

Bargaining foundations for price taking in matching markets*

Matthew Elliott & Eduard Talamàs[†]

Abstract

Agents make non-contractible investments before bargaining over who matches with whom and their terms of trade. When an agent is a *price taker*—in the sense that her investments do not change her potential partners’ payoffs—she has incentives to make socially-optimal investments. Across a variety of non-cooperative bargaining models featuring dynamic entry, we show that everyone necessarily becomes a price taker as bargaining frictions vanish if and only if there is a minimal amount of competition *always present* in the market. The necessity of this condition highlights that dynamic entry need not create enough competition to guarantee price taking even if agents are arbitrarily patient. The sufficiency of this condition highlights that everyone can be a price taker even in markets that appear extremely thin at every point in time.

1 Introduction

The extent to which markets create appropriate investment incentives is a fundamental question in economics (e.g., Becker 1962; Williamson 1975; Grout 1984; Grossman and Hart 1986; Acemoglu 1997; Acemoglu and Shimer 1999; Cole, Mailath, and Postlewaite 2001a; Hatfield, Kojima, and Kominers 2019; Akbarpour, Kominers, Li, and Milgrom 2020). In this

**Date Printed*: February 16, 2021.

[†]Elliott: University of Cambridge. Talamàs: IESE Business School and CEPR. We thank Marina Agranov, Benjamin Golub, George Mailath, Georg Nöldeke, Andrew Postlewaite, Larry Samuelson and József Sákovics for very useful conversations, and various audiences for helpful feedback. We acknowledge financial support from Cambridge-INET, the European Research Council under the grant “EMBED” (#757229), and the Foundations of Human Behavior Initiative. Talamàs also acknowledges support from the the Warren Center for Network & Data Sciences and the Rockefeller Foundation (#2017PRE301).

paper, we investigate the conditions under which agents are *price takers* in dynamic decentralized matching markets—in the sense that their investment choices before entering the market do not affect the prices that they face once they enter the market. When this is the case, everyone is the residual claimant on the returns of her investments, so private and social incentives to invest are perfectly aligned (e.g., Rogerson 1992; Makowski and Ostroy 1995; Nöldeke and Samuelson 2015).

Guaranteeing that everyone is a price taker in matching markets without dynamic entry generally requires extremely thick markets containing infinitely many agents of each type (e.g., Leonard 1983, Gretsky, Ostroy, and Zame 1999, Cole, Mailath, and Postlewaite 2001b; Elliott 2015). In contrast, we consider markets featuring dynamic entry, and we show that everyone necessarily becomes a price taker as bargaining frictions vanish if and only if there is a minimum amount of competition *always present* in the market. For example, in our benchmark model—in which there are arbitrarily many different types of agents and at most one match (consisting of two agents of different types) can form in each period—the relevant condition is that there are always at least two agents of each type present in the market. More generally, the relevant condition is that the stock of agents of any given type in the market is always strictly larger than the flow of this type of agents out of the market.

A core insight from the literature considering investment incentives, derived from results across a variety of settings (e.g., Grossman and Hart 1986; Hart and Moore 1990; Hosios 1990), is that how surplus is shared in equilibrium *ex post* drives investment incentives *ex ante*. This means that, fixing the investment choices of others, the bargaining environment or rule for dividing surplus impacts an individual's investment incentives. We show this is not true across a range of bargaining environments in dynamic markets (with agents coming and going) as long as there is always a minimal amount of competition present. As is standard, we find that payoffs vary with bargaining power and the details of the bargaining protocol. For example, in the context of labor markets, as workers become more patient relative to firms, their wages increase. But, irrespective of this and in contrast to the received wisdom, when a minimal amount of competition is always present, we find that everyone is the residual claimant on the returns from their investments in the limit as bargaining frictions vanish.

On the one hand, the fact that the presence of a minimal amount of competition in each period is necessary to guarantee price taking highlights that dynamic entry *per se* does not create enough competition to ensure that agents are price takers as frictions vanish. In particular, in our non-cooperative bargaining framework, the prospect of future competition is not a perfect substitute for present competition even if agents are arbitrarily patient. As

we illustrate with an example in section 2.2, this is because—when the minimal amount of competition is not always present—agents can obtain monopoly rents by waiting to match in periods in which this competition is not present.

On the other hand, the fact that the presence of a minimal amount of competition in each period is sufficient for price taking highlights that everyone can be a price taker even in markets that appear extremely thin at every point in time. To gain intuition for this result, consider an agent who invests differently from all her fellow agents of the same type. Consider first her bargaining position against agents with whom she generates more surplus than her fellows: As long as there is always a minimal amount of competition between these agents, she can play them off to make sure that they do not appropriate these potential gains. Second, consider her bargaining position against agents with whom she generates less surplus than her fellows: As long as she always faces a minimal amount of competition, they can ignore her without any payoff consequences, so they won't appropriate these potential losses either.

Our results suggest that thickness and fluidity of a market can play an important role in determining its competitiveness and efficiency. When agents of a certain type are rare and only come along infrequently, the minimal competition that we identify will not always be present, in which case agents may not be able to fully appropriate the marginal returns of their investments and, as a result, may not have incentives to invest efficiently. In contrast, when there is a regular inflow of all types of agents, the minimal competition requirements that we identify are naturally satisfied, and everyone's investments will maximize social surplus given others' investments as bargaining frictions vanish. In particular, consistent with the observation that a recent decline in US labor market fluidity has significantly reduced productivity (e.g., Decker, Haltiwanger, Jarmin, and Miranda 2020), our findings suggest a channel by which declines in labor market fluidity can reduce investment incentives.

The main contributions of this paper to the bargaining literature are threefold. First, to the best of our knowledge, this paper is the first to provide non-cooperative bargaining foundations for the canonical price taking assumption in matching markets (e.g., Cole, Mailath, and Postlewaite 2001b; Nöldeke and Samuelson 2015; Mailath, Postlewaite, and Samuelson 2017; Chiappori, Salanié, and Weiss 2017; Chiappori, Dias, and Meghir 2018; Dizdar 2018; Zhang 2020). This literature considers markets featuring a continuum of *price-taking agents* on each side to turn off appropriation problems and to investigate other sources of investment inefficiencies—like coordination failures, participation constraints, and imperfect information.¹ While our setup and focus is somewhat different, our non-cooperative

¹A branch of the search and matching literature investigates investment incentives in markets that also

bargaining foundations for price taking provide a way to gauge the conditions under which this standard competitive matching assumption is reasonable in practice.²

Second, this paper shows that considering dynamic entry significantly changes the conditions under which agents are price takers in matching markets. For example, Leonard (1983) and Gretsky, Ostroy, and Zame (1999) show that, generically, not everyone can be a price taker in finite assignment games. In the special case of unidimensional attributes and complementarities in these attributes, Cole, Mailath, and Postlewaite (2001a) provide a condition called “doubly overlapping attributes” that guarantees that everyone is a price taker in these games. In contrast, taking a non-cooperative bargaining approach and allowing for dynamic entry, we uncover a considerably less restrictive condition that is both necessary and sufficient to ensure that agents become price takers as bargaining frictions vanish without restricting attention to two-sided markets or requiring complementarities in attributes.

Finally, this paper contributes to the classical literature that investigates the extent to which the equilibrium outcomes in non-cooperative bargaining games become competitive as bargaining frictions vanish. The standard approach in this literature has been to compare the equilibrium predictions of a dynamic game as frictions vanish to the Walrasian equilibrium of an associated static economy (e.g., Rubinstein and Wolinsky 1985, 1990; Gale 1987; Binmore and Herrero 1988; Wolinsky 1988; McLennan and Sonnenschein 1991; de Fraja and Sákovics 2001; Gale and Sabourian 2005; Dávila and Eeckhout 2008; Lauer mann 2013; Polanski and Vega-Redondo 2018; Elliott and Nava 2019). Convergence to a Walrasian equilibrium in matching markets is neither necessary nor sufficient for creating good investment incentives, given others’ investments. In contrast, we ask a complementary question: Do agents become price takers as frictions vanish—in the sense that the prices that each agent faces once she enters the market are unaffected by her investment choices? When agents are price-takers in this sense they are the residual claimants on their investments and have private investment incentives, given others’ investments, that are aligned with social incentives.

Roadmap

The rest of this paper is organized as follows. Section 2 illustrates the main ideas of this paper with a relatively simple example. Section 3 describes the benchmark model and section

feature a continuum of agents on each side; see for example Acemoglu and Shimer (1999) and Bester (2013).

²A complementary literature considers how the holdup problem can be solved in bilateral matching settings even if agents are not price takers (e.g., Gul 2001; Che and Sákovics 2004).

4 presents our main result. Finally, section 5 discusses how this result extends beyond our benchmark model and its implications for investment efficiency.

2 Example

Section 2.1 illustrates how the lack of price taking in finite matching markets without dynamic entry can lead to a holdup problem. Section 2.2 illustrates how the existence of a minimal amount of competition in every period is both necessary and sufficient to guarantee that agents are price takers—and hence that holdup is not a problem—in the limit as frictions vanish.

2.1 The holdup problem in a market without dynamic entry

There are two buyers and two sellers. Each buyer can match with at most one seller, and vice versa. One agent (buyer b_1 , say) can make non-contractible investments before entering the market. As Figure 1 illustrates, the surplus that b_1 generates when matching with any seller depends on her investment choice, which is binary: If she chooses to *invest*, which costs her $1/2 < c < 1$, then her matching surplus with each seller is 2. If she chooses to *not invest*, then she does not pay any investment cost and her matching surplus with each seller is 1. Every buyer-seller match that does not involve b_1 generates 2 units of surplus. The unit surplus generated by the investment is larger than its cost c , so efficiency requires that b_1 invests.

Once b_1 has made her investment choice, bargaining occurs according to the following standard protocol (e.g., Elliott and Nava 2019). In each period $t = 1, 2, \dots$, one of the four agents is selected uniformly at random to be the proposer. If the selected agent has already left the market in a previous period, no actions are taken and no one matches in this period. Otherwise, the proposer chooses one agent on the other side of the market, and makes her a take-it-or-leave-it offer to share their gains from trade. The receiver of this offer then either accepts it, in which case the pair match and leave the market with their agreed shares; or rejects it, in which case no one matches in this period. The bargaining friction is that agents are impatient. We focus on the case in which this friction vanishes (agents' common discount factor δ goes to 1), and on strategies that only condition on the *Markov state*, which consists of the set of agents yet to match and the surpluses that they can generate.

