A}_h(\mathbf{x}, \mathbf{x}')$.[12]

2.3 General equilibrium

The preceding subsection discussed firms' environment, technology and optimal matching problem in partial equilibrium. The following section provides the remaining building blocks to close the model.

Households' consumption demand & labor supply. The representative household is risk-neutral, and their problem is completely static. Every instant, they receive flows of labor income $W_t L_t$ and dividends from firms $\bar{\Pi}_t$, and consume an aggregate of final goods C_t^d (where the superindex d distinguishes this from an index of aggregate unit costs, to be introduced later). The resulting utility flow is given by

$$U(C_t^d, L_t),$$

where L_t is the supply of hours worked. The budget constraint is given by

$$\int_0^1 P_{i,t} C_{i,t}^d \, \mathrm{d}i = W_t L_t + \bar{\Pi}_t := E_t,$$

where $C_{i,t}$ and $P_{i,t}$ are the consumption and price of final goods variety i, and E_t denotes the representative household's nominal income.

First, consider the optimal allocation of consumption. The consumption aggregate is of the standard CES form, $C_t^d := \left(\int_0^1 (C_{i,t}^d)^{\frac{\sigma-1}{\sigma}} \, \mathrm{d}i\right)^{\frac{\sigma}{\sigma-1}}$. This implies a standard demand function and price index given by $P_t = \left(\int_0^1 P_{i,t}^{1-\sigma} \, \mathrm{d}i\right)^{\frac{1}{1-\sigma}}$.

The final goods producers' profit-maximizing pricing rule is $P_{i,t} = \frac{\sigma}{\sigma-1} C_{i,t}$, where $C_{i,t}$ is the unit cost of the (stage 0) firm producing final goods variety i. Plugging this into the price index implies that $P_t = \frac{\sigma}{\sigma-1}\left(\int_0^1 C_{i,t}^{1-\sigma} \, \mathrm{d}i\right)^{\frac{1}{1-\sigma}}$. Inserting these into the standard CES demand function, we have demand for variety i expressed in terms of *relative cost*:

$$C_{i,t}^d = \left(\frac{C_{i,t}}{\bar{C}_t}\right)^{-\sigma} C_t^d, \tag{8}$$

where $\bar{C}_t := \left(\int_0^1 C_{i,t}^{1-\sigma} \, \mathrm{d}i\right)^{\frac{1}{1-\sigma}}$ is an *aggregate unit cost index.* In steady-state equilibrium, we can express it as

$$\bar{C} = \left(\sum_{\mathbf{x}} m_0(\mathbf{x}) C_0(\mathbf{x})^{1-\sigma}\right)^{\frac{1}{1-\sigma}}, \tag{9}$$

where $m_0(\mathbf{x})$ is the mass of final goods firms with state $\mathbf{x}$, and $C_0(\mathbf{x})$ is the unit cost of such a firm.

[12] A description of the computation of the transition matrix is left for a future version of the draft.

Finally, imposing goods market clearing implies that $C_{i,t}^d = Y_{i,t}$, where the RHS is the output of good i, for all t and i. Therefore, in the demand function, we can replace $C_t^d = Y_t := (\int_0^1 Y_{i,t}^{\frac{\sigma-1}{\sigma}}\,\mathrm{d}i)^{\frac{\sigma}{\sigma-1}}$. In steady-state equilibrium, we can drop time indices and rewrite the household budget constraint as

$$PY = (1+\mu_0)\bar{C}Y = E, \tag{10}$$

where $1+\mu_0 := \frac{\sigma}{\sigma-1}$.

Now, consider labor supply. I assume that the household's flow utility function is of the constant-elasticity GHH form,

$$U(C,L) = \frac{1}{1-\phi}\left(C - \Psi\frac{L^{1+\theta}}{1+\theta}\right)^{1-\phi}.$$

After standard solution steps, we obtain the usual expression for aggregate labor supply, which is a function of the (steady-state) real wage only:

$$L = \Psi^{-\frac{1}{\theta}}\left(\frac{W}{P}\right)^{\frac{1}{\theta}} = \Psi^{-\frac{1}{\theta}}\left(\frac{W}{\frac{\sigma}{\sigma-1}\bar{C}}\right)^{\frac{1}{\theta}}. \tag{11}$$

Firms' profit revisited. The demand for variety i in (8), along with market clearing and the household budget constraint (10), implies that a final goods firm in a chain $\mathbf{x}$ faces a demand function given by

$$Y_0(\mathbf{x}) = \left(\frac{C_0(\mathbf{x})}{\bar{C}}\right)^{-\sigma}\frac{E}{(1+\mu_0)\bar{C}}. \tag{12}$$

As a consequence, profit in stage h is is given by

$$\Pi_h(\mathbf{x}) = \mu_h \cdot \frac{E}{1+\mu_0}\left(\frac{\omega}{1+\mu}\right)^h\left(\frac{C_0(\mathbf{x})}{\bar{C}_0}\right)^{1-\sigma}, \tag{13}$$

where $\mu_0 = \frac{1}{\sigma-1}$ and $\mu_h = \mu$ for $h > 0$. This implies that the nominal demand factor in (7) is given by

$$d = (\sigma-1)(m_0+\bar{c}) + e. \tag{14}$$

where $e = \log(E)$ is log household nominal income and $\bar{c} = \log(\bar{C})$ is the log of the aggregate unit cost index in (9).

Labor demand. Consider the labor demand of the stage-h firm in an active chain $\mathbf{x}$, i.e. a chain that is connected until the final goods stage 0 and is thus producing. By cost minimization and Shepard's lemma, the firm's labor demand is given by

$$L_h^d = \frac{\partial[C_h(\mathbf{x})Y_h(\mathbf{x})]}{\partial W} = (1-\omega)\frac{C_h(\mathbf{x})Y_h(\mathbf{x})}{W} = (1-\omega)\left(\frac{\omega}{1+\mu}\right)^h\frac{C_0(\mathbf{x})Y_0(\mathbf{x})}{W},$$

where $C_h(\mathbf{x})$ and $Y_h(\mathbf{x})$ are the unit cost and output at stage h, respectively, and the last line uses (6) to connect total cost in stage h to that in stage 0. In the last connected stage $\bar{h}$, labor is the only input, so

$$L^d_{\bar{h}} = \frac{Y_{\bar{h}}}{Z_{\bar{h}}} = \left(\frac{\omega}{1+\mu}\right)^{\bar{h}} \frac{C_0(\mathbf{x})Y_0(\mathbf{x})}{W}.$$

Summing up over stages, the labor demand of a chain $\mathbf{x}$ that is connected until stage $\bar{h}(\mathbf{x})$ is

$$\begin{aligned}
L^d(\mathbf{x}) = \sum_{h=0}^{\bar{h}(\mathbf{x})} L^d_h(\mathbf{x}) &= \frac{C_0(\mathbf{x})Y_0(\mathbf{x})}{W}\left((1-\omega)\sum_{h=0}^{\bar{h}(\mathbf{x})-1}\left(\frac{\omega}{1+\mu}\right)^h + \left(\frac{\omega}{1+\mu}\right)^{\bar{h}(\mathbf{x})}\right) \\
&= \frac{1}{1+\mu_0}\left(\frac{C_0(\mathbf{x})}{\bar{C}}\right)^{1-\sigma}\frac{E}{W}\cdot\left(1-\frac{\mu\omega}{1+\mu-\omega}\left[1-\left(\frac{\omega}{1+\mu}\right)^{\bar{h}(\mathbf{x})}\right]\right),
\end{aligned}$$

where on the last line we've substituted for $C_0(\mathbf{x})Y_0(\mathbf{x})$ using (12). Notice that if $\mu = 0$, then the last factor, which is less than unity, disappears. Unsurprisingly, markups increase costs and drag down labor demand. The effect is larger in longer chains, with a larger $\bar{h}(\mathbf{x})$.
Summing up over the distribution of active chains, aggregate labor demand becomes

$$\bar{L}^d = \sum_{\mathbf{x}} m(\mathbf{x})L^d(\mathbf{x}) = \frac{1}{1+\mu_0}\frac{E}{W}\left(1-\frac{\mu\omega}{1+\mu-\omega}\left[1-\sum_{\mathbf{x}} m(\mathbf{x})\left(\frac{C_0(\mathbf{x})}{\bar{C}}\right)^{1-\sigma}\left(\frac{\omega}{1+\mu}\right)^{\bar{h}(\mathbf{x})}\right]\right). \tag{15}$$

In practice, I choose the wage as the numeraire, setting $W = 1$.

2.4 Model solution

In partial equilibrium, the firms' problems are fairly easy to solve using value function iteration. A key step, conceptually simple yet requiring tedious accounting, is the construction of the transition matrix $\mathbb{P}_h(\cdot)$. It involves enumerating, for each pair of chains $(\mathbf{x}, \mathbf{x}')$, all the possible events that may lead to a transition from the former chain to the latter in a vanishingly small interval of time.

First, some notation. Let the matrix $\mathbf{P}_h$ denote the collection of the transition probabilities $\mathbb{P}_h(\mathbf{x}, \mathbf{x}')$ for all pairs of states $\mathbf{x}, \mathbf{x}'$. Define the profit and value vectors $\boldsymbol{\pi}_h$ and $\mathbf{v}_h$ similarly. Finally, letting $m_h(\mathbf{x})$ denote the (steady-state) mass of stage-h firms in state $\mathbf{x}$, stack these values into a vector $\mathbf{m}_h$.

Given the transition matrix, the invariant distribution of firms over states satisfies

$$\mathbf{m}_h = \mathbf{P}'_h\mathbf{m}_h \quad \Longleftrightarrow \quad (\mathbf{I} - \mathbf{P}'_h)\mathbf{m}_h = \mathbf{0}.$$

With the notation defined above, the maximized HJB equation can be written

$$\rho\mathbf{v}_h = \boldsymbol{\pi}_h + (\mathbf{P}_h - \mathbf{I})\mathbf{v}_h.$$

This suggests an easy algorithm. Initialize $\mathbf{m}_h$ with uniform distributions, and firms' policies with "accept all"[13]. Then, iterate until convergence on the following steps:

1. Compute the transition matrix $\mathbf{P}_h$, given distributions and policies
2. Looping over (h, z) pairs:
   (a) Update the value function using
   $$\mathbf{v}_h^{\text{new}} = \left[(1+\rho+1/\Gamma)\mathbf{I} - \mathbf{P}_h\right]^{-1}(\boldsymbol{\pi}_h + \mathbf{v}_h^{\text{old}}/\Gamma)$$
   (b) Update matching/acceptance policies:[14]
   $$\mathcal{A}_h^{\text{new}}(\mathbf{x}, \mathbf{x}') = \mathbb{1}\{\mathbf{v}_h^{\text{new}}(\mathbf{x}') > \mathbf{v}_h^{\text{new}}(\mathbf{x})\}$$
   (c) Solve for the distribution from $(\mathbf{I} - \mathbf{P}_h')\mathbf{m}_h^{\text{guess}} = 0$, normalize it, and update:
   $$\mathbf{m}_h^{\text{new}} = (1 - 1/\Gamma)\mathbf{m}_h^{\text{guess}} + \mathbf{m}_h^{\text{old}}/\Gamma$$

   where Γ is some large number.

