

Cambridge Working Papers in Economics

Cambridge-INET Working Paper Series No: 2020/32

Cambridge Working Papers in Economics: 2070

NO HOLDUP IN DYNAMIC MARKETS

Matthew Elliott
(University of Cambridge)

Eduard Talamà
(IESE Business School)

In many markets, heterogeneous agents make non-contractible investments before bargaining over both who matches with whom and the terms of trade. In static markets, the holdup problem—that is, inefficient investments caused by agents receiving only a fraction of their returns—is ubiquitous. Markets are often dynamic, however, with agents entering over time. Taking a general non-cooperative investment and bargaining approach, we show that the holdup problem vanishes in markets with dynamic entry as agents become patient: While there is substantial wiggle room for bargaining to determine outcomes, every bargaining outcome gives everyone her marginal product.

No holdup in dynamic markets*

Matthew Elliott & Eduard Talamàs[†]

Abstract

In many markets, heterogenous agents make non-contractible investments before bargaining over both who matches with whom and the terms of trade. In static markets, the *holdup problem*—that is, inefficient investments caused by agents receiving only a fraction of their returns—is ubiquitous. Markets are often dynamic, however, with agents entering over time. Taking a general non-cooperative investment and bargaining approach, we show that the holdup problem vanishes in markets with dynamic entry as agents become patient: While there is substantial wiggle room for bargaining to determine outcomes, every bargaining outcome gives everyone her marginal product.

1 Introduction

Crucial investments are often sunk by the time agents bargain over prices and allocations. For example, workers and employers invest in human and physical capital well before bargaining over who will match with whom and for what wages. This can lead to holdup problems—that is, agents underinvesting because they do not expect to fully appropriate the returns from their investments (e.g., Williamson 1975; Grout 1984; Grossman and Hart 1986; Tirole 1986; Hart and Moore 1990) and severely limit the efficiency of these markets (e.g., Hosios 1990; Acemoglu 1996, 1997; Cole, Mailath, and Postlewaite 2001a; de Meza and Lockwood 2010; Elliott 2015; Felli and Roberts 2016).

**Date Printed:* May 20, 2020.

[†]Elliott: University of Cambridge. Talamàs: IESE Business School. We thank Benjamin Golub, George Mailath, Georg Nöldeke, Andrew Postlewaite and Larry Samuelson for very useful conversations, and various audiences for helpful feedback. Financial support from the European Research Council under the grant “EMBED” #757229, the INET institute at Cambridge (Elliott), and the Foundations of Human Behavior Initiative, the Warren Center for Network & Data Sciences and the Rockefeller Foundation #2017PRE301 (Talamàs) is acknowledged. All errors are ours.

The objective of this paper is to investigate the extent to which holdup is a problem in matching markets featuring dynamic entry. While the holdup problem in static markets has been extensively studied, relatively little attention has been given to markets with dynamic entry. In particular, to the best of our knowledge, this is the first paper that studies the holdup problem in thin markets with dynamic entry.

We model matching markets using a general non-cooperative investment and bargaining game with stochastic inflows and outflows of agents: Before entering the market, different types of agents have access to different non-contractible investments. The investments they make are sunk by the time they enter the market, and they shape their matching surpluses.

Agents in the market bargain according to a standard protocol in the spirit of Rubinstein (1982): In each period, one agent is randomly selected to be the proposer. The proposer chooses whom to make an offer to as well as how to split the resulting surplus. The agent receiving the offer then decides whether to accept it, in which case she matches with the proposer, and both leave the market; or reject it, in which case no match occurs in this period.

We characterize the type-symmetric Markov-perfect equilibria of this game for all sufficiently high discount factors, and we show that—in the limit as agents become patient—the payoffs of a given agent’s potential trade partners are unaffected by that agent’s presence. In other words, each agent obtains her marginal product independently of her investments, so there is no holdup problem.

This result is in stark contrast with the fact that, in static matching markets, it is usually impossible for everyone to obtain her marginal product. For example, consider a market with n buyers and n sellers. Each buyer can only match with one seller, and vice versa. Assume for simplicity that everyone matches. When the economy is static, each agent’s marginal product is the full surplus of one match, so the sum of all agents’ marginal products is twice the total surplus. In other words, it is far from possible to give everyone their marginal product—independently of the size n of the market. Indeed, in the context of the assignment game, Leonard (1983) shows that an agent obtains her marginal product if and only if she receives her highest possible payoff among all Walrasian equilibria. Hence, full appropriation only occurs when there is a unique Walrasian equilibrium that pins down all prices—a situation that, as shown by Gretsky, Ostroy, and Zame (1999), is not generic in finite markets.

A tempting but incorrect intuition for our result is that the combination of dynamic entry and patience makes the market effectively thick, pinning down the price that must obtain in a Walrasian equilibrium thereby allowing all agents to simultaneously receive their marginal

product. Indeed, as in the standard models of bargaining in stationary markets (e.g., Rubinstein and Wolinsky 1985), in our setting there is ample wiggle room for bargaining to play a meaningful role in determining outcomes—even in the limit as agents become impatient. Despite equilibrium prices varying with agents’ bargaining power (their proposer probabilities), we show that every bargaining outcome gives everyone her marginal product as agents become patient.

To build intuition for our result, consider an agent who invests differently from all her fellow agents of the same type. On the one hand, she is in a strong bargaining position against agents with whom she generates more surplus than her fellows: She can play these agents off to make sure that she appropriates the additional potential gains generated by her investment deviation. On the other hand, she is in a weak bargaining position against the agents with whom she generates less surplus than her fellows: These other agents can effectively ignore her with minimal payoff consequences, so if she ends up in such a match she has to appropriate these potential losses too. Hence, even while engaging in decentralized non-cooperative bargaining in a market that may appear thin at every point in time—and where bargaining dynamics shape the different types’ payoffs in non-trivial ways—every individual is a price taker, and hence a residual claimant of the surplus created or lost by any unilateral investment deviation.

Our results show that the sources of holdup in dynamic economies can be qualitatively different from those in static economies. This has important practical implications. For example, Davis and Haltiwanger (2014) document how US labor market fluidity—as measured by flows of jobs and workers across employers—has fallen over the last few decades, and they argue that this has significantly reduced productivity.¹ Our findings suggest an investment channel by which lower market fluidity slows matching rates thereby exacerbating holdup problems, which can lead to less investment and hence lower productivity.

Related literature

The literature investigating the efficiency of investments under competitive matching (e.g., Cole, Mailath, and Postlewaite 2001b; Peters and Siow 2002; Mailath, Postlewaite, and Samuelson 2013 and 2017; Nöldeke and Samuelson 2015; Chiappori, Salanié, and Weiss 2017; Chiappori, Dias, and Meghir 2018; Dizdar 2018) focuses on markets featuring a continuum of *price-taking agents* on each side to turn off the holdup problem and investigate other sources of investment inefficiencies—like coordination failures, participation constraints, and imper-

¹See also Molloy et al. (2016) and Decker et al. (2018), for example.

fect information. Our results provide non-cooperative foundations in finite markets for this widely used price-taking assumption, and shows that dynamic entry can be an important force behind price-taking behavior.