Conditional on b_1 investing, there is an essentially unique Markov-perfect equilibrium: Each proposer makes an acceptable offer to an agent on the other side of the market, and



(a) The case in which b_1 does not invest. (b) The case in which b_1 invests.

Figure 1: Thick and thin links correspond to matching surpluses 2 and 1 respectively.

everyone’s payoffs converge to 1 as δ goes to 1.³ As we show in Appendix A, when b_1 does not invest, she can wait for b_2 to leave at a cost that vanishes as δ goes to 1, at which point b_1 is in a bilateral monopoly with another seller, and her unique subgame-perfect equilibrium converges to $1/2$ as δ goes to 1. In other words, by not investing, b_1 can guarantee a limit payoff of $1/2$, which is larger than her limit payoff $1 - c$ when she invests. Appendix A also shows that a similar problem arises with an arbitrary number n of buyers and sellers, so this example highlights how full appropriation can fail even for general-purpose investments in arbitrarily large markets without dynamic entry.

2.2 The holdup problem in a market with dynamic entry

The above example illustrates a potential source of holdup in matching markets: An agent that underinvests has the possibility of waiting until she is in a bilateral monopoly, at which point she can share the surplus losses generated by her underinvestment. This suggests that this problem might be ameliorated when new buyers and sellers enter the market over time. We now describe an example that illustrates the conditions under which this is indeed the case.

Consider the following modification of the game described above: At the beginning of each period, if there are no agents left in the market, a new buyer-seller pair enters with probability 1 and, if one buyer-seller pair is left in the market, a new buyer-seller pair enters with probability $0 \leq \rho \leq 1$. This process of dynamic entry ensures that there are always either one or two buyer-seller pairs in the market and, when $\rho > 0$, that there is always

³In fact, the concept of *iterated conditional dominance*—which solves Rubinstein’s (1982) canonical alternating-offers game (e.g., Fudenberg and Tirole, 1991, page 128)—also pins down the payoffs in this case.

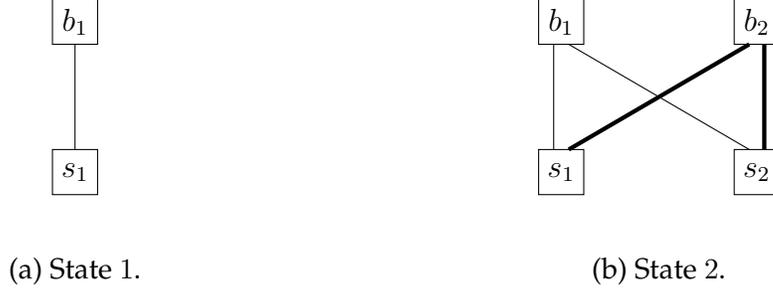


Figure 2: The two states in which non-investor b_1 is active. Conditional on no match occurring in a state 1 period, the market moves to state 2 with probability ρ . Conditional on no match occurring in a state 2 period, the market stays in state 2 with probability 1.

a strictly positive probability that there will be two buyer-seller pairs in the market in the future. Every buyer-seller match that does not involve b_1 still generates 2 units of surplus, while every buyer-seller match that involves b_1 generates 1 unit of surplus if b_1 does not invest and 2 units of surplus if b_1 invests.

As above, when b_1 invests, in every Markov-perfect equilibrium, every proposer makes an acceptable offer to an agent on the other side of the market, and everyone's payoff converges to 1 as δ goes to 1. Consider now the case in which b_1 does not invest. We describe an equilibrium in which b_1 (essentially) only matches when she is the only buyer in the market, and the sellers' limit payoffs while b_1 is in the market are strictly below 1. In particular, as long as $\rho < 1$ —so that it is *not* guaranteed that there will be two buyers and two sellers in the market at *every* point in time— b_1 's investment lifts the sellers' limit payoffs while she is in the market.

There are four different Markov states: The state 1 in which b_1 is the only buyer in the market, the state 2 in which two buyers, including b_1 , are in the market, and the two states in which b_1 has already matched and there are one and two buyer-seller pairs in the market, respectively. Figure 2 illustrates the two states in which b_1 is active.⁴ Letting w denote everyone's expected equilibrium payoff at the beginning of a period in which b_1 has already matched, we have $w = \frac{1}{4}(2 - \delta w) + \frac{3}{4}\delta w$. In particular, in this case, everyone's payoff converges to 1 as δ goes to 1. Under the following equilibrium bargaining strategies, the deviator successfully obtains monopoly power by (essentially) waiting to match until she is in a bilateral monopoly: In state 1, the two active agents make acceptable offers to each other. In state 2:

⁴We refer to the agents other than b_1 as b_2, s_1 and s_2 , even if the identity of these agents can be different in different periods.

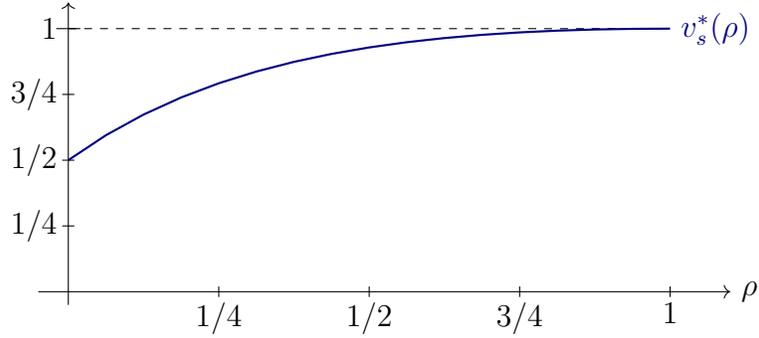


Figure 3: Sellers' limit equilibrium payoff $v_s^*(\rho)$ while b_1 is in the market conditional on b_1 not having invested.

- (i) When b_1 is selected to be the proposer, with (small) probability $0 < \pi < 1$ she makes an acceptable offer to one of the sellers (selected uniformly at random). With the remaining probability $1 - \pi$, she delays (i.e., makes an unacceptable offer).
- (ii) When b_2 is selected to be the proposer, she makes an acceptable offer to one of the sellers (selected uniformly at random), and when a seller is selected to be the proposer, she makes an acceptable offer to b_2 .

We write down the system of equations that characterizes the payoffs under these strategies in Appendix B, and we verify that, for any given $0 \leq \rho \leq 1$, there exists a threshold discount factor $\bar{\delta}(\rho) < 1$ such that, when $\delta > \bar{\delta}(\rho)$, this strategy profile is a subgame-perfect equilibrium. As δ goes to 1, the probability π that buyer b_1 matches in state 2 converges to 0. Figure 3 illustrates the limit payoff of the sellers while b_1 is active.

This example shows that dynamic entry does not necessarily ensure that everyone becomes a price taker as bargaining frictions vanish if a minimal amount of competition is not guaranteed at all times—in this case, at least two buyers and two sellers always present in the market. When sufficient competition is guaranteed (i.e., $\rho = 1$), the sellers' payoffs are independent of b_1 's investment decision. Hence b_1 is a price taker and the residual claimant on the returns of her own investments. We now turn to formalizing—in a substantially more general framework nesting this example—that guaranteeing the presence of a minimal amount of competition at every point in time is indeed sufficient to ensure that everyone becomes a price taker as bargaining frictions vanish.

3 Non-cooperative bargaining framework

There is a finite set N of types of agents. The surplus that an agent of type i and an agent of type j can generate by matching is s_{ij} .⁵ We interpret these matching surpluses as resulting from non-contractible investments that each agent must make before entering the market, and we are interested in understanding agents' incentives to deviate from these investments. Towards this goal, in this section we describe a benchmark in which no one deviates from these investments. This benchmark generalizes the bargaining model in Talamàs (2020) by relaxing the assumption that there are always the same number of agents of each type in the market.⁶ While it is common to assume stationarity for tractability, this assumption is strong when markets are relatively thin, and hence at odds with our goal of determining the limits of price taking.⁷ After characterizing the unique equilibrium outcome in this benchmark, we describe how this equilibrium outcome is affected by an arbitrary unilateral investment deviation in section 4, and we discuss the implications of this result for investment efficiency in matching markets in section 5.

3.1 The bargaining game \mathcal{G}

We extend a standard non-cooperative bargaining protocol in which proposers are randomly selected and make offers that are accepted or rejected. As we discuss later, our results generalize beyond this (still relatively broad) set up.

There are some key features of markets that we want to capture within our bargaining environment. Most importantly, we want to allow the exogenous inflow of agents into the market to depend on who is already in the market. It seems intuitive that if there is a severe shortage of one type of agent that this could induce a higher rate of entry for such agents and we want to allow for this possibility. Equally, when there are relatively few agents of a given type we might expect more entry by that type. Second, and in this case more for the sake of tractability, it is convenient to assume that the probability a given agent becomes the proposer is independent of which other agents are present. This also seems like a reasonable

⁵For simplicity, we start by assuming that only pairs of agents of different types can generate any surplus. In section 5.1, we discuss how our results extend to more general matching technologies.

⁶While the model in Talamàs (2020) allows for coalitions of arbitrary size to form and heterogeneities in preferences and proposer probabilities, for simplicity we focus on bilateral matching and on the case in which no such heterogeneities are present.

⁷For dynamic foundations for this assumption in thick markets, see, for example, Lauermaun, Nöldeke, and Tröger (2020).

benchmark.⁸

In order to capture these features we introduce a modelling device of bargaining slots. Upon entry, an agent is assigned at random to a bargaining slot associated with her type. We assume there is a finite number of slots for a given type, which amounts to assuming that the entry rate for a type of agent goes to zero when the market is saturated by agents of that type. However, we can pick any finite number as the bargaining slot capacity so we view this restriction as technical rather than economically substantive. The crucial thing is that the bargaining slots allow us to manage the size of the state space. Using bargaining slots also permits us to introduce heterogeneity among agents of the same type. For example, we might associate different bargaining slots with different proposer probabilities within type. Our results are robust to these kinds of extensions, but we omit them to reduce notation. Indeed, in the model presented, bargaining slots could simply be replaced by direct state dependent entry rates.