In practice, both values and distributions converge quite quickly.

Solving for the steady-state *general* equilibrium also involves solving for the nominal demand term d in firms' profit functions (7). The procedure for the outer loop to do so is as follows:

1. Make an initial guess for d.
2. Solve the firms' problems given the algorithm above, and compute the distribution of final goods firms $m_0(\mathbf{x})$.
3. Using the distribution, compute a new guess for d. In particular, find the unit cost index $\bar{C}$. Then solve for labor supply using (11). Given labor supply, we can solve for nominal income E from labor demand (15). Finally, compute d using (14).

3 Preliminary results

3.1 Calibration

For the purposes of this paper, I choose an illustrative calibration that allows me to highlight the key mechanisms that shape the steady-state distribution of production chains. A central avenue

[13]While the model might in principle feature multiple equilibria, in all my numerical experiments the choice of initial policies or distributions does not seem to matter for the convergence of the algorithm, nor for the equilibrium that is reached.

[14]This policy attains the maximum in the HJB equation. However, in some cases the optimal policy is not unique. For example, the firm may be indifferent because the other side in each match that would lead from state $\mathbf{x}$ to $\mathbf{x}'$ rejects the match. In such situations, any policy is trivially optimal, since it has no effect on the matching outcome. The policy here forces the firm to behave *as if they always decided the matching outcome* while only acknowledging actual matches when computing the actual transition matrix. This implies some equilibrium selection. For example, it rules out trivial Nash equilibria where no matching ever happens, since every match is rejected by both sides. These equilibrium refinement considerations will be worked out more formally in future work.

for future work is making the calibration more empirically realistic. This will entail comparing implications of the model to the data, as well as selecting the best available estimates from existing literature whenever possible.

To keep the scale of the model as small as reasonably possible, I set both the maximum number of production stages $H+1$ and the number of productivity ladders N_z to 3 [15]. Under this calibration, the model can be solved in seconds, and the agents' state spaces are small enough that equilibrium objects such as value functions and transition matrices can be easily examined to obtain insights on firm behavior.

The stochastic environment facing the firms is controlled by four key parameters: the match arrival rate λ, the exogenous separation rate δ, and the probabilities of upward and downward jumps on the productivity ladder (p_u and p_d, respectively). I set these in a way that emphasizes matching as the key source of uncertainty. In particular, since I set the arrival rate of exogenous separations to $\delta = 0.01$, the match arrival rate $\lambda = 0.15$ is 15 times larger. Similarly, since productivity shocks in the model are relatively large jumps along the very sparse ladder, I calibrate their frequency to be relatively small. Since

Table 1: Baseline calibration

Parameter	Interpretation	Value
$H+1$	Number of stages	3
N_z	Number of productivity ladders	3
ρ	Discount rate	0.05
δ	Exog. separation rate	0.01
λ	Match arrival rate	0.15
σ	Final demand elasticity	3
μ	Markup	0.025
q	Relationship quality	0.025
ω	Intermediate share	0.5
Ψ	Labor disutility shifter	2
θ	Labor disutility elasticity	4
z_1	Low productivity	0
z_3	High productivity	0.5
-	Productivity step	0.25
p_u	Prob. of upward jump in z	0.01
p_d	Prob. of downward jump in z	0.01

Most of the other parameters appear to be much less consequential for the equilibrium distribution, so I set them to illustrative baseline values. An exception is the final demand elasticity σ, which controls strongly how higher productivity in the chain translates to greater demand from the final goods retailers. It therefore affects firms' incentives to climb the cost-efficiency ladder by matching with more efficient chains. It also pins down the markup earned

[15] Furthermore, real-world evidence points to most production chains being relatively short: Dhyne and Duprez (2015) show that in the Belgian economy, firms typically participate in production chains with a maximum of 4 or 5 (domestic) stages.

by final goods producers. Calibrating this elasticity accurately is an goal for the future[16].

3.2 Steady-state implications of search behavior

Figure 2 plots the steady state distribution of firm types – defined by stage, productivity pairs – across the state space. A note on the presentation is in order: Vectors on the x-axis denote the states of the firm, ordered by upstreamness, but *excluding* the firm's own stage and state. Thus, for example, the state "$(1, \emptyset)$" of a stage-1 firm with productivity level 2 denotes that this firm is a member of the chain $(1, 2, \emptyset)$. The states are ordered by their productivity.

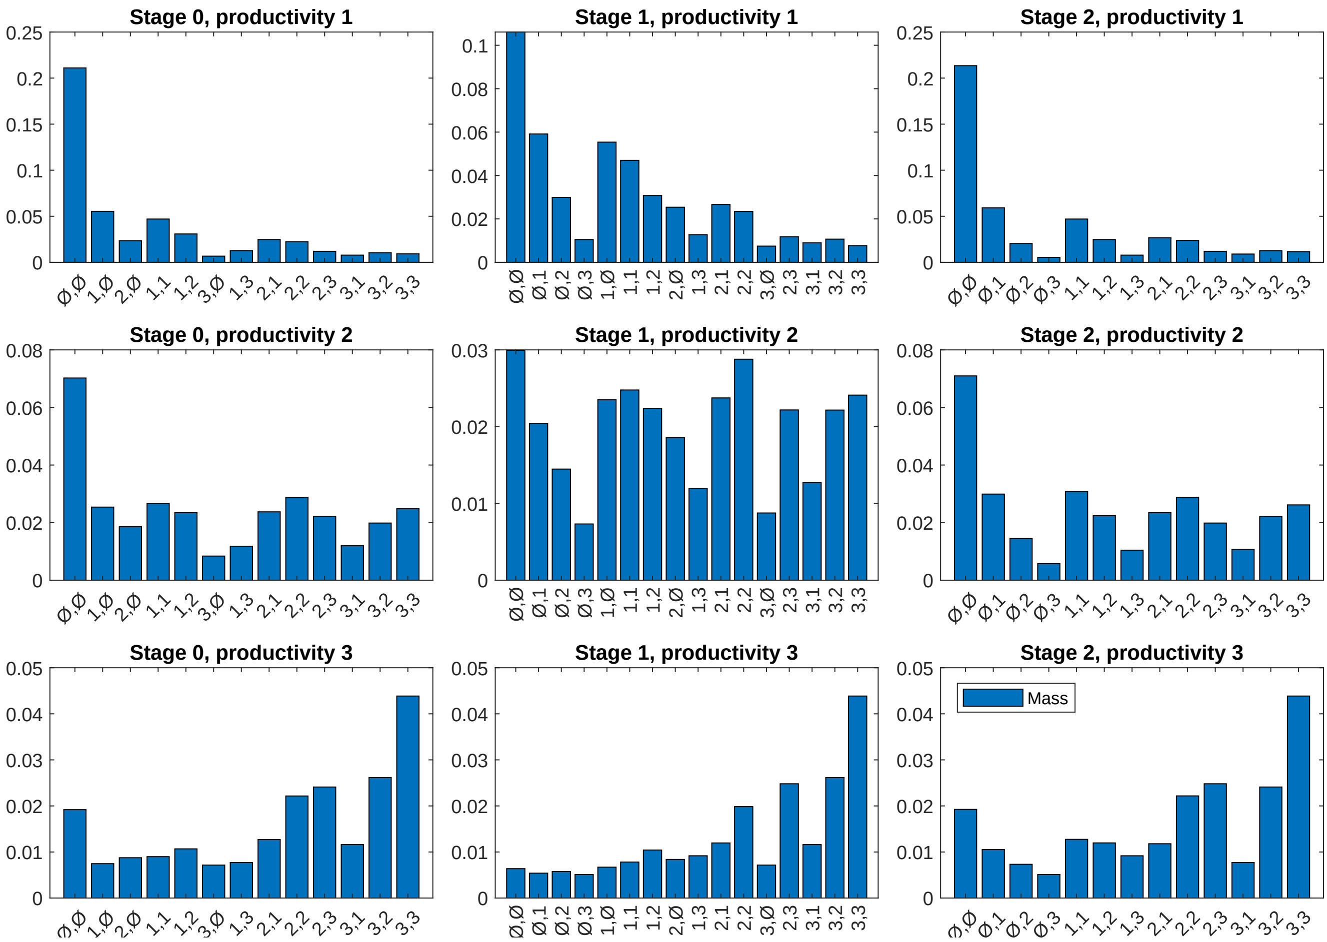

Figure 2: Distribution of firms across the state space, for different firm types

An immediate striking result is that the highest-productivity firms in every stage tend to sort into matches with each other, while the lowest-productivity firms often remain completely *unmatched*. Even the lowest-productivity firms in the middle stage receive some boost to their chances of matching from their key role in the elaboration of upstream inputs for usage by downstream firms. Figure 3, which shows the distribution of profitability across chains, drives

[16] In the context of the current partial-equilibrium version of the model, the elasticity could be based on industry-level estimates. The currently chosen value of $\sigma = 3$ is on the lower end of the range of empirical estimates.

home the point that this large fraction of unmatched firms entails significant losses in profit and output[17]. To consider one benchmark, the *static* efficient allocation would feature perfect sorting[18]. That is, all high/mid/low productivity firms would only form chains with each other, and there would be no unmatched firms.