A branch of the search and matching literature also investigates investment incentives under competitive matching. For example, Acemoglu and Shimer (1999) show that holdup is not a problem in directed search environments where firms form submarkets by committing to a posted wage. Relatedly, Bester (2013) shows that holdup need not be a problem in the steady state equilibrium of a bilateral matching market with a continuum of identical agents on each side. In contrast, in this paper we take a non-cooperative bargaining approach to study (not-necessarily-two-sided) markets with arbitrarily many different types of agents, and we show that holdup need not be a problem even when there are very few agents in the market at every point in time (on and off the equilibrium path).

In the case of finite markets (with unidimensional attributes and complementarities in these attributes), Cole, Mailath, and Postlewaite (2001a) provide a condition called “doubly overlapping attributes” that guarantees that there is an essentially unique Walrasian equilibrium, and that the associated prices continue to clear the market after any unilateral investment deviation. Under these conditions, agents are price takers—in the sense that no unilateral change in attributes affects the market prices—and, as a result, efficient non-contractible investments can be supported in equilibrium. We take a dynamic approach to address similar questions, and we find that essentially no restrictions on the nature of the investments and resulting matching surpluses are required to preclude holdup problems when agents are sufficiently patient. This is because our result does not rely on Walrasian equilibrium pinning down outcomes uniquely.

Finally, it has been well-known at least since the work of Rogerson (1992) and Makowski and Ostroy (1995) that efficiency requires agents appropriating their marginal products. Naturally, full appropriation also plays a central role in mechanism design (e.g., Bergemann and Välimäki 2002; Hatfield et al. 2019). In this paper, we show that full appropriation is endogenously satisfied in our general non-cooperative bargaining game when agents are patient.

Roadmap

The rest of this paper is organized as follows. In section 2, we illustrate the main ideas in the context of a simple example. In section 3, we describe the general model and, in section 4, we present and prove our main result. We relegate some details of the analysis in section 2 to Appendix A, and relatively standard results that we use to prove our main

result to Appendix B and Appendix C.

2 Example

We start, in subsection 2.1, by reviewing the standard holdup problem in the context of an investment and bargaining game featuring two buyers and two sellers.² Then, in subsection 2.2, we describe a homologous market with sequential entry, and we illustrate how, in this case, there is no holdup problem. For simplicity, in the version of this example featuring sequential entry, we assume that each agent that leaves the market is immediately replaced by a replica.³

Later on, we show the analogous *no holdup* result in a more general setting that allows (i) stochastic entry (a relaxation of the replica assumption), (ii) arbitrarily many types of agents, and (iii) a rich investment technology that allows for general and type-specific investments.

2.1 Holdup in a market without sequential entry

Let us start by describing a simple non-cooperative game featuring a standard holdup problem. There are two identical buyers, b_1 and b_2 , and two identical sellers, s_1 and s_2 , with a common discount factor δ . In the first period $t = 0$, they simultaneously make investments. They can choose either to *invest* or to *not invest*. Their investments shape their matching surpluses: When a buyer and a seller match in any period $t = 1, 2, \dots$, they generate

$$(1) \quad \begin{array}{ll} 2 \text{ units of surplus} & \text{if } \textit{both} \text{ have invested,} \\ 1 \text{ unit of surplus} & \text{if } \textit{only one of them} \text{ has invested, and} \\ 0 \text{ units of surplus} & \text{if } \textit{none of them} \text{ has invested.} \end{array}$$

Not investing costs zero, and investing costs c , with $1/2 < c < 1$. Hence, efficiency requires that everyone invests if the discount factor δ is sufficiently close to 1.

We focus on the case in which investments at time 0 are not contractible. This requires specifying how the outcome (that is, who matches with whom and how the resulting surplus is shared) is determined as a function of the realized investments. We take a non-cooperative

²Holdup is generally also a problem in larger static markets and in unbalanced markets (e.g., Gretsky, Ostroy, and Zame 1999 and Cole, Mailath, and Postlewaite 2001a)

³Our general model relaxes this assumption, which has been widely used in the dynamic matching and bargaining literature; see for example Rubinstein and Wolinsky (1985), Manea (2011), Nguyen (2015), Polanski and Vega-Redondo (2018) and Talamàs (2019b).

approach: Once the agents have sunk their investments, they bargain 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 at random to be the proposer: Each buyer has a probability $\frac{\beta}{2}$ of being selected, and each seller has a probability $\frac{1-\beta}{2}$ of being selected, where $0 \leq \beta \leq 1$ reflects the relative bargaining power of buyers. If the selected agent has already matched in a previous period, no trade occurs 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 with the agreed shares; or rejects it, in which case no trade occurs in this period.

This game features a standard *holdup problem*: Each agent pays the full costs of her investment at time $t = 0$, but does not fully appropriate the resulting increase in surplus in the matching stage, limiting her incentives to invest efficiently. Indeed, focusing on (Markov) strategies that only condition on the surpluses that the agents that are yet to match can generate, we now argue that there does not exist any Markov-perfect equilibrium featuring efficient investments. For brevity, we consider the case $\beta = 1/2$, and we focus on the case in which agents are arbitrarily patient.

Towards a contradiction, suppose that an efficient equilibrium exists. Given that the aggregate surplus is bounded above by 4, at least one of the agents has a limit gross payoff that is bounded above by 1. Suppose without loss of generality that the limit gross equilibrium payoff of b_1 is bounded above by 1. Consider a deviation by b_1 to *not invest*. We show that b_1 's limit gross payoff under this deviation is bounded below by $1/2$ —i.e., this deviation reduces her limit gross payoff by at most half of the corresponding reduction in gross aggregate surplus. Since this deviation involves no investment costs, b_1 's limit *net* payoff under this deviation is also bounded below by $1/2$, which is strictly higher than $1 - c$ (the upper bound on her limit equilibrium net payoff). Hence, this deviation is profitable.

The key observation driving the argument is that, when everyone but b_1 invests, b_2 does not delay in equilibrium.⁴ Hence, b_1 can just wait until b_2 matches, and then share the remaining unit of surplus approximately equally with the remaining seller—as specified by the unique subgame perfect equilibrium at that point. As a result, her payoff is bounded below by $1/2$ in the limit as δ goes to 1. Intuitively, the deviator can hold out until her competitor leaves, at which point she faces a bilateral monopoly situation—where the sur-

⁴While the fact that b_2 does not delay in equilibrium seems intuitive enough (it is difficult to imagine how her bargaining position can improve after b_1 leaves), proving this formally requires some care. We relegate the details of the argument to subsection A.1 in Appendix A.

plus loss generated by her deviation is shared with another agent, while she pockets all the associated savings.

2.2 No holdup in a market with sequential entry

Now consider a homologous market featuring *sequential entry*. In the first period $t = 0$, a continuum of identical buyers and a continuum of identical sellers simultaneously make non-contractible investments. As before, they can choose either to *invest* or to *not invest*, and their investments determine the surplus of each match—as specified by (1). Each agent that invests has to pay the investment cost c in the period in which she enters the market. As in the market without sequential entry, when agents are sufficiently patient, efficiency requires that everyone invests.