We now introduce the model formally. There is a common discount factor $\delta < 1$, perfect information, and common knowledge of the game. There are infinitely many periods $t = 1, 2, \dots$. For each type i , there are n_i bargaining slots. In any given period, each slot of a given type is either empty or occupied by one agent of that type. We refer to the agents occupying the slots in any given period as the *active agents* in that period, and we denote the set of slots of type i by \mathcal{L}_i . It will be convenient to abuse terminology by referring to an active agent that sits in slot ℓ as “agent ℓ ”, and to the set of active agents as a subset of $\mathcal{A} := \cup_{i \in N} \mathcal{L}_i$.

In each period $t = 1, 2, \dots$, one slot is selected uniformly at random. If the slot is empty, no match occurs in this period. Otherwise, its occupant becomes the *proposer*. The proposer chooses an active agent of another type and makes her a take-it-or-leave-it offer specifying how to split the surplus that she can generate by matching with her. The receiver of this offer either accepts or rejects. If she accepts, then they match and exit the market with their agreed shares. Otherwise, no one matches in this period.

We assume that dynamic entry ensures that there is always at least one active agent of each type, so that every agent can always choose to make an offer to an agent of any given type. The process of dynamic entry is stochastic: For each type i , at the beginning of each period in which there are $1 \leq s \leq n_i$ empty bargaining slots of type i , a number $s' \leq s$ of entrants of type i are randomly assigned to different empty slots of their corresponding type, where the number s' of empty slots that are filled is drawn according to a probability

⁸For example, a firm might have a fixed capacity for interviewing and making offers to workers that does not depend on how many workers or vacancies there are in the market.

distribution $q(i, s) : \{0, 1, \dots, s\} \rightarrow [0, 1]$.

3.2 Strategies and equilibrium

The strategy σ_a of agent a in \mathcal{A} specifies, for all possible histories at which she is active, the offer that she makes when she is the proposer (which consists of a slot and how much to offer to the agent occupying it), and her response after any possible offer that she can receive. We focus on Markov strategies that only condition on the *Markov state*, which specifies which slots are occupied. Formally, letting \mathcal{S}_i denote the set of all subsets of slots in \mathcal{L}_i , the set of Markov states is $\mathcal{S} := \prod_{i \in N} \mathcal{S}_i$. A *Markov-perfect equilibrium* is a profile of Markov strategies $\{\sigma_a\}_{a \in \mathcal{A}}$ that constitutes a Nash equilibrium in every subgame.

3.3 Equilibrium characterization

The main departure of this benchmark from the model in Talamàs (2020) is that we allow for a more general dynamic entry process that does not necessarily lead to a constant set of active agents over time, and hence leads to multiple Markov states. While this multiplicity of Markov states could potentially lead to complex equilibria that condition on how many agents of each type are in the market, in this benchmark case the equilibrium strategies are in fact insensitive to the state. Indeed, Proposition 3.1 shows that the equilibrium characterization in Talamàs (2020) remains valid in this case: There is still a unique payoff profile that is consistent with Markov-perfect equilibrium, and these payoffs are homogeneous within types and do not depend on the Markov state.

Proposition 3.1. *There is a profile w in \mathbb{R}^N such that, in every Markov perfect equilibrium of \mathcal{G} , the expected equilibrium payoff of each agent of type i at the beginning of each period is w_i .*

Proof. Consider a Markov-perfect equilibrium, and let $w_a(s)$ denote agent a 's expected payoff at the beginning of a period in which the state is s (before any agents enter in this period). On the equilibrium path, in state s every agent a accepts every offer that gives her at least $\delta w_a(s)$. For each type i , let \bar{w}_i and \underline{w}_i be the maximum and minimum $w_a(s)$ across all agents a of type i , all states s , and all Markov-perfect equilibria. Given that each agent's proposer probability is $\frac{1}{n}$,

$$\bar{w}_i \leq \frac{1}{n} \max_{j \neq i} (s_{ij} - \delta \underline{w}_j) + \frac{n-1}{n} \delta \bar{w}_i \quad \text{and} \quad \underline{w}_i \geq \frac{1}{n} \max_{j \neq i} (s_{ij} - \delta \bar{w}_j) + \frac{n-1}{n} \delta \underline{w}_i.$$

Combining these two inequalities gives $\bar{w}_i = \underline{w}_i$ for every type i . To see this, consider a type i for which $\bar{w}_i - \underline{w}_i$ is largest, and let j maximize $s_{ij} - \delta \underline{w}_j$. We have that $\bar{w}_i - \underline{w}_i \leq$

$\frac{1}{n}(\bar{w}_j - \underline{w}_j) + \frac{n-1}{n}\delta(\bar{w}_i - \underline{w}_i)$ which implies that $0 \leq \bar{w}_i - \underline{w}_i \leq \chi(\bar{w}_j - \underline{w}_j)$ for $\chi < 1$, so $\bar{w}_i = \underline{w}_i$. Hence, in every Markov-perfect equilibrium, the equilibrium payoff $w_a(s)$ must be the same across all agents a of type i and all states s , and these payoffs $(w_i)_{i \in N}$ must solve

$$(1) \quad w_i = \frac{1}{n} \max_{j \neq i} (s_{ij} - \delta w_j) + \frac{n-1}{n} \delta w_i \text{ for all } i \text{ in } N.$$

As shown in Talamàs (2020), system (1) admits a unique solution. \square

Our focus is on understanding the extent to which agents become price takers in this setting as the discount factor δ goes to 1. Corollary 3.2 follows from the equilibrium characterization in Talamàs (2020), and shows that the limit equilibrium payoffs are the unique Nash credible profile, defined as follows.

Definition 3.1. For every pair $i, j \in N$ and every profile θ in $\mathbb{R}_{\geq 0}^N$, let the *Nash bargaining solution* $\alpha(ij, \theta)$ be the unique solution of

$$(2) \quad \underset{(s_i, s_j)}{\operatorname{argmax}} \ s_i s_j \text{ subject to } s_{ij} \geq s_i + s_j \text{ and } (s_i, s_j) \geq (\theta_i, \theta_j)$$

if $s_{ij} \geq \theta_i + \theta_j$, and (θ_i, θ_j) otherwise.

Definition 3.2. The payoff profile β in \mathbb{R}^N is *Nash credible* if

$$\beta_i = \max_{j \in N} \alpha(ij, \beta_{-i}) \text{ for every } i \text{ in } N,$$

where β_{-i} denotes the profile β after setting its i^{th} entry to 0.

Corollary 3.2. *As the discount factor δ goes to 1, the unique equilibrium payoff profile w converges to the unique Nash credible profile v .*

In words, the limit payoff v_i of a given type i is the maximum that this type can justify as resulting from the Nash bargaining solution in some match while honoring the others' outside options—determined by the maximum that they can themselves justify in this way. Going forward it will be useful to draw a contrast between the individual investment incentives of an agent and the investment incentives that would be present if all agents of the same type could perfectly coordinate their deviations. The payoff characterization above implies that, in general, the agents of any given type are—as a group—not price takers. If all agents of a given type i make investments that increase the surplus s_{ij} of the match that determines their limit payoff (i.e., $v_i = \alpha(ij, v_{-i})$), then if j 's payoff is also determined by the match ij (i.e., $v_j = \alpha(C, v_{-j})$), these investments lift the limit payoffs of all agents of type

j as well. Hence the agents of type i do not fully appropriate the marginal returns of their investments from a joint deviation.

However, in this paper we are interested in understanding whether individual agents—as opposed to all the agents of a given type as a group—are price takers. For this, we need to understand how the equilibrium outcome changes after a *unilateral* investment deviation by one agent of an arbitrary type—instead of investment deviations that change the surplus that all the agents of a given type can generate. We investigate this question next, and we obtain a sharp answer that is not sensitive to the details of the bargaining protocol and that is qualitatively different from the answer to the related question of whether all the agents of a given type are price takers as a group.

4 Unilateral investment deviation

We are interested in understanding agents’ incentives to deviate from the investments that lead to the matching surpluses of the benchmark game above. To investigate this question, we consider an agent d —which we think of as a *deviator*—of an arbitrary type i that can change her matching surpluses with others. In particular, agent d chooses a number s_{dj} for each type $j \neq i$, which is the surplus that d can generate by matching with any agent of type j . Taking d ’s choices as given, we characterize the equilibrium of game $\hat{\mathcal{G}}$, which is exactly like game \mathcal{G} except that, in period $t = 1$, agent d is active. The set of active agents is now a subset of $\hat{\mathcal{A}} := \mathcal{A} + d$, and each Markov state specifies which slots are filled by non-deviators and whether the deviator d is active or not.⁹ Formally, letting $\hat{\mathcal{S}}_i$ denote the set of all subsets of $\mathcal{L}_i + d$ that contain at most n_i elements, the set of Markov states is $\hat{\mathcal{S}} := \prod_{j \in N-i} \mathcal{S}_j \times \hat{\mathcal{S}}_i$.

4.1 Characterization of the equilibrium in $\hat{\mathcal{G}}$

We start by stating that the equilibrium of interest exists.

Proposition 4.1. *The game $\hat{\mathcal{G}}$ admits a Markov-perfect equilibrium.*

Proof. The argument is analogous to the one behind the equilibrium existence proof in Elliott and Nava (2019), and we relegate it to Appendix C. \square

Once the deviator d leaves the market, game $\hat{\mathcal{G}}$ is identical to game \mathcal{G} , so we know that the limit equilibrium payoffs from that point on are given by v . However, in general, we don’t

⁹For notational simplicity, we write $\mathcal{A} + d$ for $\mathcal{A} \cup \{d\}$, etc.

know whether Markov-perfect equilibrium payoffs in $\hat{\mathcal{G}}$ are unique or homogeneous within types or states while d is active. Nevertheless, our main result, Theorem 4.2 below, shows that if the process of dynamic entry ensures that there are always at least two active agents of each type, then the limit payoff of every agent other than d must converge to her limit payoff once the deviator d has left the market. Hence, in the limit as bargaining frictions vanish, the deviator d 's choices do not affect the prices that she has to pay others to match once she enters the market, so she fully appropriates the marginal returns of her investment deviation.