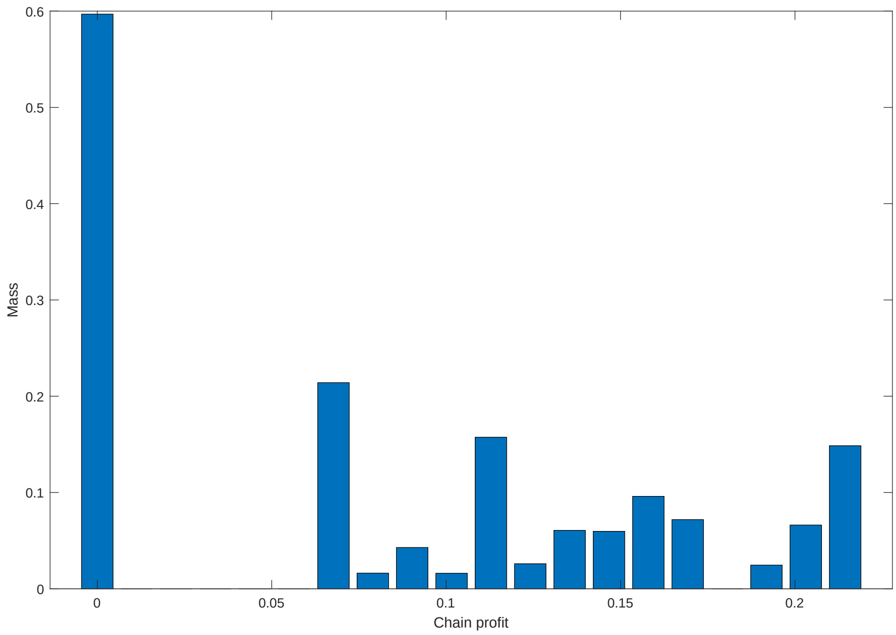

Figure 3: Profit distribution of chains

This sorting behavior is also reflected in the transition matrices of low-productivity firms, which are depicted on the first row of Figure 4. The only most probably way out of the unmatched state "$(\emptyset, \emptyset)$" for them, regardless of the stage, is to match with very short, low profit chains. Examining the value functions in Figure 5 furthermore shows that basically all firm types prefer to match with almost any other chain except a low-productivity singleton. Indeed, unmatched low-productivity firms will tend to only match with other unmatched low-productivity firms. But as soon as one member of such a chain receives the opportunity to join a more profitable chain, or a productivity shock, it is likely that the chain falls apart again. In contrast, higher-productivity chains find it easy to obtain new members.

[17]Of course, this result also depends heavily on the assumption that firms that have no customers do not exit, but wait for matches to appear. It is also by-product of the baseline specification where each production stage has a unit mass of firms, as well as of the productivity distribution.

[18]Characterization of the *dynamic* efficient allocation with exogenous separations is work in progess.

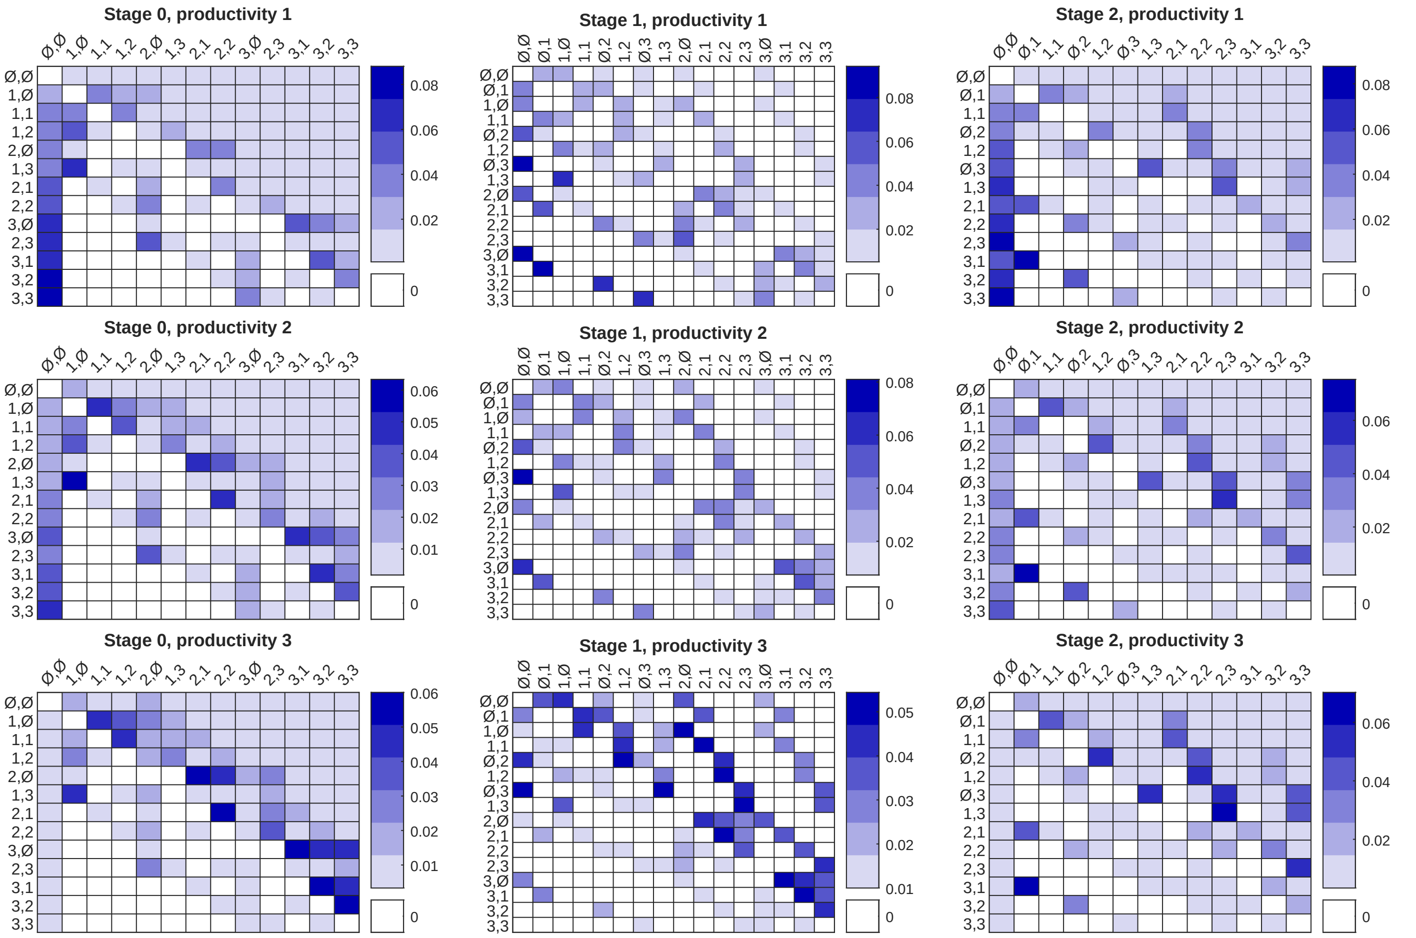

Figure 4: Transition matrix heatmaps across firm types

Figure 5 also reveals another interesting pattern. The states in the figure, to be interpreted identically as in Figure 2 above, are ordered by their productivity. However, firms' value functions over the states are clearly non-monotonic. That is, firms' forward-looking considerations have an effect on their matching behavior, so that they do not always prefer to join the production chain guaranteeing the highest flow profit. These reversals are particularly pronounced for the most upstream firms, but occur throughout the chain. These non-monotonicities between value and productivity are also clearly visible in the policy function or "acceptance matrix" plot, shown in the Appendix (Figure A.1).

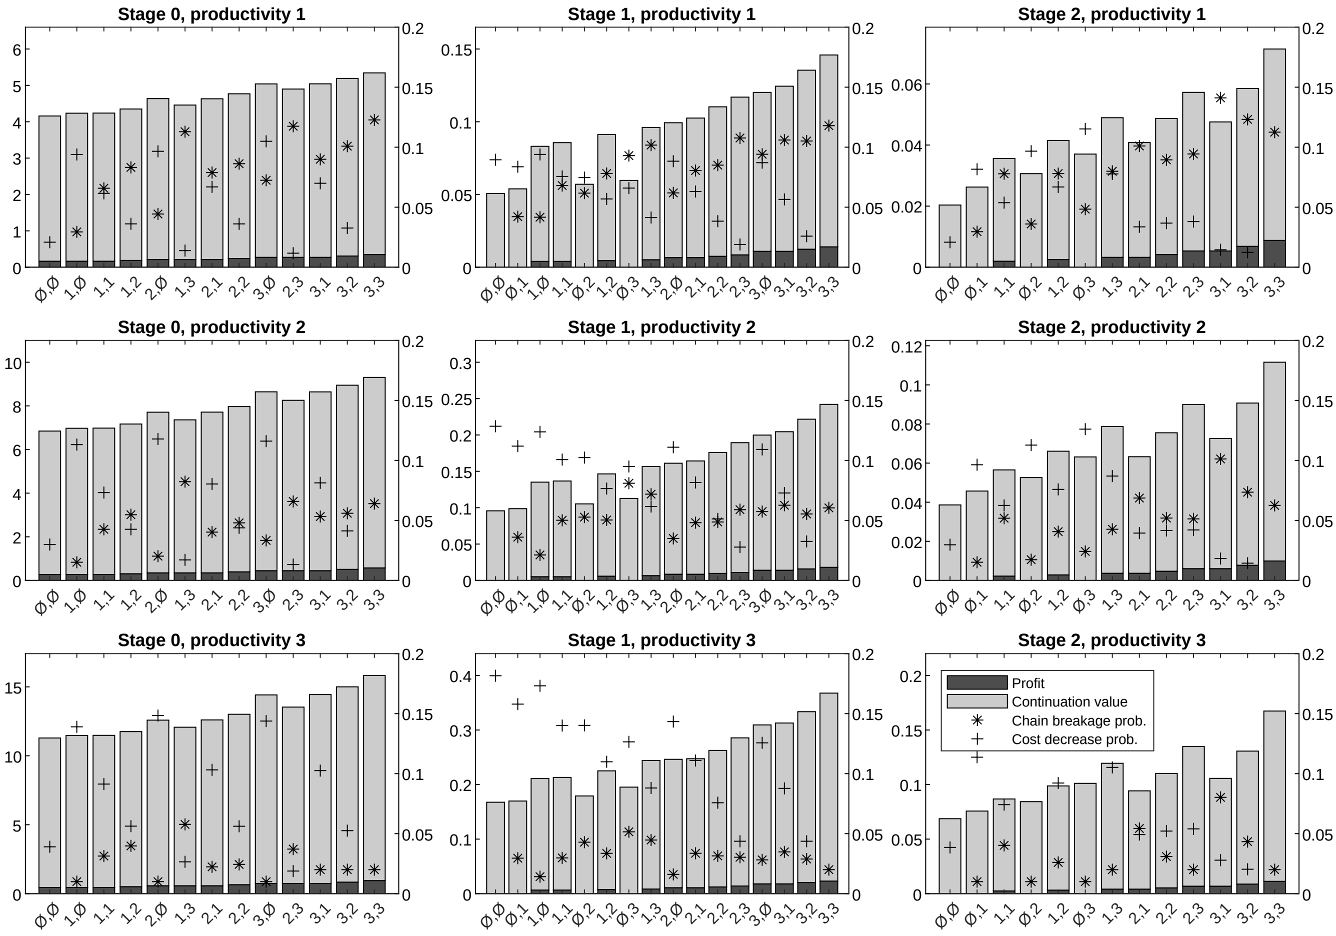

Figure 5: Decomposition of the value function into profit and continuation value across firm types and states (left axis), along with measures of chain instability (right axis)

To shed light on the reasons behind these non-monotonicities, Figure 5 also includes crude measures of the instability of the chain. On the one hand, the asterisks indicate the *breakage probability* of the chain (on the right axis), which is measured as the probability that one or more of the other members of the chain separate within the next period. Many of the non-monotonicities are related to differences in this probability. For example, consider an average-productivity (level 2) firm in the most downstream stage 0. They value being matched to a single high-productivity intermediate supplier (state $(3, \emptyset)$) over joining a longer, more profitable chain (such as states $(2, 3)$ or $(3, 1)$), because the shorter chain offers more stability.