Once the agents have sunk their investments, they bargain according to the following standard protocol (e.g., Talamàs 2019b): In each period $t = 1, 2, \dots$, there are two active buyers and two active sellers. In particular, in the first period $t = 1$, two buyers and two sellers are selected uniformly at random to be active and, every time a buyer and a seller trade, they leave the market, and a new buyer-seller pair is drawn uniformly at random (from those that are yet to become active) to replace them. The agents in the market bargain exactly as in the case without sequential entry described above. Hence, in each period, both the bargaining protocol and the matching surpluses are exactly as in a subgame that starts in period $t = 1$ of the game without sequential entry described above.

We argue that, in stark contrast to the setting without sequential entry, holdup is not a problem in this game when agents are patient. To illustrate this as simply as possible, we first consider the case of equal bargaining powers ($\beta = 1/2$). Focusing on Markov strategies that condition only on the profile of investments made at $t = 0$ and the investments of the active agents, we show that there exists an efficient Markov-perfect equilibrium (in which everyone invests).⁵

In Appendix B, we show that there exists a Markov-perfect equilibrium of the subgame that starts at $t = 1$ for any profile of investments made in period $t = 0$. Hence, in order to show that there exists an efficient Markov-perfect equilibrium, it is enough to describe Markov-perfect equilibria of the subgames that start with (i) everyone that is yet to enter the market having invested and (ii) all but one of the active agents having invested; and to show that any unilateral investment deviation from a strategy profile in which everyone invests

⁵Furthermore, when agents are arbitrarily patient, every type-symmetric Markov-perfect equilibrium is efficient (see subsection A.2 in Appendix A).

and that is consistent with these equilibria is unprofitable.

Let us start by considering the case of equal bargaining powers ($\beta = 1/2$). First, let us describe Markov-perfect equilibrium strategies for the subgames where everyone that is yet to trade (active or inactive) has invested: Each agent accepts every offer that gives her at least w , and each proposer offers w to an agent on the other side of the market, who accepts. Each agent must be indifferent between accepting and rejecting an offer that gives her w ; that is,

$$(2) \quad w = \delta \left(\frac{1}{4} \underbrace{(2-w)}_{\text{proposer's payoff}} + \frac{3}{4} \underbrace{w}_{\text{non-proposer's payoff}} \right), \text{ or, equivalently, } w = \frac{\delta}{2-\delta}.$$

Second, let us describe Markov-perfect equilibrium strategies for every subgame in which all but one agent, who is active, has invested. As soon as the deviator leaves, switch to the strategy just described. While the deviator is active: Each non-deviator (i) offers w (defined by Equation 2) to some other non-deviator, who accepts with probability one, and (ii) accepts an offer if and only if it gives her at least w . The deviator (i) offers w to some non-deviator, who accepts with probability one, and (ii) accepts an offer if and only if it gives her at least w' . The deviator obtains $1-w$ when she is the proposer, so her cutoff w' must satisfy

$$w' = \delta \left(\frac{1}{4} \underbrace{(1-w)}_{\text{deviator's payoff when proposer}} + \frac{3}{4} \underbrace{w'}_{\text{deviator's payoff when non-proposer}} \right),$$

or, rearranging,

$$(3) \quad w' = w \frac{2-2\delta}{4-3\delta} = w \left(1 - \frac{2-\delta}{4-3\delta} \right) = w - \frac{\delta}{4-3\delta}.$$

Hence, this is indeed an equilibrium, since Equation 3 implies that $2-w > 1-w'$, which implies, in turn, that the best that the non-deviators can do is to obtain $2-w$ when they are the proposers. We conclude that the deviator saves c in the period before she enters, but has an associated expected loss of $\frac{1}{4} + \frac{3}{4} \frac{\delta}{4-3\delta}$. Hence, her deviation is not profitable when $\delta \geq \bar{\delta}$, where $\bar{\delta} := \frac{4}{3} - \frac{1}{3c} < 1$.

The key behind the argument above is that, when everyone invests, the price that a deviator has to pay in order to match with an agent on the other side of the market is not affected by her deviation. As a result, each agent faces the full negative consequences of her own deviation. This conclusion does not depend on the buyers' relative bargaining power β . In general, the limit gross equilibrium payoffs when everyone invests are 2β for buyers, and

$2(1 - \beta)$ for sellers, and, just as above, a unilateral investment deviation does not affect the non-deviators' payoffs. In other words, while different relative bargaining powers shape the limit equilibrium payoffs, they do not affect the fact that agents face the full consequences of their individual investment decisions.

3 Model

There is a finite set I of types of agents, and a continuum of agents of each type. The type of an agent determines her investment opportunities and her resulting gains from trade, as specified below. All the agents have a common discount factor $0 \leq \delta < 1$, common knowledge of the game and perfect information about all the events preceding any of their decision nodes in the game.

3.1 Investment

In the first period $t = 0$, all the agents simultaneously choose their investments: Each agent of type i chooses an investment from a finite set $K_i \subset \mathbb{R}^{m_i}$, where $m_i \geq 1$. An agent of type i with investment profile x_i and an agent of type $j \neq i$ with investment profile x_j produce $y(x_i, x_j) > 0$ units of surplus when they match, and the costs of their investments are $c(x_i)$ and $c(x_j)$, respectively. An agent of type i with investment profile x_i generates $y(x_i, x_i) > 0$ in isolation (this can capture her exogenous outside options, for example). Each agent pays her investment cost in the period in which she enters the market.

Remark 3.1. *Given that the function y determines the surplus of each match only as a function of the investment profiles of its members, this formulation encodes all the heterogeneities among types via their investment opportunities. This can capture arbitrary heterogeneity among different types of agents. For example, suppose that there are two seller types, i' and i'' , and two buyer types, j' and j'' , and further that i' is a much better fit for type j' than j'' is, while i'' is a much better fit for type j'' than j' is. To capture this situation, we can simply take the surplus $y(x_i, x_j)$ associated with any investment profile $(x_i, x_j) \in (K_{i'} \times K_{j'}) \cup (K_{i''} \times K_{j''})$ to be high relative to the associated investment costs, and the surplus $y(x_i, x_j)$ associated with any investment profile $(x_i, x_j) \in (K_{i''} \times K_{j'}) \cup (K_{i'} \times K_{j''})$ to be low relative to the associated investment costs.*

Remark 3.2. *We assume that all investments are decided before any bargaining occurs for two reasons. First, this highlights that our mechanism does not rely on intertwining the investment and bargaining stages (as is the case in Che and Sakovics 2004, for example). Second, this substantially*

simplifies the analysis by allowing us to leverage existing results in the non-cooperative bargaining literature (e.g., Elliott and Nava 2019 and Talamàs 2019b).