Theorem 4.2. *Consider a sequence of discount factors converging to 1 and associated Markov-perfect equilibria of $\hat{\mathcal{G}}$ with payoffs converging to $(\hat{v}_a(s))_{a \in \hat{\mathcal{A}}, s \in \hat{\mathcal{S}}}$. If the process of dynamic entry ensures that in each period there are at least two active agents of each type, then for every Markov state s and every agent $a \neq d$ of any type j , $\hat{v}_a(s) = v_j$.*

Proof. Let $\hat{w}_a(s)$ denote agent a 's expected payoff at the start of a period in which the state is s (before anyone enters in this period). On the equilibrium path, when the state is s , every agent accepts every offer that gives her at least $\delta \hat{w}_a(s)$. Let $\hat{v}_a(s)$ denote the limit of $\hat{w}_a(s)$ as δ goes to 1. We start by showing that

$$(3) \quad \hat{v}_{a_j}(s) + \hat{v}_{a_k}(s) \geq s_{jk} \text{ for all states } s \text{ and any two agents } a_j, a_k \neq d \text{ of types } j \text{ and } k \neq j,$$

because this implies that it is enough to show that

$$(4) \quad \hat{v}_{a_j}(s) \leq v_j \text{ for every state } s \text{ and every agent } a_j \text{ of any type } j.$$

Indeed, (4) implies that, for every agent $a_j \neq d$ of any type j , letting k maximize $s_{jk} - v_k$ and a_k be an arbitrary agent of type k ,

$$v_j = s_{jk} - v_k \leq s_{jk} - \hat{v}_{a_k}(s) \leq \hat{v}_{a_j}(s) \text{ for every state } s.$$

where the last inequality follows from (3).

To see (3), consider a state s with the lowest $\hat{v}_{a_j}(s) + \hat{v}_{a_k}(s)$, and suppose for contradiction that $\hat{v}_{a_j}(s) + \hat{v}_{a_k}(s) < s_{jk}$. Once the deviator d leaves, the limit payoffs are given by v , and $v_j + v_k \geq s_{jk}$, so the deviator d is active in state s , and either $\hat{v}_{a_j}(s) < v_j$ or $\hat{v}_{a_k}(s) < v_k$. Assume without loss of generality that $\hat{v}_{a_j}(s) < v_j$. Starting from state s , consider agent a_j 's strategy that deviates from her equilibrium strategy while d is active as follows: Reject every offer, and offer slightly above $\hat{v}_{a_k}(s')$ to agent a_k as soon as she is the proposer in any state s' in which agent a_k is active. Such offers are accepted with probability one, and the cost of waiting to either be able to make this offer or to see the deviator d leave converges to 0 as δ

goes to 1. Since both the expected value of $s_{jk} - \hat{v}_{a_k}(s')$ (which a_j obtains in the limit if she is ends up being able to make such an offer) and v_j (which a_j obtains in the limit if the deviator d leaves before a_j is able to make such an offer) are strictly larger than $\hat{v}_{a_j}(s)$, we conclude that a_j 's limit expected payoff under this deviation is strictly above $\hat{v}_{a_j}(s)$ as δ goes to 1, a contradiction.

It only remains to show (4). For any type j , consider an agent $a \neq d$ and a state r such that

$$(5) \quad \hat{v}_a(r) \geq \hat{v}_{a_j}(r') \quad \text{for every agent } a_j \neq d \text{ of type } j \text{ and every state } r'.$$

Given (5) and the fact that j is chosen arbitrarily, it is enough to show that $\hat{v}_a(r) \leq v_j$. Suppose for contradiction that $\hat{v}_a(r) > v_j$. First note that

$$(6) \quad \hat{v}_d(r) + \hat{v}_{a_j}(r) \geq s_{dj} \quad \text{for every agent } a_j \neq d \text{ of type } j \neq i.$$

To see this, suppose for contradiction that $\hat{v}_d(r) < s_{dj} - \hat{v}_{a_j}(r)$ for some agent a_j of type j . Starting from state r , consider agent d 's strategy that consists of rejecting every offer, and offering slightly above $\hat{v}_{a_j}(r')$ to agent a_j as soon as agent d is selected to be the proposer in any state r' in which agent a_j is active. Such offers are accepted with probability one, and the waiting cost of this strategy converges to 0 as δ goes to 1. Moreover, given (5), the expected value of $s_{dj} - \hat{v}_{a_j}(r')$ is strictly larger than $\hat{v}_d(r)$, so d 's limit expected payoff of following this strategy is strictly above $\hat{v}_d(r)$ as δ goes to 1, a contradiction.

By (3) and (6), a can never obtain more than a limit payoff of $\hat{v}_a(r)$ by matching, so starting from state r , for all discount factors δ sufficiently close to 1, agent a must match before d leaves with probability one (otherwise, $\hat{v}_a(r)$ would be a strictly convex combination of v_j and numbers that are weakly smaller than $\hat{v}_a(r)$, a contradiction of $\hat{v}_a(r) > v_j$). Consider an arbitrary state r' in which, for all δ sufficiently close to 1, agent a matches with positive probability in state r' , obtaining $\hat{v}_a(r') = \hat{v}_a(r)$. The fact that a matches with positive probability in state r' for all δ sufficiently close to 1 implies that either $\hat{v}_{a_k}(r') + \hat{v}_a(r') \leq s_{jk}$ for some agent a_k of some type $k \neq j$, or that $\hat{v}_d(r') + \hat{v}_a(r') \leq s_{dj}$, which combined with (3) and (6), respectively, give that $\hat{v}_{a_k}(r') + \hat{v}_a(r') = s_{jk}$ in the former case and $\hat{v}_d(r') + \hat{v}_a(r') = s_{dj}$ in the latter case. The assumption that there are always at least two active agents of each type ensures that there is an agent $a_j \neq a$ of type j . For any such agent a_j , we have that $\hat{v}_{a_j}(r')$ is a strictly convex combination of v_j and numbers that are weakly smaller than $\hat{v}_a(r)$, so $\hat{v}_{a_j}(r') < \hat{v}_a(r')$, and hence either $\hat{v}_{a_k}(r') < s_{jk} - \hat{v}_{a_j}(r')$ or $\hat{v}_d(r') < s_{dj} - \hat{v}_{a_j}(r')$, a contradiction of (3) or (6). \square

Theorem 4.2 shows that the necessary condition identified by the example in section 2.2 to guarantee that everyone is a price taker as bargaining frictions vanish—namely, that there

are always at least two active agents of each type—is also sufficient to guarantee that every agent is a price taker in our benchmark model. We discuss this minimal competition requirement in more detail in section 5.1 below.

5 Discussion

This section discusses: (i) how our results extend to more general matching technologies as well as to alternative bargaining protocols; (ii) the implications of our main result on price-taking for investment efficiency in matching markets; and (iii) what this suggests about investment incentives in different markets.

5.1 Extensions

For simplicity, our benchmark model above assumes that productive matches are bilateral, but our main result goes through under more general matching technologies. For example, letting $C \subseteq N$ be the set of types that are represented in a given set of agents, suppose that the surplus that these agents generate when they match is $y(C)$. In this case, a proposer of type j can choose any coalition $C \subseteq N - j$, as well as an active agent of each type k in C , and proposes how to split the surplus $y(C + j)$ that they can generate by matching. The selected agents then respond in a pre-specified sequence, and this match forms in this period if and only if all of them accept. If, once the deviator has left, we restrict attention to stationary strategies, an argument analogous to the one behind Proposition 3.1 shows that the equilibrium payoffs in this case satisfy

$$w_i = \frac{1}{n} \max_{i \in C \subseteq N} \left(y(C) - \sum_{j \in C-i} \delta w_j \right) + \frac{n-1}{n} \delta w_i \text{ for all } i \text{ in } N,$$

which, as shown by Talamàs (2020), admits a unique solution. In this case, a straightforward extension of the argument in the proof of Theorem 4.2 shows that agent d 's choices do not affect the limit prices that she faces once she enters the market either as long as there are always at least two agents of each type in the market.

Our price taking result also extends to the case in which coalitions that contain more than one agent of each type are feasible. For example, consider the case in which productive coalitions can contain up to $m_i \geq 1$ agents of each type i . We can embed this case in the framework above by enlarging the set N of types to an artificial set \hat{N} that contains m_i copies of each type i . In this case, Proposition 3.1 requires that there are at least m_i agents of each

type i in N , and Theorem 4.2 requires that there are at least $m_i + 1$ agents of each type i in N . This extension highlights that the key condition behind our price taking result is that there is always a minimal amount of competition present in the market, in the sense that in each period there is at least one agent of each type that is in the market who does not match. In the baseline case in which at most one agent of each type can match, this requires that there are at least two active agents of each type. More generally, the relevant condition is that the inflow of agents into the market is sufficiently high so that the stock of agents in the market is always larger than the flow of agents out of the market.

For concreteness, we have considered a particular bargaining protocol in which proposers can strategically choose which matches to propose. This has provided us with a simple benchmark for the case in which no one chooses to deviate, because it has allowed us to leverage the equilibrium characterization in Talamàs (2020). However, taking as given the limit payoffs once the deviator has left, our price taking result goes through under any bargaining protocol in which every agent has a strictly positive probability of being able to make an offer to any other agent in each period. For example, consider the following alternative protocol, which is a version of the standard random-matching protocol used in the literature of bargaining in stationary markets (e.g., Rubinstein and Wolinsky 1985, Manea 2011, Nguyen 2015). Suppose that the bargaining protocol is exactly as in the coalitional extension of our model above except that, in each period, not only a slot is selected uniformly at random to be the proposer, but also a coalition $C \subseteq N$, and the proposer, of type j say, can only choose an active agent of each type in $C - j$, and make a take-it-or-leave-it offer specifying how to split the surplus that she can generate by matching with them. Continuing to restrict attention to stationary strategies once the deviator has left, an argument analogous to the one behind Proposition 3.1 shows that the system of equilibrium payoffs under this protocol must be such that, for each agent a of type j and each state s , $w_a(s) = w_j$, where $(w_j)_{j \in N}$ solves

$$w_i = \frac{1}{n} \sum_{i \in C \subseteq N} \frac{1}{2^n} \max \left[\delta w_i, y(C) - \sum_{j \in C-i} \delta w_j \right] + \frac{n-1}{n} \delta w_i \text{ for every } i \text{ in } N,$$

which, as shown by Nguyen (2015), admits a unique solution. As emphasized by Talamàs (2020), both the coalitions that form and how the resulting surplus is split under this alternative protocol are different. However, an argument analogous to the one behind the proof of Theorem 4.2 shows that our price taking result holds under this protocol as well. In other words, while this alternative bargaining protocol makes a difference for the point predictions of the theory, it does not change the fact that each agent is a price taker in the limit as bargaining frictions vanish as long as there are always at least two active agents of each

type.