Further up the chain, an upstream input producer in stage 2, regardless of their productivity level, will prefer to join a chain with a less profitable final goods firm of productivity level 2 and a moderately profitable stage-1 input producer (such as in states $(2, 2)$ or $(2, 3)$) over joining a more profitable chain with a productive final goods firm but an unproductive input supplier (state $(3, 1)$), because the latter chain has a vastly higher probability of losing the final goods firm to another chain. The implication is that when either one of these example firm types is a member of the more stable yet lower-profit chain, they will turn down offers to join these more profitable but unstable chains. That is, firms consider the *fragility* of their supply chains when searching for better matches.

Figure 5 also plots another measure of instability – this time in a more positive sense. The *cost decrease probability* (marked with plus symbol, with values also on the right axis), is the probability of transitioning to a state where one or more of the (possibly unoccupied) slots on the production chain is taken by a more productive firm than in the status quo. In particular, the most upstream firms may sometimes prefer to match with a productive mid-stage firm that lacks a final goods producer. Even though the chain makes zero profit, it is highly likely that it will obtain a very profitable match in the future. Therefore, firms sometimes pass on the opportunity to make an immediate profit if the lower-profit chain exhibits higher *growth potential*.

The strength of these mechanisms naturally depends on the model parametrization, whose effects I include results on in the appendix. Particularly important are the size of the productivity ladders, which determines the reward from being matched to highly productive trading partners. Naturally, when the reward from being matched with the best possible firms is higher, firms care less for the risk (Experiment 1 in the appendix). Another important parameter is the arrival rates of matches, exogenous separations and jump productivity shocks. In particular, a higher arrival rate of matches (Experiment 2) increases the possibility of finding a more profitable chain from the perspective of the most productive firms. The other side of the coin is that chain arrangements between less productive firms become unstable and vulnerable to poaching of the best-performing firms.

3.3 Aggregate dynamics: Effect of a link-destruction shock

In this section, I present preliminary results regarding aggregate dynamics in the model. As an example, I consider the transitional dynamics in response to an MIT shock that randomly destroys 10% of all firm-level links in the economy. Figure 6 shows how such a shock affects the distribution of chains with a certain number of links. As is obvious, a significant fraction of supply chains with 2 or 3 stages lose one or more of their connections, which results in shorter chains or entirely unmatched firms.

Figure 7 show the effect on aggregate variables. The effect here is an immediate 2.9% percent reduction in output, from which the economy recovers slowly as firms rebuild their linkages. Meanwhile, the unit cost index increases, since shorter chains are on average less efficient. As a result, output demand weakens, so firms also reduce their labor demand. In the current calibration, there is no definite notion of a time period, so I avoid drawing strong conclusions about the shock's persistence. Also, the impact effect of the shock is dependent on the calibration – a wider productivity ladder implies that supply chain connections are more valuable, meaning the impact effect is greater. Interestingly, firms do not find it necessary to adjust their matching policies very much in response to the shock, since it does not change the relative ranking of expected future profits in different supply chains.

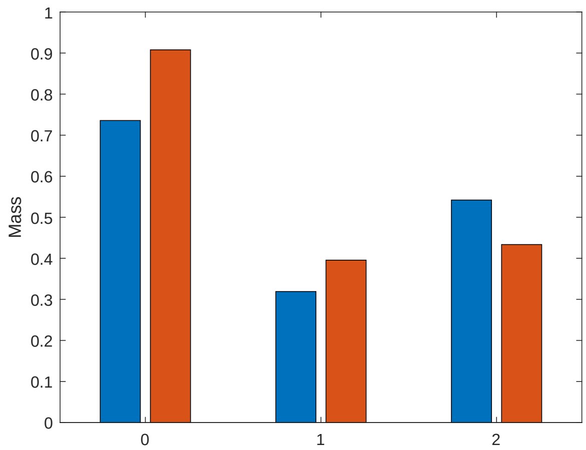

Figure 6: Immediate effect of randomly destroying 10% of all links on the distribution of chains with a certain number of links

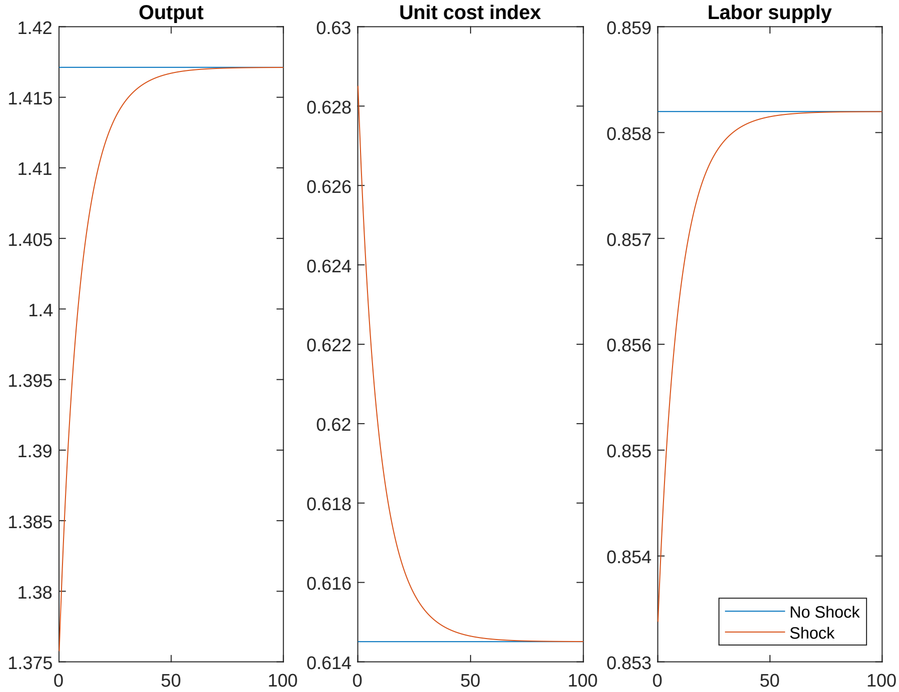

Figure 7: Effects of the link destruction shock on aggregate variables

4 Future work

In this paper, I have presented a stripped down, tractable model of the dynamics of supply chains, where firms face search frictions when looking for customers and suppliers. The preliminary investigation of steady-state behavior contained in this paper represents only the tip of the iceberg when it comes to the potential of the model. Indeed, even delving deeper into the mechanisms at work in the steady-state matching process is ongoing work. For example, the notion of mismatch in a chain, which is closely related to its stability or lack of thereof, is related to the variance of productivity levels in the chain, but not entirely explained by it. It will be useful to find simple, intuitive metrics for mismatch, as well as for illustrating the fragility and growth potential of chains.

The quantitative significance of this matching behaviour is also an open question. One dimension of answering it will involve finding a calibration that is as empirically relevant as possible, as well as tweaking some features of the model to better match the data. Another task will be exploring how firms' matching decisions interact with shocks. As example presented here, I used the model used to conduct a simulation where a certain percentage of the randomly chosen linkages in the steady state distribution is broken, and firms slowly reform their supply chains. Beyond that, it is also possible to implement a general equilibrium approach to study other aggregate shocks following the approach of Boppart et al. (2018).

However, it is possible to go further than that. While the current solution procedure in the absence of aggregate shocks is tractable when each firm has full information regarding their own chain and the distribution of potential matches, it becomes costly when allowing for longer chains. Furthermore, it would be time-consuming when resolving out-of-steady-state dynamics, where the distribution of the economy's production chains becomes a state variable. Hence, the solution procedure to be applied in the completed paper, which I am currently developing, will take a different approach. Following Krusell et al. (1998), I assume that firms use a limited set of "moments" in forecasting, and the forecasting functions are estimated using simulated data. In contrast to the plain vanilla Krusell-Smith algorithm, however, the object to be forecast in my setup consists not only of aggregate variables, but also includes the state of the firm's *own* chain. This algorithm scales well with longer chains and larger state spaces, and can be used to allow for aggregate and stage-specific productivity shocks. It also opens the door to generalizations of the model, such as allowing for multiple sellers and buyers per firm. It is also consistent with anecdotal evidence that very few firms have complete knowledge of their entire supply chains (Elliott et al., 2022).

References

Acemoglu, D. and Azar, P. D. (2020). Endogenous production networks. *Econometrica*, 88(1):33–82.

Acemoglu, D. and Tahbaz-Salehi, A. (2020). Firms, Failures, and Fluctuations: The Macroeconomics of Supply Chain Disruptions. mimeo.

Antràs, P. (2023). An 'austrian' model of global value chains. Working Paper 30901, National Bureau of Economic Research.

Antràs, P. and de Gortari, A. (2020). On the Geography of Global Value Chains. *Econometrica*, 88(4):1553–1598.

Atalay, E., Hortaçsu, A., Roberts, J., and Syverson, C. (2011). Network structure of production. *Proceedings of the National Academy of Sciences*, 108(13):5199–5202.

Boppart, T., Krusell, P., and Mitman, K. (2018). Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control*, 89(C):68–92.

Carvalho, V. M. and Voigtländer, N. (2014). Input diffusion and the evolution of production networks. Working Paper 20025, National Bureau of Economic Research.