3.2 Non-cooperative bargaining

Once everyone chooses her investment in period $t = 0$, bargaining occurs in discrete periods $t = 1, 2, \dots$. For each type i , there are $n_i \geq 2$ bargaining slots. In any given period, each slot of a given type can be occupied by one agent of that type, or be empty. We refer to the agents occupying the slots in any given period as the *active agents* in that period, and we denote the total number of slots by $n := \sum_{i \in I} n_i$.⁶

In each period $t = 1, 2, \dots$, one slot is selected uniformly at random (i.e., each slot is selected with probability $1/n$). If the slot is empty, no trade occurs in this period. Otherwise, its occupant becomes the *proposer*. The proposer a chooses an active agent b (which can be herself) and makes her a take-it-or-leave-it offer specifying a split of the surplus $y(x_a, x_b)$, where x_a and x_b denote agents a and b 's investment profiles, respectively. The receiver of this offer can then *accept* or *reject*. If she accepts, then a and b exit the market with the agreed shares, vacating their respective bargaining slots. Otherwise no trade occurs (and no bargaining slots are vacated) in this period.

3.3 Stochastic entry

For each type i and each $s \leq n_i$, at the beginning of each period that starts with s empty bargaining slots of type i , a number $s' \leq s$ is drawn according to a stationary probability distribution q_s^i . Then, s' agents are drawn uniformly at random from those agents of type i that are yet to become active, and these are randomly assigned to different empty slots of type i . We restrict attention to markets that never become extremely small, in the following sense.

Assumption 3.3. *There are always at least two active agents of each type.*

Assuming that there is always at least one active agent of each type simplifies the analysis by guaranteeing that, when all the agents of the same type choose the same investments, payoffs are uniquely determined by our notion of Markov-perfect equilibrium (Proposition C.2). Assuming that there are always at least two active agents of each type further

⁶Assuming the existence of such bargaining slots is a useful modelling device. This assumption need not be interpreted literally.

simplifies the analysis by guaranteeing that a deviating agent can directly play off the agents of other types, and that no unilateral investment deviation shrinks the relevant bargaining opportunities of the non-deviators.

Remark 3.4. *Assumption 3.3 holds under fairly mild conditions on the stochastic inflow process. It holds, for example, if there is at least one active agent of each type in period $t = 1$, and $q_{n_i-1}^i(0) = 0$ for each type i . It also holds under the classical replica framework of the literature on non-cooperative bargaining in stationary markets (e.g., Rubinstein and Wolinsky 1985 and 1990; Gale 1987; de Fraja and Sákovicš 2001, Manea 2011, Lauermann 2013, Nguyen 2015, Polanski and Vega-Redondo 2018, Talamàs 2019b) as long as there are two or more agents of each type.*

3.4 Histories, strategies and equilibrium

There are three kinds of *histories*. We denote by h_t a history of the game up to—but not including—time t . We denote by $(h_t; i)$ the history that consists of h_t followed by agent i being selected to be the proposer at time t . We denote by $(h_t; i \rightarrow j; s)$ the history that consists of $(h_t; i)$ followed by agent i offering a share s to agent j . A *strategy* σ_i for agent i specifies her investment and, for all possible histories h_t , the offer $\sigma_i(h_t; i)$ that she makes following the history $(h_t; i)$ and her response $\sigma_i(h_t; j \rightarrow i; s)$.

The strategy profile σ is a *type-symmetric Markov-perfect equilibrium* if it induces a Nash equilibrium in every subgame, all the agents of any given type follow the same strategy, and each agent a 's bargaining strategy conditions only on (i) the investment profile, (ii) the set $\{y(x_b, x_c) \mid \text{agents } b, c \text{ active}\}$ of surpluses among the active agents, (iii) the set $\{y(x_a, x_b) \mid \text{agent } b \text{ active}\}$ of surpluses that she can create with the active agents, (iv) for each type i that is such that not all agents of type i yet to enter have chosen the same investment, the number of vacant slots of type i , and (v) the going proposal (in the case of a response).⁷

4 No holdup in equilibrium

Theorem 4.1 below shows that an investment profile $(x_i)_{i \in I}$ can be implemented as a type-symmetric Markov-perfect equilibrium for all sufficiently high discount factors if and only if it is constrained efficient—in the sense that no agent, *taking others' payoffs as given and free to choose whom to match with*, has a profitable investment deviation. In particular, in every

⁷The number of vacant slots of type i is payoff relevant only when not all agents of type i yet to enter have chosen the same investment.

type-symmetric Markov-perfect equilibrium, agents become price takers as they become arbitrarily patient, and hence they appropriate the full returns of their investments.

As background for this result, note that, for each type-symmetric investment profile $x := (x_i)_{i \in I}$ and each type i , there exists $V_i(x) > 0$ such that, in every subgame-perfect equilibrium of the subgame that starts at $t = 1$ with the investment profile x , $V_i(x)$ is the expected equilibrium (gross) payoff at the beginning of each period of each agent of type i (Proposition C.2). We denote the limit of $V_i(x)$ as δ goes to 1 by $V_i^*(x)$.⁸

Theorem 4.1. *A type-symmetric Markov-perfect equilibrium with investment profile $x := (x_i)_{i \in I}$ exists for all sufficiently high discount factors if and only if*

$$(4) \quad x_i \in \underset{z_i \in K_i}{\operatorname{argmax}} \left[\max \left(y(z_i, z_i), \max_{j \in I} [y(z_i, x_j) - V_j^*(x)] \right) - c(z_i) \right] \text{ for each } i \text{ in } I.$$

Proof. Necessity: Fix a type-symmetric Markov-perfect equilibrium σ with investment profile $(x_i)_{i \in I}$. Let v_i and w_i denote the (gross) expected equilibrium payoff of each active agent of type i in a period in which she is and she is not the proposer, respectively. Given that σ is Markov perfect, each agent gets—when she is the proposer—the maximum amount that she can obtain while leaving the receiver indifferent between accepting and rejecting (unless she chooses to match with herself). Hence,

$$v_i = \max \left(y(x_i, x_i), \max_{j \in I} [y(x_i, x_j) - w_j] \right).$$

Given that each agent is selected to be the proposer with probability $1/n$ and that, in equilibrium, no agent is ever offered more than her expected equilibrium payoff, we have that $w_i = \delta \left(\frac{1}{n} v_i + \frac{n-1}{n} w_i \right)$. Rearranging gives

$$w_i = \chi v_i = \chi \max \left(y(x_i, x_i), \max_{j \in I} [y(x_i, x_j) - w_j] \right) \text{ where } \chi := \frac{\delta}{n - \delta(n-1)} \rightarrow 1 \text{ as } \delta \rightarrow 1.$$

Hence, it is enough to show that, for any investment deviation from the equilibrium σ by an agent d , and for any agent $a \neq d$ (of type i , say), a 's expected equilibrium payoff \hat{w}_a when rejecting an offer from d gets arbitrarily close to w_i as δ goes to 1. Indeed, given that the set of investments is finite, and that w_i converges to $V_i^*(x)$ for each type i , when δ is sufficiently close to 1 each agent a must then choose her investment z_a to maximize

$$\max \left(y(z_a, z_a), \max_{j \in I} [y(z_a, x_j) - V_j^*(x)] \right) - c(z_a).$$

⁸Talamàs (2019a) describes a simple algorithm that computes the profile $V^*(x)$ for every investment profile $x := (x_i)_{i \in I}$, and characterizes $V^*(x)$ in terms of the classical Nash bargaining solution.