We conclude that our price taking result applies to a wide variety of markets of interest. For example, in labor markets in which homogeneous workers and firms match one to one, both workers and firms are guaranteed to be price takers as bargaining frictions vanish if and only if there are always at least two workers and two firms available to match—as in the case of $\rho = 1$ in the example of section 2.2. If workers and firms are instead heterogeneous, and each firm can hire at most m workers, then both workers and firms are guaranteed to be price takers as bargaining frictions vanish if and only if there are always at least two firms of each type and at least m workers of each type available to match.

5.2 Implications for investment efficiency

For simplicity, we have obtained our price taking result in a framework in which an arbitrary agent can pursue an arbitrary investment deviation. We now outline the implications of this result for the more involved case in which all agents that enter the market in any given period simultaneously choose their investments.

Suppose that each agent of type i chooses an investment from a set K_i of elements in \mathbb{R}^{m_i} , where $m_i \geq 1$. The cost of investment x_i is given by $c(x_i)$. The matching surplus of a two of agents of types i and j and whose investments are given by k in $K_i \times K_j$ is given by $y(k)$.

Let us start by assuming that agents choose their investments before observing any of the others' actions. This is a reasonable benchmark to analyze situations in which agents must sink their investments well before knowing who will be their potential matching partners. For simplicity, let us focus on type-symmetric equilibria—in which all agents of the same type choose the same investment. Proposition 3.1 implies that, for each type-symmetric investment profile $x := (x_i)_{i \in N}$ and each type i , there exists $w_i(x) > 0$ such that, conditional on every agent of type i choosing investment x_i , the expected equilibrium (gross) payoff at the beginning of each period of each agent of type i is $w_i(x)$. We denote the limit of $w_i(x)$ as δ goes to 1 by $v_i(x)$.

In this case, our main result, Theorem 4.2, implies that, if there are always at least two active agents of each type, in order to be able to implement the investment profile $x := (x_i)_{i \in N}$ in a type-symmetric equilibrium for all sufficiently high discount factors, each agent of each type i must find it optimal to choose x_i taking as given the others' limit payoffs; that

is,

$$(7) \quad x_i \in \operatorname{argmax}_{z_i \in K_i} \left[\max_{j \neq i} [y(x_j \times z_i) - v_j(x)] - c(z_i) \right] \text{ for each } i \text{ in } N.$$

In other words, the equilibrium investment profile x in the limit as frictions vanish must be *constrained efficient*—in the sense that no agent, *taking others' limit payoffs as given and free to choose whom to match with*, has a profitable investment deviation.

Note that no constraints of note had to be placed on the investment technology. For example, investments can be complements, substitutes or some combination. Of course, constrained efficiency does not rule out the possibility of miss-coordination on inefficient equilibria, which will sometimes be a possibility. It only guarantees that a planner would not want to change one agent's investment, holding fixed the investments of others.

We conclude with a caveat: The assumption that agents choose their investments before observing the previous investments choices is an important one, because it ensures that an investment deviation by one agent does not trigger further investment deviations by other agents. Indeed, Appendix D describes a simple example that shows how agents need not be price takers if they are able to identify investment deviations before choosing their investments, because in this case, an agent's investment deviation can trigger further investment deviations by others.

5.3 Interpretation

Modeling market clearing with a fixed, finite number of agents in the absence of entry suggests that bargaining outcomes will be sensitive to the details of the bargaining protocol, and these in turn will influence investment incentives. Moreover, it will not be possible to simultaneously give good investment incentives to both sides of the market—not all agents can be the residual claimants on their investments. We show that these conclusions don't hold when there is dynamic entry if and only if a minimal competition condition is met.

To interpret this result—and to use it to provide some insight into investment incentives in actual matching markets—requires taking a view on the types of markets in which the minimal competition condition is more likely to be met. In markets where there is a lot of heterogeneity among agents in terms of their characteristics and investment opportunities, there will be a lot of different types of agents and the minimal competition condition will be harder to satisfy. This suggests that the condition is more likely to be satisfied in low skill markets than high skill markets. Similarly, markets that are thick and/or fluid (with a high

inflow rate) are more likely to satisfy the minimal competition condition all else equal than markets that are thin and/or viscous.

Appendices

A Details omitted from subsection 2.1

Suppose for contradiction that there is a Markov-perfect equilibrium in which b_1 does not invest and b_2 does not strictly benefit from making acceptable offers in a period before anyone has matched. Specifically, when selected to propose, b_2 weakly prefers to delay than to make either seller an acceptable offer; that is, letting w_a denote a 's expected equilibrium payoff at the beginning of this period, $\delta w_{b_2} \geq 2 - \delta w_{s_1}$ and $\delta w_{b_2} \geq 2 - \delta w_{s_2}$. Note also that $\delta w_{b_1} + \delta w_{s_1} < 1$ or $\delta w_{b_1} + \delta w_{s_2} < 1$, since otherwise the sum of everyone's payoffs in this period would be strictly larger than the maximum total surplus 3 that they can jointly generate. In either case, we obtain $\delta w_{b_2} > 1 + \delta w_{b_1} \geq 1$, which contradicts our assumption above that b_2 delays. Indeed, for b_2 to be willing to delay, her expected equilibrium payoff must increase when the state of the market changes—i.e., once b_1 matches. In particular, for b_2 to be willing to delay we must have $w_{b_2} \leq \delta w$, where w is the expected payoff of b_2 and the remaining seller at the beginning of each period once b_1 has left, and satisfies $w = \frac{1}{4}(2 - \delta w) + \frac{3}{4}\delta w$, and hence $\delta w < 1$.

We now extend the model presented in section 2.1 to include $n \geq 2$ buyers and n sellers. The probability that any given agent is selected to be the proposer is $\frac{1}{2n}$. As in the case of $n = 2$ considered above, once b_1 matches, there is an essentially unique Markov-perfect equilibrium: At the beginning of each period, each agent yet to match has the same expected payoff w , each agent accepts an offer if and only if it gives her at least δw , everyone's proposals offer δw to an agent on the other side of the market, and w satisfies

$$w = \frac{1}{2n}(2 - \delta w) + \frac{2n - 1}{2n}\delta w.$$

We now describe strategies that constitute a Markov-perfect equilibrium (for all sufficiently high discount factors δ) of every subgame in which b_1 has not invested and is yet to match. In any subgame in which there are $\ell = 1, 2, \dots, n$ buyers yet to match, including b_1 :

- (i) [Response strategies] Buyer b_1 accepts an offer if and only if it gives her at least $\delta w_{b_1}^\ell$, each buyer $b \neq b_1$ yet to match accepts an offer if and only if it gives her at least δw_b^ℓ , and each seller yet to match accepts an offer if and only if it gives her at least δw_s^ℓ .
- (ii) [Proposing strategies] For $\ell = 1$, buyer b_1 offers δw_s^ℓ to the seller yet to match. For $\ell = 2$, buyer b_1 offers δw_s^ℓ to a seller yet to match (chosen uniformly at random) with (small)

probability π , and delays (i.e., makes an unacceptable offer) otherwise. For $\ell > 2$, buyer b_1 delays with probability 1.

For $\ell \geq 1$, each buyer $b \neq b_1$ offers δw_s^ℓ to a seller yet to match (chosen uniformly at random), and each seller offers δw_b^ℓ to a buyer $b \neq b_1$ yet to match (chosen uniformly at random).

For $\ell = 1$:

1. We have that $w_{b_1}^1 = \frac{1}{2n}(1 - \delta w_{b_1}^1) + \frac{2n-1}{2n}\delta w_{b_1}^1$ and $w_{b_1}^1 = w_s^1$.

For $\ell = 2$:

2. Since b_1 is indifferent between making acceptable offers and moving to the next period, we have that $1 - \delta w_s^2 = \delta w_{b_1}^2$.
3. The expected equilibrium payoff $w_{b_1}^2$ of buyer b_1 is

$$\frac{1}{2n}\delta w_{b_1}^2 + \frac{3}{2n}\delta w_{b_1}^1 + \frac{2(n-2)}{2n}\delta w_{b_1}^2,$$

where the three terms of this expression correspond to the following:

1. The fact that b_1 delays implies that her expected equilibrium payoff when she is the proposer (which occurs with probability $\frac{1}{2n}$) is $\delta w_{b_1}^2$.
2. The expected equilibrium payoff of b_1 when a buyer other than b_1 or a seller that are yet to match are the proposers (which occurs with probability $\frac{3}{2n}$), a match that does not involve b_1 occurs with probability 1, so b_1 's expected equilibrium payoff is $\delta w_{b_1}^1$.
3. With the remaining probability $\frac{2(n-2)}{2n}$, an agent that has already matched is the proposer, in which case no match occurs this period and b_1 's expected equilibrium payoff is $\delta w_{b_1}^2$.
4. The expected equilibrium payoff w_s^2 of a seller s that is yet to match is

$$\frac{1}{2n}(2 - \delta w_b^2) + \frac{1}{2n}\delta w_s^1 + \frac{1}{2n}\left(\frac{1}{2}\delta w_s^2 + \frac{1}{2}\delta w_s^1\right) + \frac{1}{2n}\left(\pi\left(\frac{1}{2}\delta w_s^2 + \frac{1}{2}\delta w\right) + (1 - \pi)\delta w_s^2\right) + \frac{2(n-2)}{2n}\delta w_s^2,$$

where the five terms of this expression correspond to the following:

1. The expected equilibrium payoff of seller s when she is the proposer (which occurs with probability $\frac{1}{2n}$) is $2 - \delta w_b^2$.
2. The expected equilibrium payoff of seller s when another seller that is yet to match is the proposer (which occurs with probability $\frac{1}{2n}$) is δw_s^1 .