Dhyne, E. and Duprez, C. (2015). Has the crisis altered the Belgian economy's DNA ? *Economic Review*, ii:31–43.

Duprez, C. and Magerman, G. (2018). Price updating in production networks. Working Paper Research 352, National Bank of Belgium.

Eaton, J., Eslava, M., Jinkins, D., Krizan, C. J., and Tybout, J. R. (2021). A search and learning model of export dynamics. Working Paper 29100, National Bureau of Economic Research.

Elliott, M., Golub, B., and Leduc, M. V. (2022). Supply network formation and fragility. *American Economic Review*, 112(8):2701–47.

Ersahin, N., Giannetti, M., and Huang, R. (2022). Trade Credit and the Stability of Supply Chains. mimeo.

Fernandez-Villaverde, J., Mandelman, F., Yu, Y., and Zanetti, F. (2021). Search Complementarities, Aggregate Fluctuations and Fiscal Policy. mimeo.

Ghassibe, M. (2021). Endogenous Production Networks and Non-Linear Monetary Transmission. mimeo.

Gourio, F. and Rudanko, L. (2014). Can Intangible Capital Explain Cyclical Movements in the Labor Wedge? *American Economic Review*, 104(5):183–188.

Heise, S. (2021). Firm-to-Firm Relationships and the Pass-Through of Shocks: Theory and Evidence. mimeo.

Huneeus, F. (2020). Production Network Dynamics and the Propagation of Shocks. mimeo.

Krusell, P., Smith, A. A., and Jr. (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy*, 106(5):867–896.

Lenoir, C., Martin, J., and Mejean, I. (2022). Search Frictions in International Goods Markets. *Journal of the European Economic Association*, 21(1):326–366.

Lim, K. (2018). Endogenous Production Networks and the Business Cycle. mimeo.

Miyauchi, Y. (2021). Matching and Agglomeration: Theory and Evidence from Japanese Firm-to-Firm Trade. mimeo.

Roldan-Blanco, P. and Gilbukh, S. (2021). Firm dynamics and pricing under customer capital accumulation. *Journal of Monetary Economics*, 118(C):99–119.

Taschereau-Dumouchel, M. (2020). Cascades and Fluctuations in an Economy with an Endogenous Production Network. mimeo.

A Additional figures and tables

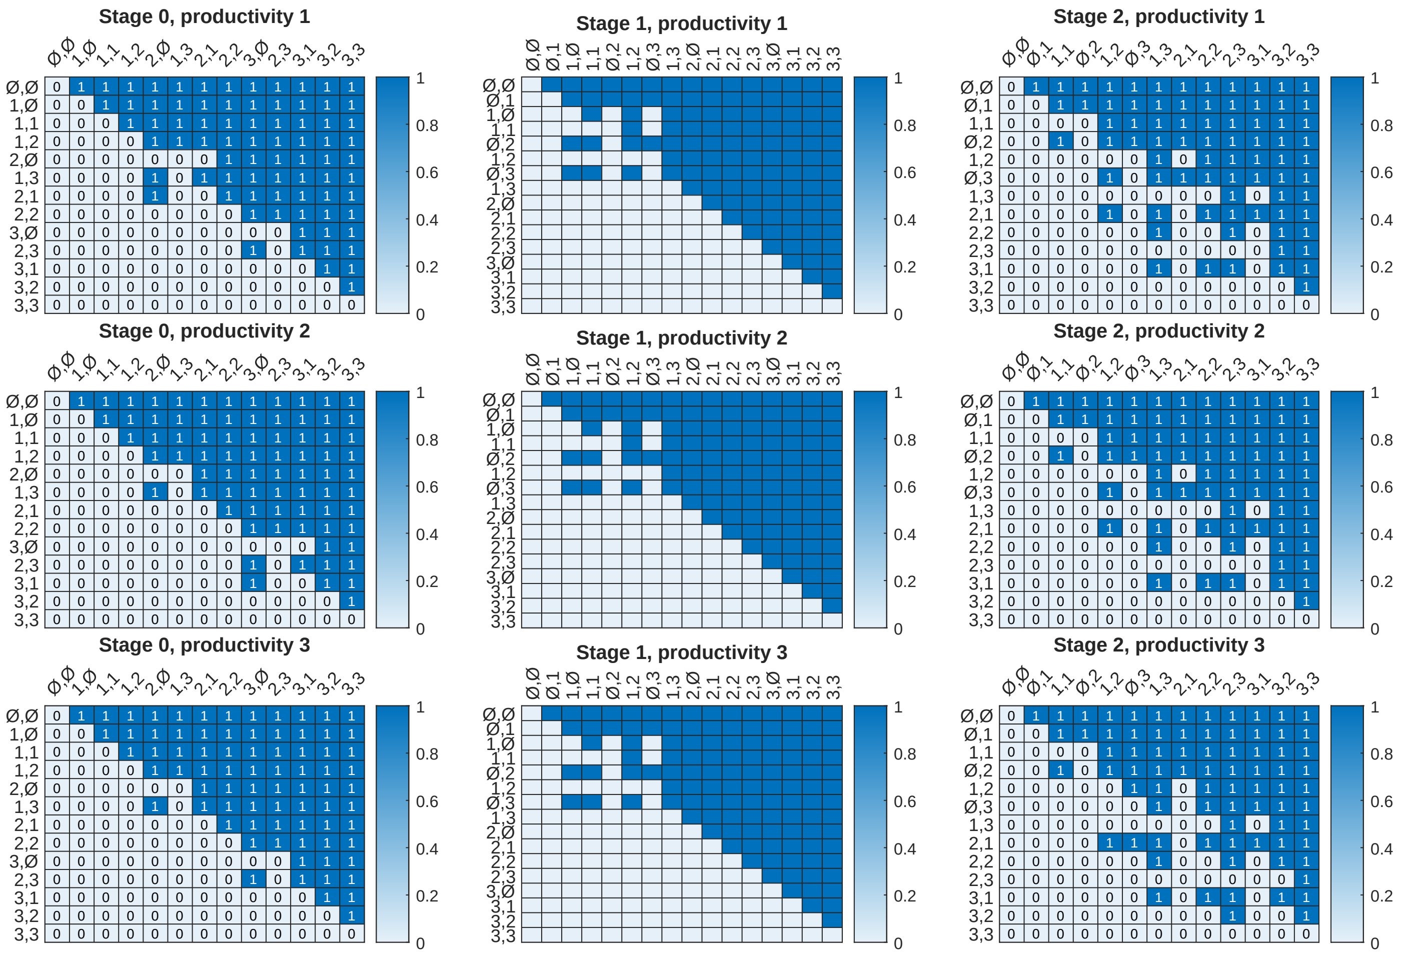

Figure A.1: Acceptance policies by firm type

A.1 Experiment 1: wider productivity ladder

Table A.1: Calibration for Experiment 1, wider productivity ladder

Parameter	Interpretation	Value
$H+1$	Number of stages	3
N_z	Number of productivity ladders	3
ρ	Discount rate	0.05
δ	Exog. separation rate	0.01
λ	Match arrival rate	0.15
σ	Final demand elasticity	3
μ	Markup	0.025
q	Relationship quality	0.025
ω	Intermediate share	0.5
Ψ	Labor disutility shifter	2
θ	Labor disutility elasticity	4
z_1	Low productivity	0
z_3	High productivity	$0.5 \rightarrow 1.0$
-	Productivity step	$0.25 \rightarrow 0.5$
p_u	Prob. of upward jump in z	0.01
p_d	Prob. of downward jump in z	0.01

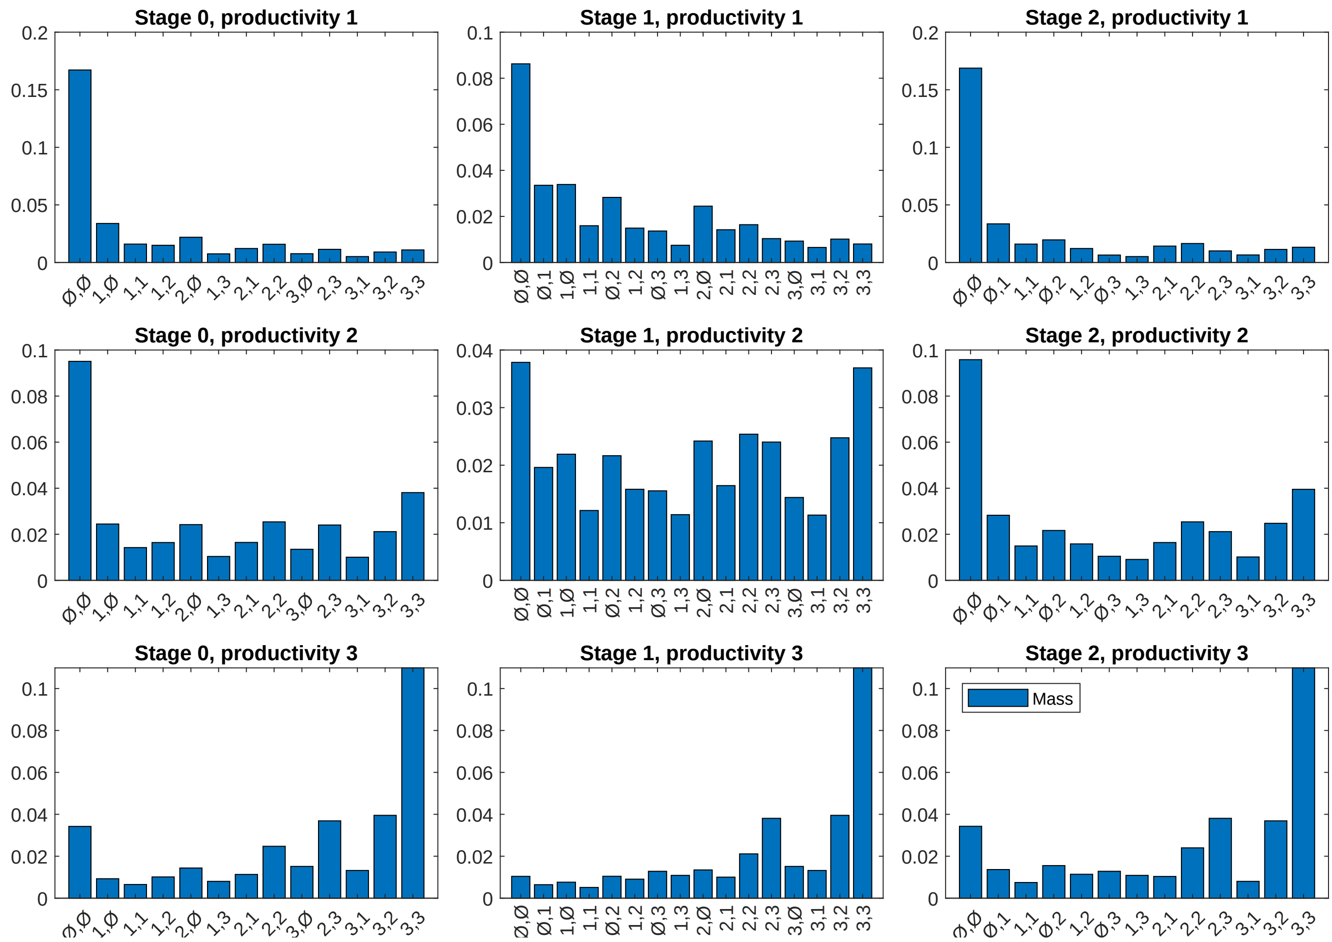

Figure A.2: Distribution of firms across the state space, for different firm types, experiment 1