Suppose that an agent d of type k deviates from σ by investing $x_d \neq x_k$. Assumption 3.3 ensures that the Markov state does not change while the deviator d is active.⁹ Hence, given that σ is Markov perfect, for each agent a we can let \hat{w}_a be her expected equilibrium payoff when rejecting an offer while d is active. Furthermore, when d is the proposer, she offers \hat{w}_a to some agent a , who accepts with probability one.¹⁰

Fix an arbitrary type i , and let $a \neq d$ be an agent of type i such that there exists an agent $c \neq a$ with whom the deviator trades with positive probability in equilibrium (Assumption 3.3 ensures that we can find such an agent). We argue that $\hat{w}_a - w_i$ converges to 0 as δ goes to 1. Since σ is type symmetric, this implies that, for each agent b of type i , $\hat{w}_b - w_i$ also converges to 0 in the limit as δ goes to 1.

Let $\epsilon > 0$. First, note that $\hat{w}_a \geq w_i - \epsilon$ for all sufficiently high discount factors. This is because agent a can always wait for the deviator to leave, and—once this happens—her expected equilibrium payoff (when rejecting an offer) is w_i . We now argue that $\hat{w}_a \leq w_i + \epsilon$ for all sufficiently high discount factors. For contradiction, suppose otherwise. Given that, as we have just argued, for each type j and each agent b of type j other than the deviator, $\hat{w}_b \geq w_j - \epsilon$ for all sufficiently high discount factors, agent a must be making offers to the deviator for all sufficiently high discount factors. For each such discount factor δ , letting $\pi > 0$ be the probability that the deviator trades with someone other than a when a is not the proposer, we have that

$$\hat{w}_a = \delta \left[\frac{1}{n} (y(x_a, x_d) - \hat{w}_d) + \frac{n-1}{n} (\pi w_i + (1-\pi) \hat{w}_a) \right]$$

and, given that d can always make offers to a ,

$$\hat{w}_d \geq \delta \left[\frac{1}{n} (y(x_a, x_d) - \hat{w}_a) + \frac{n-1}{n} \hat{w}_d \right].$$

If the weak inequality holds with equality, it is easy to check that \hat{w}_a gets arbitrarily close to w_i as δ goes to 1, a contradiction. Otherwise, \hat{w}_a is strictly smaller than w_i for all large enough δ , also a contradiction.

Sufficiency: Proposition B.1 shows that there exists a type-symmetric Markov perfect equilibrium in the subgame starting at $t = 1$ for every choice of agents' investments. Hence, for

⁹Indeed, neither the set $\{y(x_a, x_b) \mid \text{agents } a, b \text{ active}\}$ of surpluses among the active agents nor the set $\{y(x_d, x_b) \mid \text{agent } b \text{ active}\}$ of surpluses that the deviator d can generate with the active agents change while the deviator d is active.

¹⁰Note that agent d deviates at the investment stage only, so σ still governs her bargaining strategy. Given that σ is Markov, and that the environment is stationary from the point of view of the deviator d , she can obtain a strictly bigger amount when she is the proposer than when she is the receiver, so she leaves the market—by matching to herself or to someone else—with probability one when she is the proposer.

each investment profile z , we can pick an equilibrium $\sigma(z)$ of the subgame that starts at $t = 1$. Given any investment profile $(x_i)_{i \in I}$, define a strategy profile as follows: All the agents of type i invest x_i , and each agent's bargaining strategy given any investment profile z is as specified by $\sigma(z)$. This strategy profile is a Markov-perfect equilibrium if no agent has incentives to deviate at the investment stage ($t = 0$), which—as argued in the necessity part of the proof—is guaranteed (for all sufficiently high discount factors) by condition (4). \square

The absence of holdup in equilibrium does not imply that every equilibrium involves efficient investments, nor that every efficient investment profile can be implemented in equilibrium. For example, coordination failures can sustain inefficient investment profiles in equilibrium. To see this, consider the example described in subsection 2.2, with the following modification: The surplus of any match is 0 unless both sides invest (in which case this surplus is 2). It is still efficient that everyone invests (for sufficiently high discount factors), but there is an (inefficient) equilibrium in which no one invests. Also, while the absence of holdup guarantees constrained-efficient investments (as defined above), it does not guarantee that all efficient investments can be sustained in equilibrium. For instance, in the example described in subsection 2.2, when buyers' bargaining power β is small, there is no equilibrium in which everyone invests, even though efficiency calls for everyone investing independently of β . Indeed, given that the buyers' limit (gross) equilibrium payoff when everyone invests is 2β , when $\beta < \frac{\epsilon}{2}$, they do not find it worthwhile to invest.¹¹

5 Conclusion

In the context of a general non-cooperative investment and bargaining game, this paper shows that dynamic entry can solve the holdup problem—even in markets that are thin at every point in time. In particular, in stark contrast to the standard conclusion reached under a static view of markets, we show that everyone simultaneously obtains her marginal product as frictions vanish, thus eliminating the holdup problem. This provides non-cooperative foundations for the standard price taking assumption in matching markets, and shows that dynamic entry can significantly ameliorate the holdup problem.

As is standard in the literature, bargaining outcomes are sensitive to the bargaining protocol. But, remarkably, our no holdup result is not. For example, even if different proposer probabilities affect the different types' payoffs in the limit as agents become patient, they

¹¹Cole, Mailath, and Postlewaite (2001a; 2001b) and Nöldeke and Samuelson (2015), for example, investigate these sources of inefficiency in a competitive matching environment that precludes holdup problems.

do not affect the fact that every individual obtains her marginal product in this limit. Our main finding also extends to more general bargaining protocols in which agents have a more restricted choice of whom to make offers to, or where agents match at random (as in Rubinstein and Wolinsky (1985) and most of the subsequent literature).

References

- ACEMOGLU, D. (1996): "A microfoundation for social increasing returns in human capital accumulation," *The Quarterly Journal of Economics*, 111, 779–804.
- (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.
- BERGEMANN, D. AND J. VÄLIMÄKI (2002): "Information acquisition and efficient mechanism design," *Econometrica*, 70, 1007–1033.
- BESTER, H. (2013): "Investments and the holdup problem in a matching market," *Journal of Mathematical Economics*, 49, 302–311.
- 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.
- DAVIS, S. J. AND J. HALTIWANGER (2014): "Labor market fluidity and economic performance," *National Bureau of Economic Research*.