3. The expected equilibrium payoff of seller s when a buyer other than b_1 that is yet to match is the proposer (which occurs with probability $\frac{1}{2n}$) is $\frac{1}{2}\delta w_s^2 + \frac{1}{2}\delta w_s^1$ because s receives an acceptable offer with probability $\frac{1}{2}$ and the proposer matches with a different seller with the remaining probability $\frac{1}{2}$.
 4. The expected equilibrium payoff of seller s when b_1 is the proposer (which occurs with probability $\frac{1}{2n}$) is $\pi\left(\frac{1}{2}\delta w_s^2 + \frac{1}{2}\delta w\right) + (1-\pi)\delta w_s^2$, because b_1 makes an acceptable offer with probability π (in which case she makes seller s an acceptable offer with probability $\frac{1}{2}$ and she matches with another seller with the remaining probability $\frac{1}{2}$) and delays with the remaining probability $1 - \pi$.
 5. The expected equilibrium payoff of seller s when an agent that has already matched is the proposer (which occurs with probability $\frac{2(n-2)}{2n}$) is δw_s^2 .
5. The expected equilibrium payoff w_b^2 at the start of a period of a buyer $b \neq b_1$ that is yet to match is $\frac{1}{2n}(2 - \delta w_s^2) + \frac{1}{n}\delta w_b^2 + \frac{1}{2n}(\pi\delta w + (1 - \pi)\delta w_b^2) + \frac{n-2}{n}\delta w_b^2$, where the four terms of this expression correspond to the following:
1. The expected equilibrium payoff of buyer b when she is the proposer (which occurs with probability $\frac{1}{2n}$) is $2 - \delta w_s^2$.
 2. The expected equilibrium payoff of buyer b when a seller that is yet to match is the proposer (which occurs with probability $\frac{1}{n}$) is δw_b^2 because b receives an acceptable offer with probability 1.
 3. The expected equilibrium payoff of buyer b when b_1 is the proposer (which occurs with probability $\frac{1}{2n}$) is $\pi\delta w + (1 - \pi)\delta w_b^2$, because b_1 makes an acceptable offer with probability π (in which case b obtains δw) and delays with the remaining probability $1 - \pi$.
 4. The expected equilibrium payoff of buyer b when an agent that has already matched is the proposer (which occurs with probability $\frac{n-2}{n}$) is δw_b^2 .

Finally, for every $\ell = 3, \dots, n$:

6. The expected equilibrium payoff $w_{b_1}^\ell$ of buyer b_1 is

$$\frac{1}{2n}\delta w_{b_1}^\ell + \frac{2\ell-1}{2n}\delta w_{b_1}^{\ell-1} + \frac{2n-2\ell}{2n}\delta w_{b_1}^\ell,$$

where the three terms of this expression correspond to the following:

1. The fact that b_1 delays implies that her expected equilibrium payoff when she is the proposer (which occurs with probability $\frac{1}{2n}$) is $\delta w_{b_1}^\ell$.

2. The expected equilibrium payoff of b_1 when a buyer other than b_1 or a seller that are yet to match are the proposers (which occurs with probability $\frac{2\ell-1}{2n}$), a match that does not involve b_1 occurs with probability 1, so b_1 's expected equilibrium payoff is $\delta w_{b_1}^{\ell-1}$.
 3. With the remaining probability $\frac{2n-2\ell}{2n}$, an agent that has already matched is the proposer, in which case no match occurs this period and b_1 's expected equilibrium payoff is $\delta w_{b_1}^\ell$.
7. The expected equilibrium payoff w_s^ℓ of a seller s that is yet to match is $\frac{1}{2n}(2 - \delta w_b^\ell) + \frac{\ell-1}{2n}\delta w_s^{\ell-1} + \frac{\ell-1}{2n}\left(\frac{1}{\ell}\delta w_s^\ell + \frac{\ell-1}{\ell}\delta w_s^{\ell-1}\right) + \frac{1}{2n}\delta w_s^\ell + \frac{2(n-\ell)}{2n}\delta w_s^\ell$, where the five terms of this expression correspond to the following:
1. The expected equilibrium payoff of seller s when she is the proposer (which occurs with probability $\frac{1}{2n}$) is $2 - \delta w_b^\ell$.
 2. The expected equilibrium payoff of seller s when another seller that is yet to match is the proposer (which occurs with probability $\frac{\ell-1}{2n}$) is $\delta w_s^{\ell-1}$.
 3. The expected equilibrium payoff of seller s when a buyer other than b_1 that is yet to match is the proposer (which occurs with probability $\frac{\ell-1}{2n}$) is $\frac{1}{\ell}\delta w_s^\ell + \frac{\ell-1}{\ell}\delta w_s^{\ell-1}$ because s receives an acceptable offer with probability $\frac{1}{\ell}$ and the proposer matches with a different seller with the remaining probability $\frac{\ell-1}{\ell}$.
 4. The expected equilibrium payoff of seller s when b_1 is the proposer (which occurs with probability $\frac{1}{2n}$) is δw_s^ℓ .
 5. The expected equilibrium payoff of seller s when an agent that has already matched is the proposer (which occurs with probability $\frac{2(n-\ell)}{2n}$) is δw_s^ℓ .
8. The expected equilibrium payoff w_b^ℓ at the start of a period of a buyer $b \neq b_1$ that is yet to match is $\frac{1}{2n}(2 - \delta w_s^\ell) + \frac{\ell-2}{2n}\delta w_b^{\ell-1} + \frac{\ell}{2n}\left(\frac{1}{\ell-1}\delta w_b^\ell + \frac{\ell-2}{\ell-1}\delta w_b^{\ell-1}\right) + \frac{1}{2n}\delta w_b^\ell + \frac{2(n-\ell)}{2n}\delta w_b^\ell$, where the five terms of this expression correspond to the following:
1. The expected equilibrium payoff of buyer b when she is the proposer (which occurs with probability $\frac{1}{2n}$) is $2 - \delta w_s^\ell$.
 2. The expected equilibrium payoff of buyer b when a buyer other than b or b_1 that is yet to match is the proposer (which occurs with probability $\frac{\ell-2}{2n}$) is $\delta w_b^{\ell-1}$.
 3. The expected equilibrium payoff of buyer b when a seller that is yet to match is the proposer (which occurs with probability $\frac{\ell}{2n}$) is $\frac{1}{\ell-1}\delta w_b^\ell + \frac{\ell-2}{\ell-1}\delta w_b^{\ell-1}$ because b

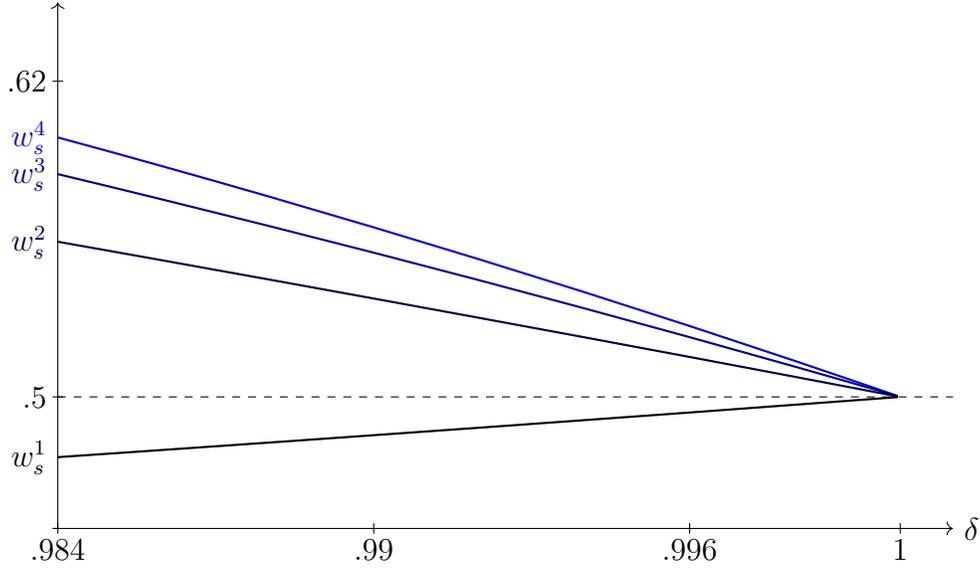


Figure 4: Sellers' equilibrium payoff w_s^ℓ as a function of δ in the case $n = 4$.

receives an acceptable offer with probability $\frac{1}{\ell-1}$ and the proposer matches with another buyer different from b_1 with the remaining probability $\frac{\ell-2}{\ell-1}$.

4. The expected equilibrium payoff of buyer b when b_1 is the proposer (which occurs with probability $\frac{1}{2n}$) is δw_b^ℓ .
5. The expected equilibrium payoff of buyer b when an agent that has already matched is the proposer (which occurs with probability $\frac{2(n-\ell)}{2n}$) is δw_b^ℓ .

It can be checked that the system of equations defined by points (1)-(8) have a unique valid solution $(w_{b_1}^1, w_s^1, \pi) \cup (w_{b_1}^\ell, w_b^\ell, w_s^\ell)_{\ell=2,3,\dots,n}$ for all δ sufficiently large. Figure 4 illustrates how w_s^ℓ converges to $1/2$ as δ goes to 1.

B Construction of equilibrium in subsection 2.2

For any state s and any agent a , let w_a^s denote a 's expected equilibrium payoff at the beginning of a period in which the state is s (before anyone enters in this period).