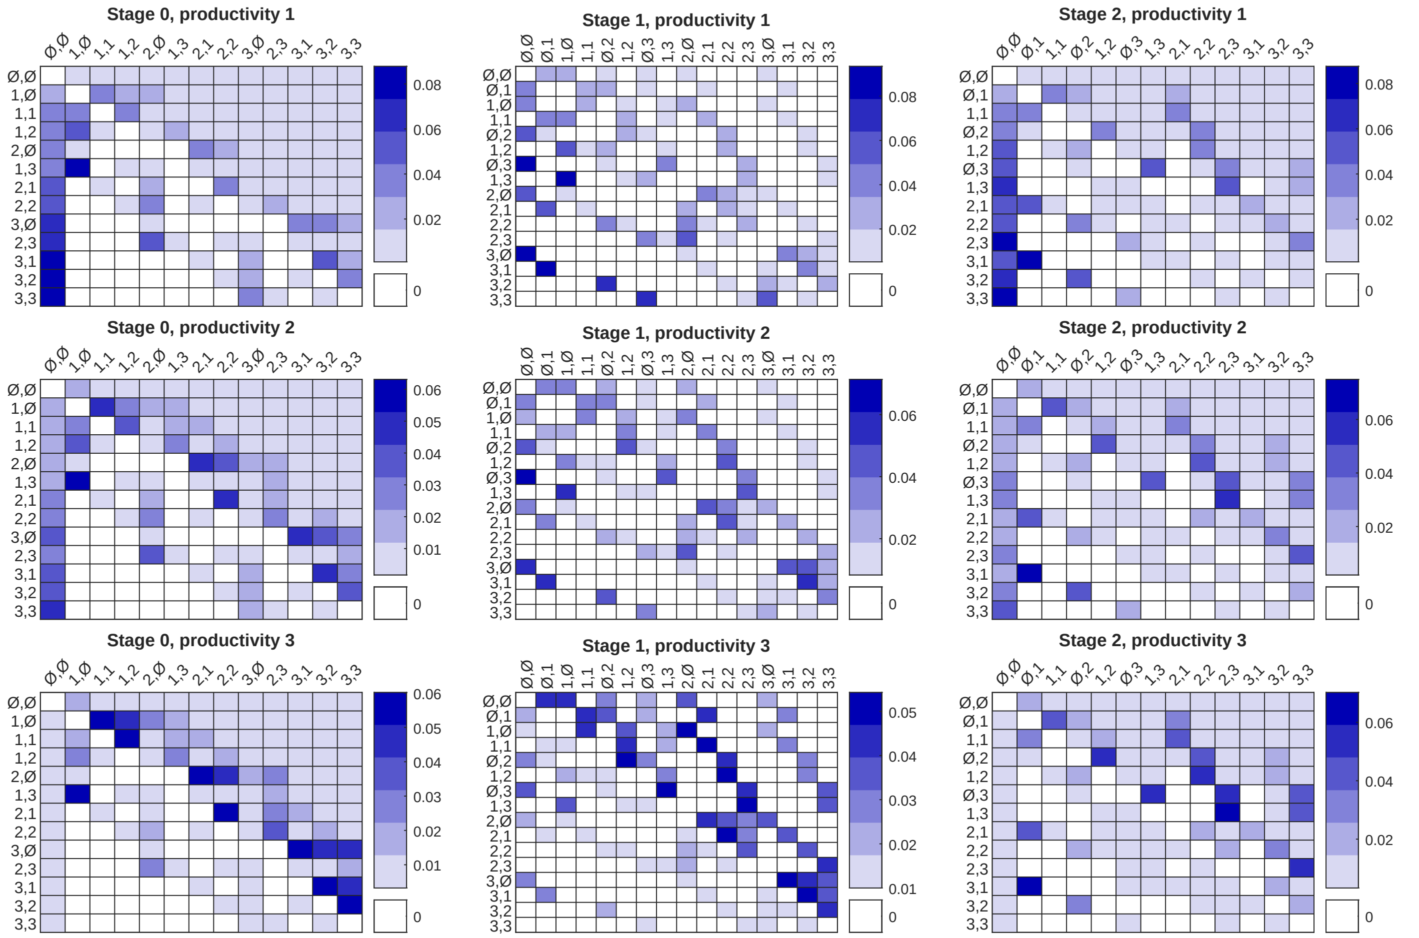

Figure A.3: Transition matrices across firm types, experiment 1

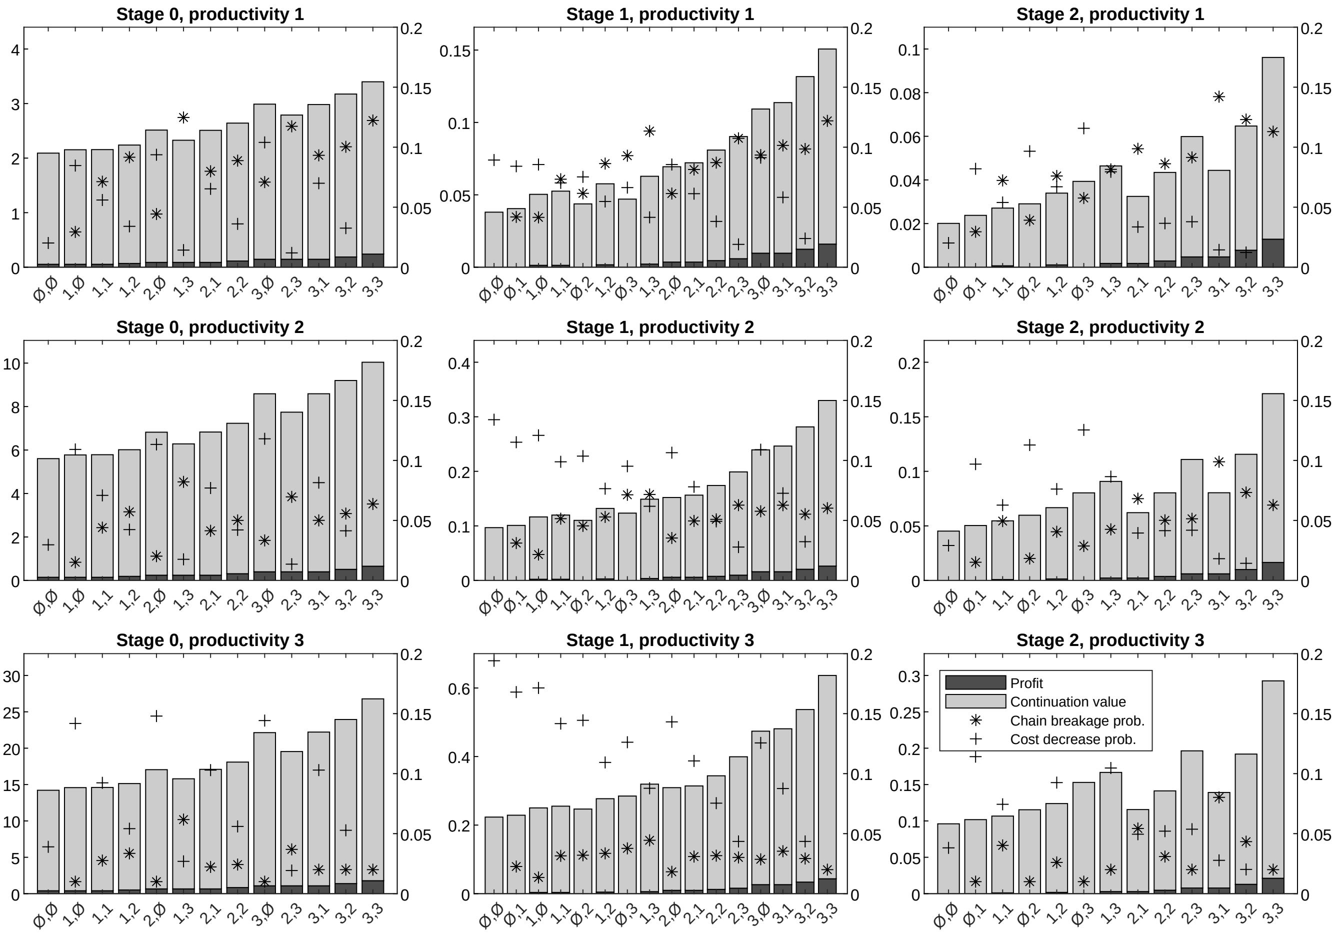

Figure A.4: Composition of the value function across firm types and states, experiment 1