- DE FRAJA, G. AND J. SÁKOVICS (2001): "Walras retrouvé: Decentralized trading mechanisms and the competitive price," *Journal of Political Economy*, 109, 842–863.
- DE MEZA, D. AND B. LOCKWOOD (2010): "Too much investment? A problem of endogenous outside options," *Games and Economic Behavior*, 69, 503–511.
- DECKER, R. A., J. C. HALTIWANGER, R. S. JARMIN, AND J. MIRANDA (2018): "Changing business dynamism and productivity: Shocks vs. responsiveness," *National Bureau of Economic Research*.
- 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.
- FELLI, L. AND K. ROBERTS (2016): "Does Competition Solve the Hold-up Problem?" *Economica*, 83, 172–200.
- 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.
- 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.
- 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.
- 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 (2013): "Pricing and investments in matching markets," *Theoretical Economics*, 8, 535–590.
- (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.
- (2017): "Bargaining in dynamic markets," *Games and Economic Behavior*, 104, 59–77.
- MOLLOY, R., R. TREZZI, C. L. SMITH, AND A. WOZNIAK (2016): "Understanding declining fluidity in the US labor market," *Brookings Papers on Economic Activity*, 2016, 183–259.
- 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.
- PETERS, M. AND A. SIOW (2002): "Competing premarital investments," *Journal of Political Economy*, 110, 592–608.
- 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. (2019a): "Nash bargaining with endogenous outside options," *Mimeo*.
- (2019b): "Price dispersion in stationary networked markets," *Games and Economic Behavior*, 115, 247–264.
- TIROLE, J. (1986): "Procurement and renegotiation," *Journal of Political Economy*, 94, 235–259.
- WILLIAMSON, O. E. (1975): *The Economic Institutions of Capitalism*, New York Free Press.

Appendices

A Details omitted from section 2

A.1 Details omitted from subsection 2.1

We show formally the key observation from the example in subsection 2.1: When everyone but b_1 invests, b_2 does not delay in equilibrium. Towards a contradiction, suppose that, when b_2 is the proposer, she delays (makes an unacceptable proposal) with probability $\pi > 0$, she makes an acceptable offer to s_1 with probability $(1 - \pi)\beta$, and she makes an acceptable offer to s_2 with probability $(1 - \pi)(1 - \beta)$.

First note that, since b_2 delays, neither the sellers nor the other buyer b_1 can delay. To see this, let w_i denote i 's expected equilibrium gross payoff in any given period (before anyone has matched) conditional on not trading in this period. It follows from $\min\{w_{b_2} + w_{s_1}, w_{b_2} + w_{s_2}\} \geq 2$ (which holds because b_2 delays) and $w_{b_1} + w_{s_2} + w_{b_2} + w_{s_1} < 3$ (which holds because the aggregate discounted surplus is below 3) that $w_{b_1} + w_{s_1} < 1$ and $w_{b_1} + w_{s_2} < 1$. Hence, in equilibrium, b_1 makes an acceptable offer with probability one to either s_1 or s_2 when she is the proposer. Moreover, by the same argument, the sellers both make acceptable offers to b_1 with probability one when they are the proposers.

Second, letting w be the amount that a seller that is yet to match is indifferent between accepting and rejecting in a subgame in which b_1 has already matched, the fact that b_2 delays implies that $w_{s_1} > w$. We have that

$$(5) \quad w = \delta \left(\frac{1}{4}(2 - w) + \frac{3}{4}w \right)$$

and that

$$(6) \quad w_{s_1} = \delta \left(\underbrace{\frac{1}{4}(1 - w_{b_1})}_{s_1 \text{ proposes}} + \underbrace{\frac{1}{4}\pi w_{s_1} + (1 - \pi)(\beta w_{s_1} + (1 - \beta)w')}_{b_2 \text{ proposes}} + \underbrace{\frac{1}{4}(\kappa w_{s_1} + (1 - \kappa)w)}_{b_1 \text{ proposes}} + \underbrace{\frac{1}{4}w}_{s_2 \text{ proposes}} \right)$$

where κ denotes the probability that b_1 makes an offer to s_1 in any period before anyone has matched when b_1 is the proposer, and w' is the quantity that a seller that is yet to match is indifferent between accepting and rejecting in a subgame in which b_2 has already matched.

Given that $2 - w \geq 1 \geq 1 - w_{b_1}$, that $w' < w$, and that $w_{s_1} = w_{s_2}$ unless $\beta = 1$ (because b_2 makes offers only to a seller with the lowest cutoff), the combination of Equation 5 and Equation 6 implies that $w_{s_1} \leq w$, a contradiction.

A.2 Details omitted from subsection 2.2

We show that—for all sufficiently high discount factors—every type-symmetric Markov-perfect equilibrium of the game described in subsection 2.2 is efficient. Suppose for contradiction that there exists a sequence \mathcal{D} of discount factors converging to 1, such that, for each $\delta \in \mathcal{D}$, there exists a type-symmetric Markov-perfect equilibrium σ in which only the sellers invest. We show that, for all sufficiently high $\delta \in \mathcal{D}$, a buyer can profitably deviate by investing. A similar argument shows that—when agents are arbitrarily patient—there exists no type-symmetric Markov-perfect equilibrium in which only the buyers invest, or in which neither the buyers nor the sellers invest.

Let $\delta \in \mathcal{D}$ and consider the associated equilibrium σ . On the equilibrium path, each agent's expected payoff when she rejects an offer satisfies

$$(7) \quad w = \delta \left(\frac{1}{4} \underbrace{(1-w)}_{\text{proposer's payoff}} + \frac{3}{4} \underbrace{w}_{\text{non-proposer's payoff}} \right), \text{ that is, } w = \frac{\delta}{4-2\delta} \rightarrow \frac{1}{2}.$$

Suppose that buyer b'_1 deviates and invests, and consider a subgame in which b'_1 is active. From the point of view of b'_1 , the environment is stationary. Hence, when she is the proposer, she makes offers that leave the receiver indifferent between accepting and rejecting, and which are accepted with probability one.¹²

Let s be a seller such that there exists a seller $s' \neq s$ with whom the deviator trades with positive probability, and let \hat{w}_s be her expected equilibrium payoff when rejecting an offer from b'_1 . We show that, in the limit as δ goes to 1, \hat{w}_s converges to $1/2$. Given that σ is Markov perfect and specifies that each agent of the same type follows the same strategy, this implies that every other seller's expected equilibrium payoff when rejecting an offer from b'_1 also converges to $1/2$, so when δ is sufficiently high this deviation is profitable (the deviator's net gain is $1 - c > 0$).

First, we argue that \hat{w}_s is *bounded below* by $1/2$ in the limit as δ goes to 1. Given that the seller s can always wait for the deviator to leave (at which point she obtains $1/2$), for each $\epsilon > 0$, \hat{w}_s is bounded below by $1/2 - \epsilon$ for all high enough $\delta \in \mathcal{D}$. A similar argument shows that the expected equilibrium payoff \hat{w}_b of each buyer $b \neq b'_1$ conditional on not being the proposer in a period in which b'_1 is active is bounded below by $1/2$ in the limit as δ goes to 1.

Second, we argue that \hat{w}_s is *bounded above* by $1/2$ in the limit as δ goes to 1. Suppose for contradiction that there exists $\epsilon > 0$ such that $\hat{w}_s \geq 1/2 + \epsilon$ for all sufficiently high $\delta \in \mathcal{D}$.

¹²Note that we are considering an *investment* deviation from σ , so the Markov-perfect equilibrium σ still governs the deviator's bargaining strategy.