1. Recall that we defined w above as every active agent's expected payoff after b_1 has matched, $w = \frac{1}{4}(2 - \delta w) + \frac{3}{4}\delta w$.
2. Buyer b_1 's expected equilibrium payoff in a period in which, after this period's entry is determined, the state is 1 is $\frac{1}{4}(1 - \delta w_s^1) + \frac{3}{4}\delta w_{b_1}^1$. In state 2, b_1 is indifferent between

making an acceptable offer and waiting for the next period, so

$$\delta w_{b_1}^2 = 1 - \delta w_s^2$$

and

$$w_{b_1}^1 = \rho \delta w_{b_1}^2 + (1 - \rho) \left(\frac{1}{4}(1 - \delta w_s^1) + \frac{3}{4} \delta w_{b_1}^1 \right).$$

3. In state 2, buyer b_2 and a seller match unless b_1 is the proposer, so

$$w_{b_1}^2 = \frac{1}{4} \delta w_{b_1}^2 + \frac{3}{4} \delta w_{b_1}^1.$$

4. The expected equilibrium payoff of seller s in state 2 is $\frac{1}{4}(2 - \delta w_{b_2}^2) + \frac{1}{4} \left(\frac{\pi}{2} \delta w + \left[1 - \frac{\pi}{2} \right] \delta w_s^2 \right) + \frac{1}{4} \left(\frac{1}{2} \delta w_s^2 + \frac{1}{2} \delta w_s^1 \right) + \frac{1}{4} \delta w_s^1$. This is because

- The expected equilibrium payoff of seller s when she is the proposer is $2 - \delta w_{b_2}^2$.
- The expected equilibrium payoff of seller s when buyer b_1 is the proposer is $\frac{\pi}{2} \delta w + \left[1 - \frac{\pi}{2} \right] \delta w_s^2$, because in this case b_1 matches with the other seller with probability $\pi/2$ (in which case the expected payoff of s is δw) and with the remaining probability, s either receives an acceptable offer or stays in the market; in either case, her expected equilibrium payoff is δw_s^2 .
- The expected equilibrium payoff of seller s when buyer b_2 is the proposer is $\frac{1}{2} \delta w_s^2 + \frac{1}{2} \delta w_s^1$, because in this case s receives an acceptable offer with probability $1/2$ (in which case her expected payoff is δw_s^2) and with the remaining probability, b_2 matches with the other seller (in which case her expected payoff is δw_s^1).
- The expected equilibrium payoff of seller s when the other seller is the proposer is δw_s^1 , because in this case the other seller matches with b_2 with probability 1.

Hence,

$$w_s^2 = \frac{1}{4}(2 - \delta w_{b_2}^2) + \frac{1}{4} \left(\frac{\pi}{2} \delta w + \left[1 - \frac{\pi}{2} \right] \delta w_s^2 \right) + \frac{1}{4} \left(\frac{1}{2} \delta w_s^2 + \frac{1}{2} \delta w_s^1 \right) + \frac{1}{4} \delta w_s^1,$$

and

$$w_s^1 = \rho w_s^2 + (1 - \rho) \left(\frac{1}{4}(1 - \delta w_{b_1}^1) + \frac{3}{4} \delta w_s^1 \right).$$

5. Buyer b_2 's expected equilibrium payoff $w_{b_2}^2$ in state 2 is

$$w_{b_2}^2 = \frac{1}{4}(2 - \delta w_s^2) + \frac{1}{4} \left(\pi \delta w + (1 - \pi) \delta w_{b_2}^2 \right) + \frac{1}{2} \delta w_{b_2}^2,$$

where the three terms of this expression correspond to the following:

- The expected equilibrium payoff of buyer b_2 when she is the proposer is $2 - \delta w_s^2$.
- The expected equilibrium payoff of buyer b_2 when buyer b_1 is the proposer is $\pi \delta w + (1 - \pi) \delta w_{b_2}^2$, because in this case b_1 matches with probability π (in which case b_2 's expected equilibrium payoff is δw) and no one matches (so the market stay in state 2) with the remaining probability.
- The expected equilibrium payoff of buyer b_2 when a seller is the proposer is $\delta w_{b_2}^2$, because every seller makes her acceptable offers.

It can be checked that the seven equations defined by points (1)-(5) have a unique valid solution $(w, w_{b_1}^1, w_s^1, w_{b_1}^2, w_{b_2}^2, w_s^2, \pi)$ for all δ sufficiently close to 1. Figure 3 in section 2 illustrates the limit of w_s^2 as δ goes to 1 for all $0 \leq \rho \leq 1$.

C Existence of a type-symmetric Markov-perfect equilibrium

We characterize the Markov perfect equilibrium of the subgame that starts at $t = 1$ with agent d active, and then use it to show that such an equilibrium exists. The argument is similar to the one in the equilibrium existence proof in Elliott and Nava (2019). We consider the general case in which each productive match contains at most one agent of each type: Each match that contains one agent of each type in $C \subseteq N$ generates $y(C)$ units of surplus.

Consider a Markov-perfect-equilibrium and its corresponding value function $V : \hat{\mathcal{A}} \rightarrow \mathbb{R}^m$, where m denotes the cardinality of $\hat{\mathcal{A}}$ and $V(A)$ gives each agent's expected equilibrium payoff at the beginning of a period that starts with active agent set A (before any agents enter in this period). Consider a subgame with an arbitrary active agent set A . By Markov perfection, agent b accepts every offer that gives her strictly more than $\delta V_b(A)$, and rejects every offer that gives her strictly less than $\delta V_b(A)$. Hence, no one offers more than $\delta V_b(A)$ to any agent b , and an agent of type j proposes to form $C \subseteq N - j$ that maximizes $y(C + j) - \sum_{k \in C} \min_{b_k \in \mathcal{L}_k} \delta V_k(A)$. When agent a of type j is the proposer, if there exists a coalition $C \subseteq N - j$ such that $y(C + j) - \sum_{k \in C} \min_{b_k \in \mathcal{L}_k} \delta V_{b_k}(A) > \delta V_a(A)$, then she makes offers only to coalitions that maximize this net surplus, and agreement obtains with probability one. Otherwise, she delays—in the sense that she makes offers that are not accepted in equilibrium. Letting $\pi^{a,A}(B)$ denote the probability that the set B of agents match when the set of active agents is A and agent a is the proposer, any agreement probability distribution $\pi^{a,A}$ that is consistent with the value function V is in the set $\Pi^{a,A}(V)$ of such distributions that satisfy that, for any set of agents B with types in $j \in C \subseteq N$, $\pi^{a,A}(B) = 0$ if $\delta V_a(A) > y(C) - \sum_{b \in B-a} \delta V_b(A)$ or $y(C) - \sum_{b \in B-a} \delta V_b(A) < \max_{D \subseteq N-j} (y(D + j) - \sum_{k \in D} \min_{b_k \in \mathcal{L}_k} \delta V_{b_k}(A))$.

For any value function V , any set of active agents A in $\hat{\mathcal{A}}$, and any agent a in A of type j , define the correspondence $f^{a,A}(V) : A \rightarrow \mathbb{R}^m$ by

$$\begin{aligned} f_a^{a,A}(V) &= \pi^{a,A}(\emptyset)\delta V_a(A) + [1 - \pi^{a,A}(\emptyset)] \max_{C \subseteq N-j} (y(C+j) - \sum_{k \in C} \min_{b_k \in \mathcal{L}_k} \delta V_{b_k}(A)) \\ f_b^{a,A}(V) &= \pi^{a,A}(\emptyset)\delta V_b(A) + \sum_{B \in \mathcal{B}_{\subseteq A}} \pi^{a,A}(B)\delta V_b(A) + \sum_{C \subseteq A-b} \pi_A(C)\delta V_b(A-C) \quad \forall b \neq a, \end{aligned}$$

for any $\pi^{a,A}$ in $\Pi^{a,A}(V)$. In other words, $f^{a,A}(V)$ gives the set of expected payoffs that are consistent with the value function V in any history in which the active agent set is A and the proposer is agent a . Letting \mathcal{V} denote the set of value functions $V : \hat{\mathcal{A}} \rightarrow \mathbb{R}^m$, consider the correspondence $F : \mathcal{V} \rightarrow \mathcal{V}$ defined by

$$(8) \quad F(V)(A) = \frac{1}{n} \sum_{a \in \hat{\mathcal{A}}} f^{a,A}(V).$$

The value function V corresponds to a Markov-perfect equilibrium payoff profile if and only if V is in $F(V)$, so it is enough to show that the correspondence F has a fixed point. This follows from Kakutani's (1941) fixed point theorem. Indeed, the domain \mathcal{V} of F is a non-empty, compact and convex subset of an Euclidean space. Moreover, since, for any A in $\hat{\mathcal{A}}$ and any a in A , the correspondence $\Pi^{a,A}$ is upper-hemicontinuous with non-empty convex images, so is the correspondence $f^{a,A}$, and hence so is F .

D Example: Agents need not be price takers when an investment deviation can trigger further investment deviations

As in section 6.3, all agents make investments upon arriving to the market. In contrast to section 6.3, agents can observe all the previous investment choices made by other agents before choosing their own investment. We construct an equilibrium in which agents are not price takers. In particular, when a buyer makes an investment deviation, it triggers further investment deviations by other buyers that enter the market after her, and this changes the sellers' limit payoffs.

There are two types of agents, buyers and sellers, and two slots per type. The process of dynamic entry is such that there are always two active buyers and two active sellers (as in the case $\rho = 1$ of the example of section 2.2). Each buyer can choose to invest or to not invest. Investing costs .55. A buyer that does not invest has a matching surplus of 1 unit with any seller, whereas a buyer that invests has a matching surplus of 2 units with any seller.

We construct an equilibrium in which no buyer invests, and the unit surplus of each match is shared equally as $\delta \rightarrow 1$. Note that these investments are not constrained efficient.

In particular, if buyers were price takers, they would find it optimal to invest, because each would appropriate the unit value generated by her investment, which is larger than the cost of this investment.