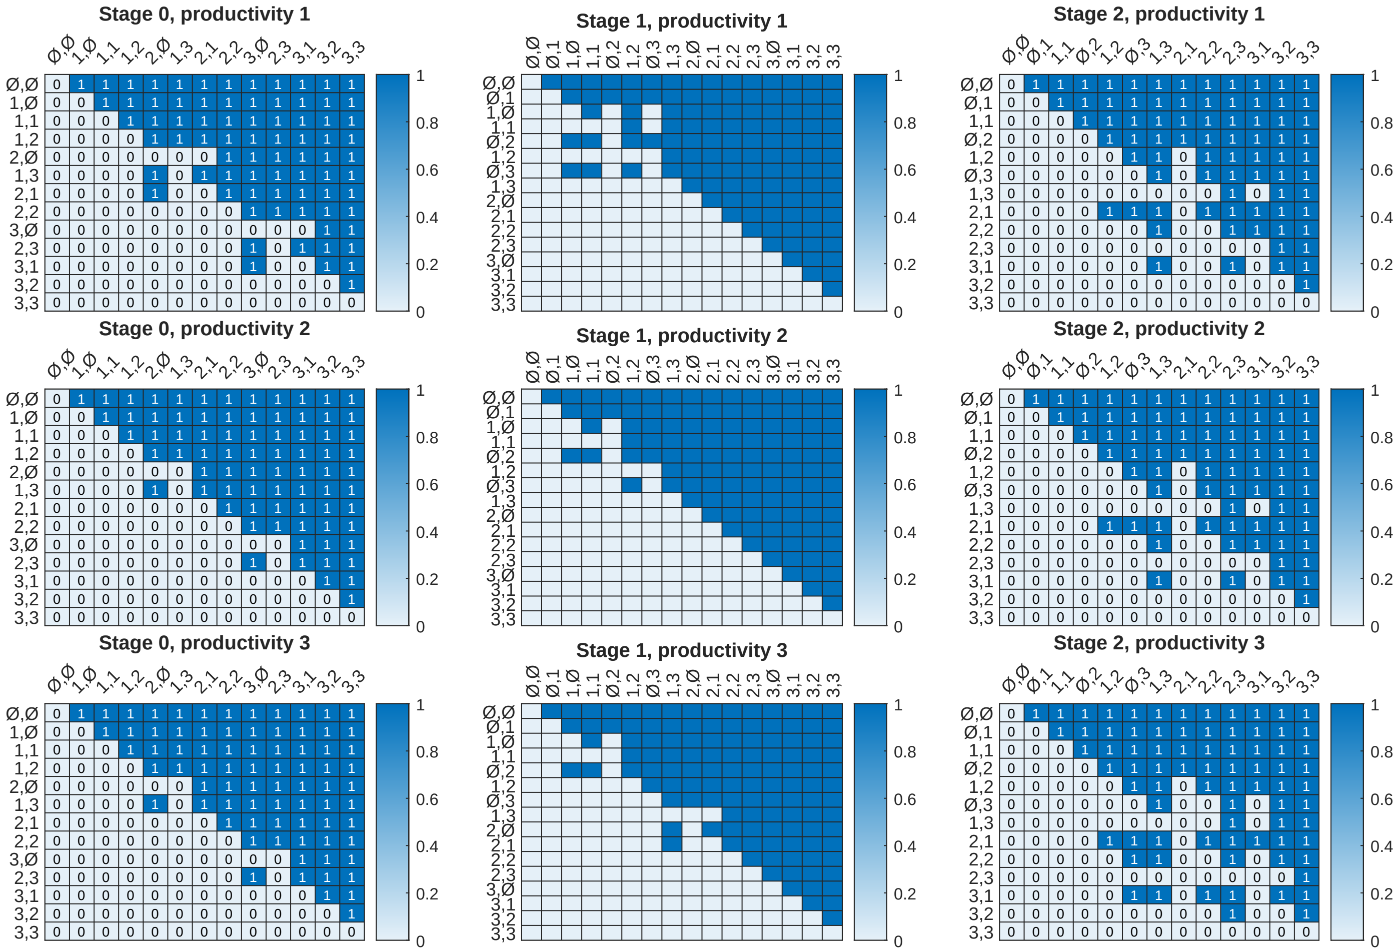

Figure A.5: Acceptance policies by firm type, experiment 1

A.2 Experiment 2: higher match arrival rate, λ

Table A.2: Calibration for Experiment 2, higher match arrival rate

Parameter	Interpretation	Value
$H+1$	Number of stages	3
N_z	Number of productivity ladders	3
ρ	Discount rate	0.05
δ	Exog. separation rate	0.01
λ	Match arrival rate	$0.15 \to 0.25$
σ	Final demand elasticity	3
μ	Markup	0.025
q	Relationship quality	0.025
ω	Intermediate share	0.5
Ψ	Labor disutility shifter	2
θ	Labor disutility elasticity	4
z_1	Low productivity	0
z_3	High productivity	0.5
-	Productivity step	0.25
p_u	Prob. of upward jump in z	0.01
p_d	Prob. of downward jump in z	0.01

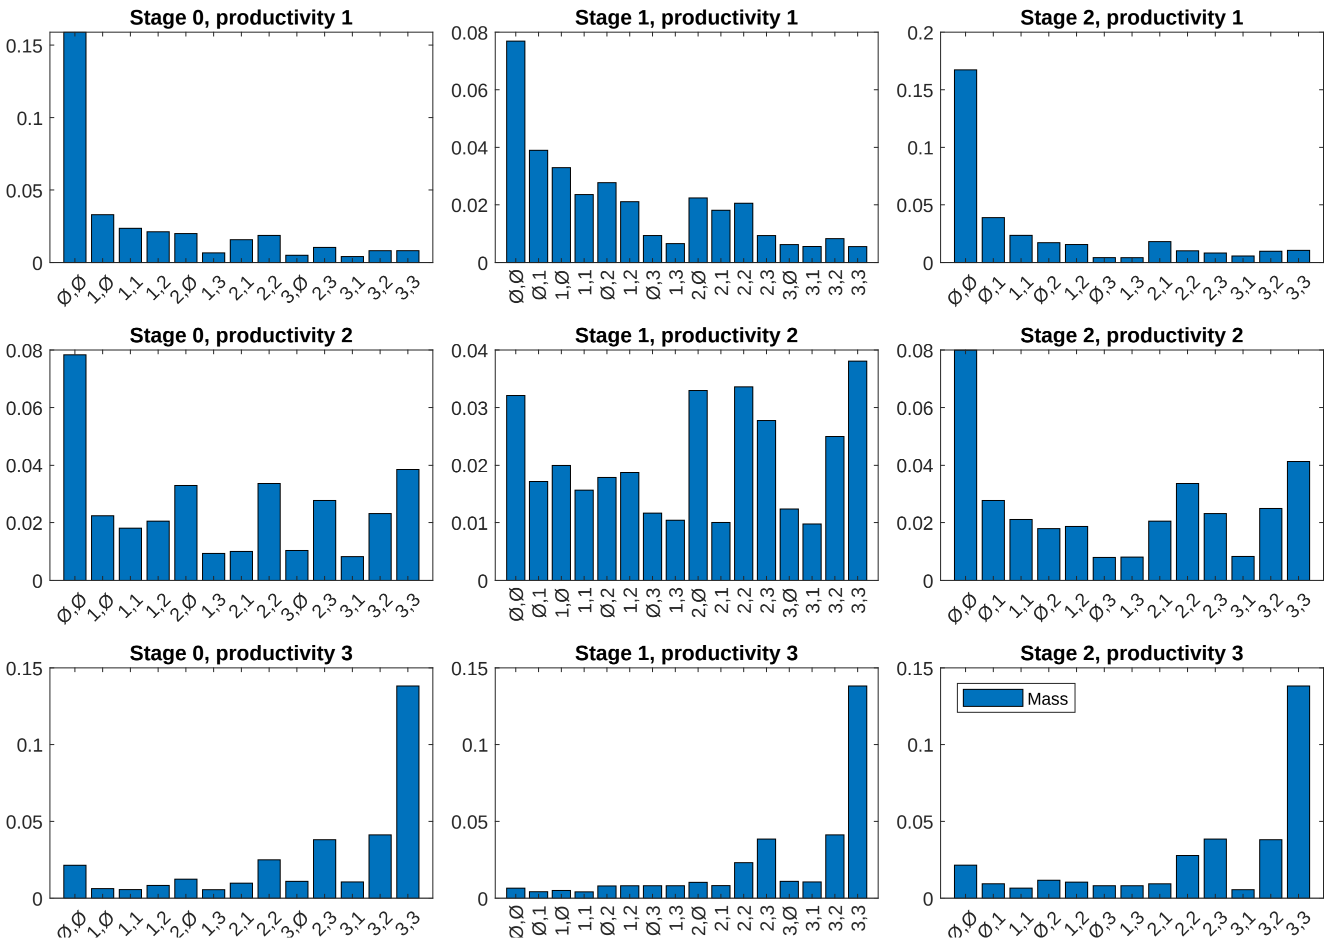

Figure A.6: Distribution of firms across the state space, for different firm types, experiment 2