Given that, as we have just argued, for each buyer $b \neq b'_1$, \hat{w}_b is bounded below by $1/2$ in the limit as δ goes to 1, seller s must be making offers to the deviator b'_1 for all sufficiently high $\delta \in \mathcal{D}$. For each such discount factor δ , letting $\pi > 0$ be the probability that the deviator trades with someone other than s when s is not the proposer, we have that

$$\hat{w}_s = \delta \left[\frac{1}{4}(2 - \hat{w}_{b'_1}) + \frac{3}{4}(\pi w + (1 - \pi)\hat{w}_s) \right] \text{ and } \hat{w}_{b'_1} = \delta \left[\frac{1}{4}(2 - \hat{w}_s) + \frac{3}{4}\hat{w}_{b'_1} \right],$$

which implies that \hat{w}_s converges to $1/2$ as δ goes to 1, a contradiction.

B Existence of a type-symmetric Markov-perfect equilibrium

Proposition B.1 below is analogous to the Markov-perfect equilibrium existence proof in Elliott and Nava (2019).

Proposition B.1. *For every investment profile x , there exists a strategy profile that is a type-symmetric Markov-perfect equilibrium of the subgame starting in period $t = 1$ with investment profile x .*

Proof. Let the *kind* of an agent be determined by her type and her investment profile. Without loss of generality, we can assume that the investment sets $\{K_i\}_{i \in I}$ do not overlap, so we can identify the set of agent kinds by $K := \cup_{i \in I} K_i$, which is finite because each K_i is itself finite. Let m denote the number of elements of K . We abuse terminology by referring to $i \in K$ as “agent i .” Let \mathcal{K} be the finite set of all possible profiles of agents that can be active in the market at any given time. We characterize the Markov perfect equilibrium of the subgame that starts at $t = 1$ with any given investment profile, and then use it to show that such an equilibrium exists.

Consider a Markov-perfect-equilibrium strategy profile and its corresponding value function $V : \mathcal{K} \rightarrow \mathbb{R}^m$, where $V(K)$ gives each agent’s expected equilibrium payoff in any period at the beginning of a period that starts with active agent set K (before any agents become active this period). Consider a subgame with active agent set $K \in \mathcal{K}$, and let s_{ij} denote the surplus that agents i and j generate when they match in this subgame. By Markov perfection, agent j accepts any offer that gives her strictly more than $\delta V_j(K)$, and rejects any offer that gives her strictly less than $\delta V_j(K)$. This implies that no one offers more than $\delta V_j(K)$ to any agent j . Therefore, a proposer i makes offers with positive probability only to j that maximizes her net payoff $s_{ij} - \delta V_j(K)$. Hence, when $i \in K$ is the proposer, the expected

payoff of $k \in K \setminus \{i\}$ is

$$\sum_{j \in K \setminus \{i,k\}} \pi_{ij} \delta V_k(K \setminus \{i,j\}) + \left(1 - \sum_{j \in K \setminus \{i,k\}} \pi_{ij}\right) \delta V_k(K)$$

where π_{ij} denotes the probability that i and j agree to trade. When i is the proposer, if there exists $j \in K$ such that $\delta(V_i(K) + V_j(K)) < s_{ij}$, then she makes offers only to $j \in K$ for which $s_{ij} - \delta V_j(K)$ is maximum, and agreement obtains with probability one. Otherwise, she delays—in the sense that she makes offers that are not accepted in equilibrium. We denote the probability that $i \in K$ delays by π_{ii} . Thus, any agreement probability profile $\pi_i(K) \in \Delta(K)$ —corresponding to the histories in which i is the proposer—that is consistent with the value function V must be in

$$\Pi^{i,K}(V) = \left\{ \pi_i \in \Delta(K) \left| \begin{array}{l} \pi_{ii} = 0 \quad \text{if } \delta V_i < \max_{j \in K \setminus \{i\}} \{s_{ij} - \delta V_j(K)\}, \\ \pi_{ik} = 0 \quad \text{if } s_{ik} - \delta V_k(K) < \max\{\delta V_i(K), \max_{j \in K \setminus \{i\}} \{s_{ij} - \delta V_j(K)\}\} \end{array} \right. \right\}.$$

For any value function V , any $K \in \mathcal{K}$ and any agent $i \in K$, define $f^{i,K}(V) : K \rightarrow \mathbb{R}^m$ by

$$\begin{aligned} f_i^{i,K}(V) &= \pi_{ii} \delta V_i(K) + (1 - \pi_{ii}) \max_{j \in K \setminus \{i\}} \{s_{ij} - \delta V_j(K)\} \\ f_k^{i,K}(V) &= (\pi_{ii} + \pi_{ik}) \delta V_k(K) + \sum_{j \in K \setminus \{i,k\}} \pi_{ij} \delta V_k(K \setminus \{i,j\}) \quad \forall k \neq i, \end{aligned}$$

for any $\pi_i \in \Pi^{i,K}(V)$. That is, $f_i^{i,K}(V)$ gives the set of expected payoffs that are consistent with the value function V in any history in which active agent set is K and the proposer is agent i . Letting \mathcal{V} denote the set of value functions $V : \mathcal{K} \rightarrow \mathbb{R}^m$, consider the correspondence $F : \mathcal{V} \rightarrow \mathcal{V}$ defined by

$$(8) \quad F(V)(K) = \frac{1}{n} \sum_{i \in K} f^{i,K}(V), \text{ for all value functions } V \text{ and all } K \in \mathcal{K}.$$

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

C For online publication: Unique equilibrium payoffs

Proposition C.2 shows that, as long as there is always at least one agent of each type active in the market, the notion of subgame-perfect equilibrium pins down the payoffs of all agents conditional on their (type-symmetric) investment strategies. Moreover, these payoffs are independent of the details of the process by which bargaining slots are filled. This is a slight generalization of the analogous result in Talamàs (2019b), where it is assumed that exactly one agent of each type is active in the market at each point in time. Proposition C.2 holds under the following assumption, which is weaker than Assumption 3.3.

Assumption C.1. *There is always at least one active agent of each type.*

Proposition C.2. *Fix an investment profile $x := (x_i)_{i \in I}$, and suppose that Assumption C.1 holds. For every type i , there exists a value $V_i(x) > 0$ such that, in every subgame-perfect equilibrium with investment profile x , the expected equilibrium payoff of each active agent of type i at the beginning of each period is $V_i(x)$.*

The proof of Proposition C.2 is identical to the corresponding result in Talamàs (2019b), which is itself similar to the proof of the analogous result in Manea (2017) in the context of a model with random matching (as opposed to the framework with strategic choice of partners that we focus on in this paper). We provide this proof here for completeness.

Proposition C.2 follows from Proposition C.3, since every subgame-perfect equilibrium of a game with perfect information (as the one we study) survives the process of iterated conditional dominance (Theorem 4.3 in Fudenberg and Tirole 1991).

Following Fudenberg and Tirole (1991, page 128), we define iterated conditional dominance on the class of multi-stage games with observed actions as follows.