There are three possible Markov states: The state 0 in which none of the two active buyers has invested, the state 1 in which only one of the two active buyers has invested, and the state 2 in which both buyers have invested. We describe an equilibrium in which:

1. In state 0, there is no delay: The proposer makes an acceptable offer to an agent on the other side of the market, and this offer is accepted with probability 1. Moreover, state 0 is absorbing; that is, the buyer that enters after a match occurs in state 0 never invests.
2. In state 1:
 - (i) Only the buyer that has not invested (b_2 , say) delays, and she does so with probability 1. The sellers make acceptable offers to the buyer that has invested (b_1 say).
 - (ii) If b_1 leaves, the market stays in state 1 (that is, the buyer that replaces b_1 chooses to invest) with probability π , and it moves to state 0 (that is, the buyer that replaces b_1 chooses to not invest) with probability $1 - \pi$.
 - (iii) If b_2 leaves, the market moves to state 2 (that is, the buyer that replaces b_2 chooses to invest) with probability 1.
3. In state 2 there is no delay: The proposer makes an acceptable offer to an agent on the other side of the market, and this offer is accepted with probability 1. Moreover, state 2 is absorbing; that is, the buyer that enters after a match in state 2 always invests.

Everyone's equilibrium payoff w^0 in state 0 satisfies:

$$w^0 = \frac{1}{4}(1 - \delta w^0) + \frac{3}{4}\delta w^0.$$

Similarly, everyone's equilibrium payoff w^2 in state 2 satisfies:

$$w^2 = \frac{1}{4}(2 - \delta w^2) + \frac{3}{4}\delta w^2.$$

In state 1, buyer b_2 delays with probability 1, which implies that

$$w_{b_2}^1 = \frac{1}{4}\delta w_{b_2}^1 + \frac{3}{4}(\pi\delta w_{b_2}^1 + (1 - \pi)\delta w^0) \quad \text{and} \quad w_{b_1}^1 = \frac{1}{4}(2 - \delta w_s^1) + \frac{3}{4}\delta w_{b_1}^1.$$

Moreover, the fact that the buyer that replaces b_1 mixes between investing and not investing implies that she is indifferent between these two choices; that is,

$$\delta w_{b_1}^1 - .55 = \delta w^0.$$

Finally, the equilibrium payoff of any seller s in state 1 satisfies:

$$w_s^1 = \frac{1}{4}(2 - \delta w_{b_1}^1) + \frac{1}{4} \left(\frac{1}{2} \delta w_s^1 + \frac{1}{2} (\pi \delta w_s^1 + (1 - \pi) \delta w^0) \right) + \frac{1}{4} \delta w_s^1 + \frac{1}{4} (\pi \delta w_s^1 + (1 - \pi) \delta w^0)$$

where the first term corresponds to the event that s is the proposer, the second term corresponds to the event that b_1 is the proposer, the third term corresponds to the event that b_2 is the proposer, and the fourth term corresponds to the event that the other seller is the proposer.

It can be verified that the solution to the six equations above satisfies $w^0 \rightarrow 1/2$, $w^2 \rightarrow 1$, $w_{b_1}^1 \rightarrow 1.1$, $w_{b_2}^1 \rightarrow .15$, $w_s^1 \rightarrow .9$ and $\pi \rightarrow 1$ as $\delta \rightarrow 1$. Hence, the strategies described constitute an equilibrium for all sufficiently high discount factors δ . In particular, $\delta w^0 = \delta w_{b_1}^1 - c$ and $\delta w^2 - c > \delta w_{b_2}^1$, so the buyers' investment choices are optimal. Moreover, $\delta w_{b_2}^1 + \delta w_s^1 > 1$, so it is optimal for b_2 to delay in state 1. We conclude that no one ever investing (that is, staying in state 0 forever) is part of an equilibrium for all δ sufficiently large.

References

- ACEMOGLU, D. (1997): "Training and innovation in an imperfect labour market," *The Review of Economic Studies*, 64, 445–464.
- ACEMOGLU, D. AND R. SHIMER (1999): "Holdups and efficiency with search frictions," *International Economic Review*, 40, 827–849.
- AKBARPOUR, M., S. D. KOMINERS, S. LI, AND P. R. MILGROM (2020): "Investment Incentives in Near-Optimal Mechanisms," *Available at SSRN 3544100*.
- BECKER, G. S. (1962): "Investment in human capital: A theoretical analysis," *Journal of political economy*, 70, 9–49.
- BESTER, H. (2013): "Investments and the holdup problem in a matching market," *Journal of Mathematical Economics*, 49, 302–311.
- BINMORE, K. G. AND M. J. HERRERO (1988): "Security equilibrium," *The Review of Economic Studies*, 55, 33–48.
- CHE, Y.-K. AND J. SÁKOVICS (2004): "A dynamic theory of holdup," *Econometrica*, 72, 1063–1103.
- CHIAPPORI, P.-A., M. C. DIAS, AND C. MEGHIR (2018): "The Marriage Market, Labor Supply, and Education Choice," *Journal of Political Economy*, 126, S26–S72.
- CHIAPPORI, P.-A., B. SALANIÉ, AND Y. WEISS (2017): "Partner choice, investment in children, and the marital college premium," *The American Economic Review*, 107, 2109–67.
- COLE, H., G. MAILATH, AND A. POSTLEWAITE (2001a): "Efficient non-contractible investments in finite economies," *The BE Journal of Theoretical Economics*, 1, 1–34.
- (2001b): "Efficient non-contractible investments in large economies," *Journal of Economic Theory*, 101, 333–373.
- DÁVILA, J. AND J. EECKHOUT (2008): "Competitive bargaining equilibrium," *Journal of Economic Theory*, 139, 269–294.
- DE FRAJA, G. AND J. SÁKOVICS (2001): "Walras retrouvé: Decentralized trading mechanisms and the competitive price," *Journal of Political Economy*, 109, 842–863.

- DECKER, R. A., J. HALTIWANGER, R. S. JARMIN, AND J. MIRANDA (2020): "Changing Business Dynamism and Productivity: Shocks versus Responsiveness," *The American Economic Review*, 110, 3952–90.
- DIZDAR, D. (2018): "Two-sided investment and matching with multidimensional cost types and attributes," *American Economic Journal: Microeconomics*, 10, 86–123.
- ELLIOTT, M. (2015): "Inefficiencies in networked markets," *American Economic Journal: Microeconomics*, 7, 43–82.
- ELLIOTT, M. AND F. NAVA (2019): "Decentralized bargaining in matching markets: Efficient stationary equilibria and the core," *Theoretical Economics*, 14, 211–251.
- FUDENBERG, D. AND J. TIROLE (1991): *Game theory*, MIT Press.
- GALE, D. (1987): "Limit theorems for markets with sequential bargaining," *Journal of Economic Theory*, 43, 20–54.
- GALE, D. AND H. SABOURIAN (2005): "Complexity and competition," *Econometrica*, 73, 739–769.
- GRETSKY, N. E., J. M. OSTROY, AND W. R. ZAME (1999): "Perfect competition in the continuous assignment model," *Journal of Economic Theory*, 88, 60–118.
- GROSSMAN, S. J. AND O. D. HART (1986): "The costs and benefits of ownership: A theory of vertical and lateral integration," *Journal of Political Economy*, 94, 691–719.
- GROUT, P. A. (1984): "Investment and wages in the absence of binding contracts: A Nash bargaining approach," *Econometrica*, 449–460.
- GUL, F. (2001): "Unobservable Investment and the Hold-Up Problem," *Econometrica*, 69, 343–376.
- HART, O. AND J. MOORE (1990): "Property Rights and the Nature of the Firm," *Journal of Political Economy*, 98, 1119–1158.
- HATFIELD, J. W., F. KOJIMA, AND S. D. KOMINERS (2019): "Strategy-proofness, investment efficiency, and marginal returns: An equivalence," *Becker Friedman Institute for Research in Economics Working Paper*.
- HOSIOS, A. J. (1990): "On the efficiency of matching and related models of search and unemployment," *The Review of Economic Studies*, 57, 279–298.

- KAKUTANI, S. (1941): "A generalization of Brouwer's fixed point theorem," *Duke Mathematical Journal*, 8, 457–459.
- LAUERMANN, S. (2013): "Dynamic matching and bargaining games: A general approach," *The American Economic Review*, 103, 663–689.
- LAUERMANN, S., G. NÖLDEKE, AND T. TRÖGER (2020): "The Balance Condition in Search-and-Matching Models," *Econometrica*, 88, 595–618.
- LEONARD, H. B. (1983): "Elicitation of honest preferences for the assignment of individuals to positions," *Journal of Political Economy*, 91, 461–479.
- MAILATH, G. J., A. POSTLEWAITE, AND L. SAMUELSON (2017): "Premuneration values and investments in matching markets," *The Economic Journal*, 127, 2041–2065.
- MAKOWSKI, L. AND J. M. OSTROY (1995): "Appropriation and efficiency: A revision of the first theorem of welfare economics," *The American Economic Review*, 808–827.
- MANEA, M. (2011): "Bargaining in stationary networks," *The American Economic Review*, 101, 2042–2080.
- MCLENNAN, A. AND H. SONNENSCHNEIN (1991): "Sequential bargaining as a noncooperative foundation for Walrasian equilibrium," *Econometrica*, 59, 1395–1424.
- NGUYEN, T. (2015): "Coalitional bargaining in networks," *Operations Research*, 63, 501–511.
- NÖLDEKE, G. AND L. SAMUELSON (2015): "Investment and competitive matching," *Econometrica*, 83, 835–896.
- POLANSKI, A. AND F. VEGA-REDONDO (2018): "Bargaining frictions in trading networks," *The BE Journal of Theoretical Economics*, 18.
- ROGERSON, W. P. (1992): "Contractual solutions to the hold-up problem," *The Review of Economic Studies*, 59, 777–793.
- RUBINSTEIN, A. (1982): "Perfect equilibrium in a bargaining model," *Econometrica*, 97–109.
- RUBINSTEIN, A. AND A. WOLINSKY (1985): "Equilibrium in a market with sequential bargaining," *Econometrica*, 1133–1150.
- (1990): "Decentralized trading, strategic behaviour and the Walrasian outcome," *The Review of Economic Studies*, 57, 63–78.

TALAMÀS, E. (2020): "Nash bargaining with endogenous outside options," *Mimeo*.

WILLIAMSON, O. E. (1975): *The Economic Institutions of Capitalism*, New York Free Press.

WOLINSKY, A. (1988): "Dynamic markets with competitive bidding," *The Review of Economic Studies*, 55, 71–84.

ZHANG, H. (2020): "An Investment-and-Marriage Model with Differential Fecundity," *Journal of Political Economy*, *accepted*.