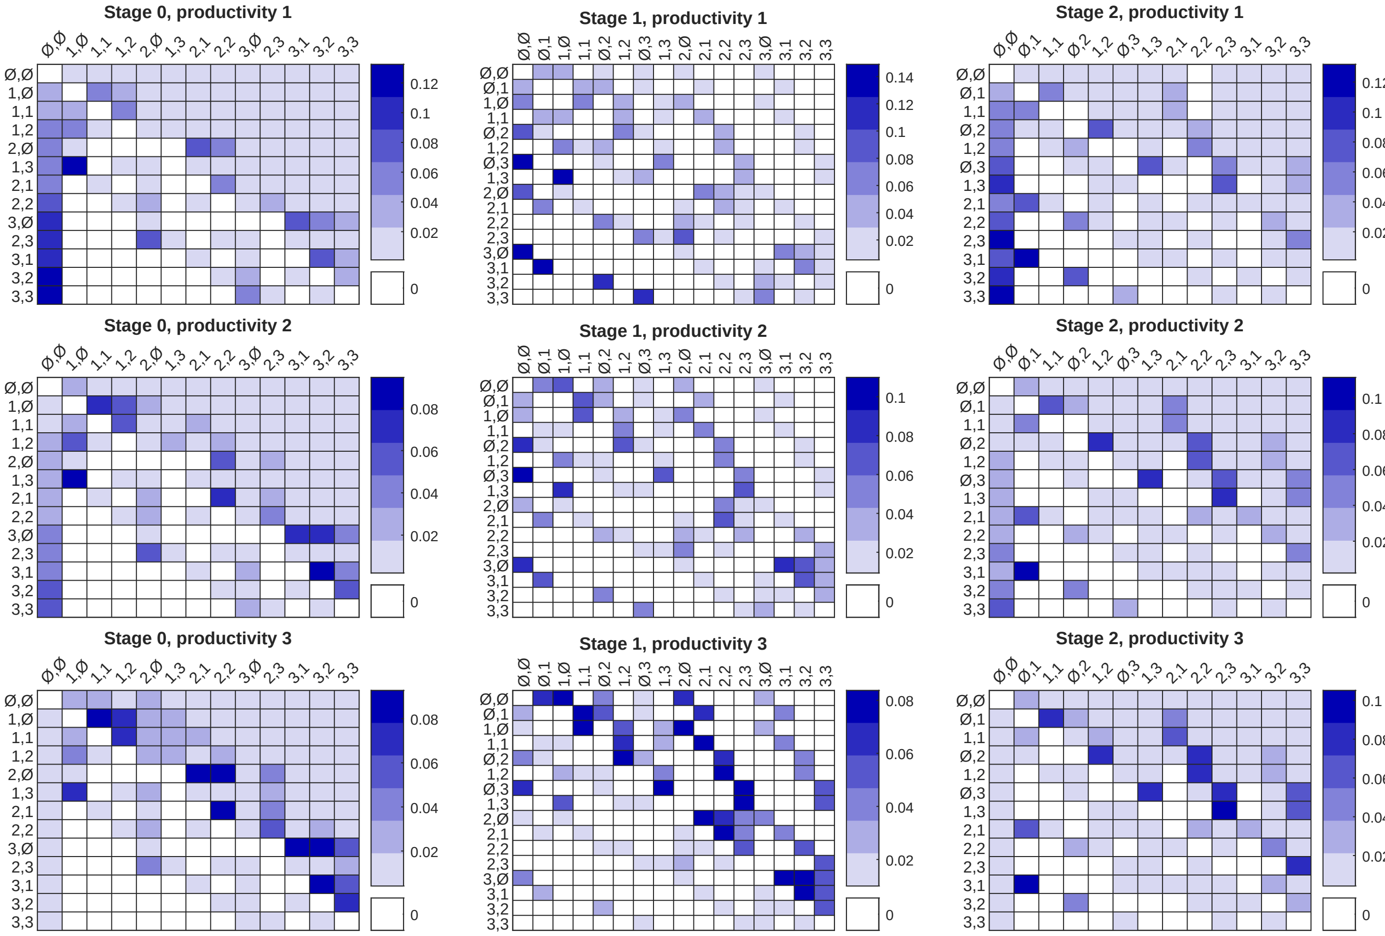

Figure A.7: Transition matrices across firm types, experiment 2

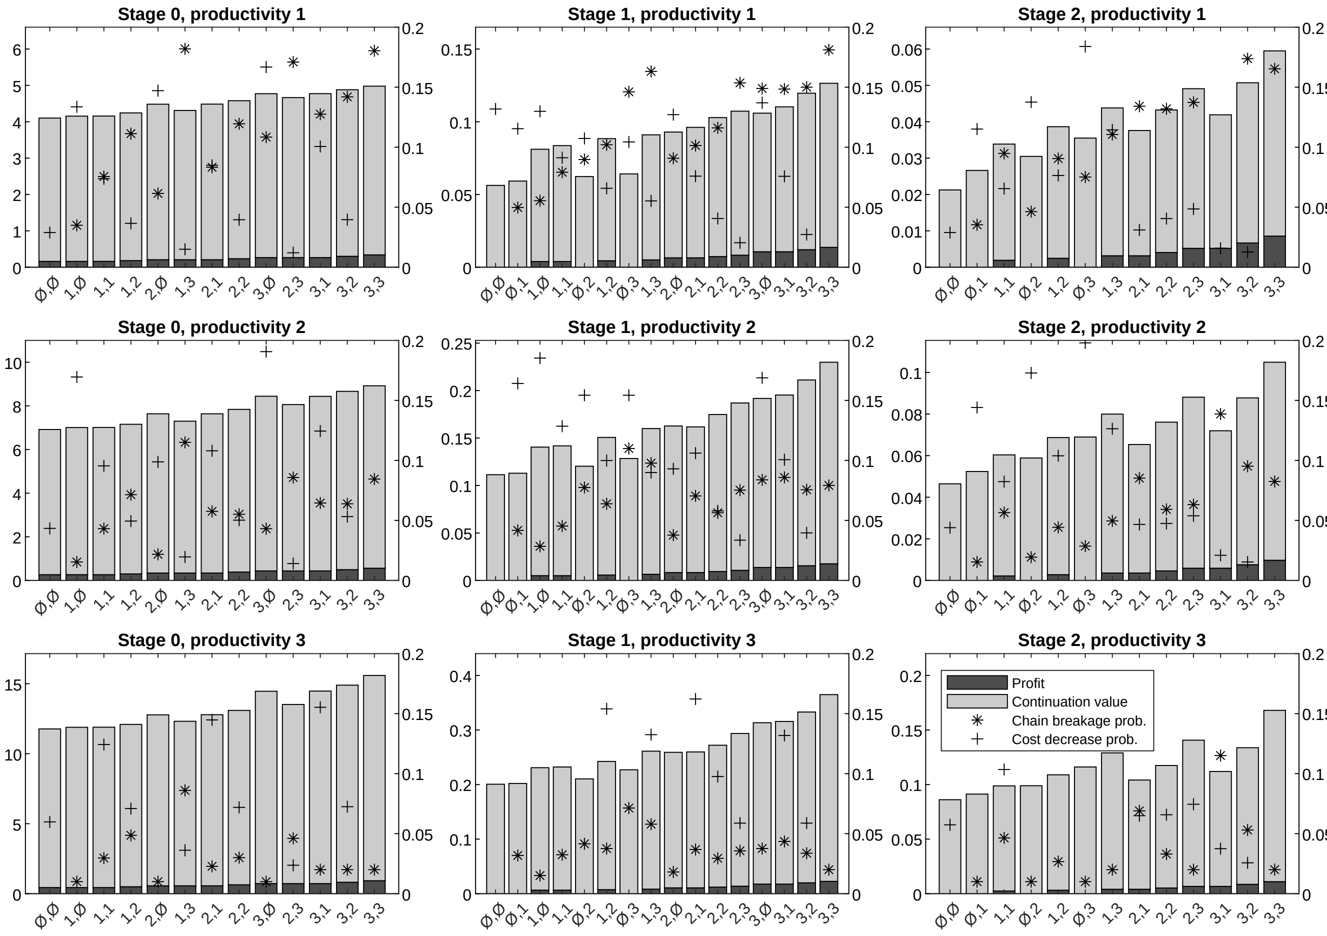

Figure A.8: Composition of the value function across firm types and states, experiment 2

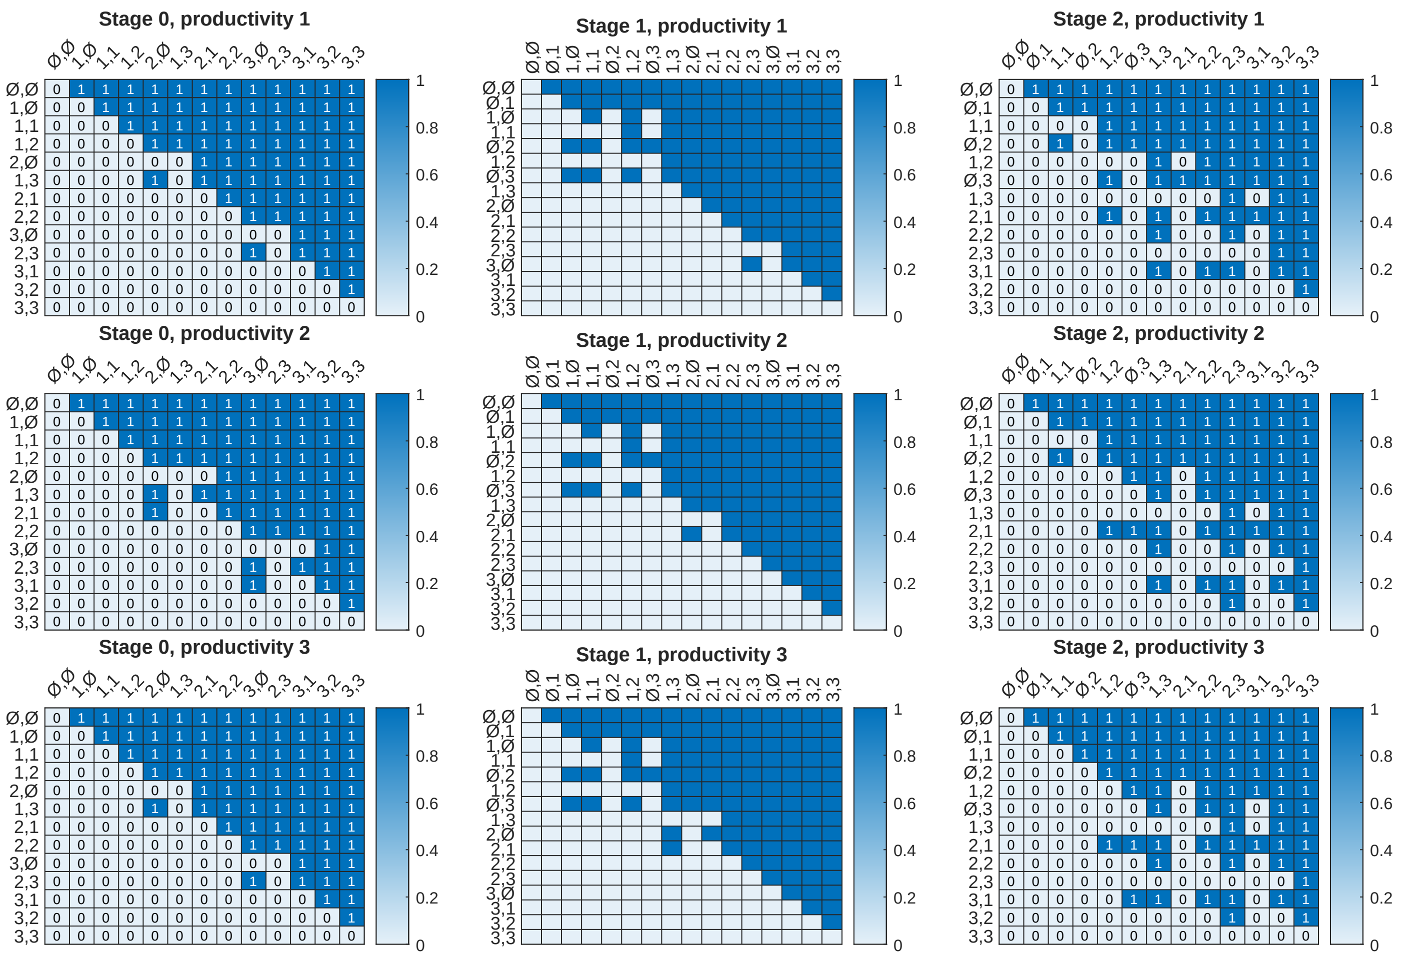

Figure A.9: Acceptance policies by firm type, experiment 2