Definition C.1. Action a_i^t available to some agent i at information set H_t is *conditionally dominated* if every strategy of agent i that assigns positive probability to action a_i^t in the information set H_t is strictly dominated. *Iterated conditional dominance* is the process that, at each round, deletes every conditionally-dominated action given the strategies that have survived all the previous rounds.

Fudenberg and Tirole (1991) show how iterated conditional dominance solves the alternating-offers bilateral model of Rubinstein (1982). Manea (2017) shows how iterated conditional dominance also solves a wide class of models similar to the one considered in this article. We prove Proposition C.3 using the techniques developed in Manea (2017).

Proposition C.3. Fix an investment profile $(x_i)_{i \in I}$. For every type i , there exists $w_i > 0$ such that, in every game in which Assumption C.1 holds, after the process of iterated conditional dominance, every agent of type i always accepts (rejects) an offer that gives her strictly more (less) than w_i .

Proof. The proof consists of two steps. First, we define recursively two sequences $(m_i^k)_{i \in I}$ and $(M_i^k)_{i \in I}$, and show by induction on k that after every step s of iterated conditional dominance (see below for a formal definition of such a step), each agent of type i always rejects every offer that gives her strictly less than δm_i^s and always accepts every offer that gives her strictly more than δM_i^s . Second, we show that both sequences $(m_i^k)_{i \in I}$ and $(M_i^k)_{i \in I}$ converge to the same point $(w_i)_{i \in I}$.

We denote the surplus $y(x_i, x_j)$ that a buyer of type i and a seller of type j generate when they match by s_{ij} .

(i) Iterated Conditional Dominance Procedure

Let us start by reviewing how the process of iterated conditional dominance works in the present context. For simplicity, we break up the procedure into steps $0, 1, \dots$, with each step containing three rounds.

Step 0.

Round 0a. Note that a strategy that ever accepts with positive probability a negative share is strictly dominated by the strategy *reject all offers and make only offers that give me a positive share*. These are all the actions that are eliminated in Round 0a. Hence, after this round every agent of type i always rejects every offer that gives her strictly less than δm_i^0 , where

$$(9) \quad m_i^0 := 0.$$

Round 0b. Given the actions that survive round 0a, each agent of type i has an expected payoff (at the beginning of the period, before the proposer has been chosen) of at most M_i^0 , where

$$(10) \quad M_i^0 := \max_j \{s_{ij}\}.$$

because, by assumption, no agent of type j can ever offer any agent of type i a payoff higher than s_{ij} , and, by the actions eliminated in round 0a, no agent ever accepts a negative payoff. Hence, every strategy κ of an agent of type i that ever rejects with positive probability an offer a that gives her strictly more than δM_i^0 is strictly dominated by a similar strategy κ' that specifies *accept a with probability π* in every instance in which κ specifies *reject a with probability π* . These are all the actions that are eliminated in Round 0b; so after this round every agent of type i always accepts every offer that gives her strictly more than δM_i^0 .

Round 0c. Given the actions that survive rounds 0a and 0b, every strategy κ of every agent of type i that ever makes an offer with positive probability that gives $y > \delta M_j^0$ to an agent of type j is strictly dominated by a similar strategy κ' that specifies *offer* $y - \epsilon > \delta M_j^0$ to agent j with probability π in every instance in which κ specifies *offer* y to an agent of type j with probability π , since every agent of type j must accept both y and $y - \epsilon$. These are all the actions that are eliminated in round 0c; after this round *no agent ever makes an offer giving* $y > \delta M_j^0$ to any agent of type j .

Proceeding inductively, imagine that, after step $s = k \in \mathbb{Z}_{\geq 0}$, we have concluded (as we have just done for the case $s = 0$) that every agent of type i :

1. rejects every offer that gives her strictly less than δm_i^s ,
2. has an expected payoff (at the beginning of each period) of at most M_i^s ,
3. accepts every offer that gives her strictly more than δM_i^s , and
4. does not make offers that give strictly more than δM_j^s to any agent of type j .

We now show that points (1) to (4) also hold at step $s = k + 1$.

Step $k + 1$.

We refer to strategies that assign positive probability only to actions that have survived all previous rounds of iterated conditional dominance as “surviving strategies.”

Round (k+1)a. Given the surviving strategies, it is conditionally dominated for any agent of type i to ever accept an offer that gives her a surplus strictly lower than δm_i^{k+1} , where m_i^{k+1} is defined as follows:

$$(11) \quad m_i^{k+1} := \frac{1}{n} \max \left(\max_{j \in N} (s_{ij} - \delta M_j^k), \delta m_i^k \right) + \frac{n-1}{n} \delta m_i^k$$

To see this, consider a period- t subgame where an agent of type i has to respond to an offer $x < \delta m_i^{k+1}$. We argue that, for sufficiently small $\epsilon > 0$, accepting this offer is conditionally dominated by the following plan of action—which is designed to give her a time- t expected payoff that approaches δm_i^{k+1} as ϵ goes to 0: *Reject all offers received at dates $t' \geq t$. When selected to be the proposer at time t' , offer $\delta M_j^{k+t+1-t'} + \epsilon$ if $t' \in [t+1, t+k+1]$ and $\max_{j \in N} (s_{ij} - \delta M_j^{k+t+1-t'}) > \delta m_i^{k+t+1-t'}$, and make an unacceptable offer otherwise (e.g. offer a negative amount to some agent).*

Note that since $t' \geq t + 1$, we have that $k + t + 1 - t' \leq k$. Hence, by the induction hypothesis, all agents j accept the offer $\delta M_j^{k+t+1-t'} + \epsilon$ at period $t' \in [t + 1, t + k + 1]$. Moreover, note that

Equation 11 can be written as

$$(12) \quad m_i^{k+1} = \begin{cases} \delta m_i^k & \text{if } \max_{j \in N} (s_{ij} - \delta M_j^k) \leq \delta m_i^{k+t+1-t'} \\ \frac{1}{n} \max_{j \in N} (s_{ij} - \delta M_j^k) + \frac{n-1}{n} \delta m_i^k & \text{otherwise} \end{cases}$$

and an analogous equation can be used to expand the term m_i^k in Equation 12, and then m_i^{k-1} in the resulting equation, and so on until reaching $m_i^0 = 0$. It is clear from the resulting formula for m_i^{k+1} that, under the surviving strategies, the strategy constructed above generates an expected period- t payoff for i of δm_i^{k+1} as $\epsilon \rightarrow 0$. Hence, letting $\epsilon > 0$ be sufficiently small, this strategy conditionally dominates accepting x in period t . These are the actions eliminated in round (k+1)a; after this round *no agent of type i ever accepts any offer that gives her a surplus lower than δm_i^{k+1} .*

Round (k+1)b. Given the surviving strategies, it is conditionally dominated for any agent of type i to reject an offer that gives her strictly more than δM_i^{k+1} , where M_i^{k+1} is defined by

$$(13) \quad M_i^{k+1} := \frac{1}{n} \max \left(\max_{j \in N} (s_{ij} - \delta m_j^k), \delta M_i^k \right) + \frac{n-1}{n} \delta M_i^k$$

To prove this, we show that for each agent of type i , all surviving strategies deliver expected payoffs of at most M_i^{k+1} at the beginning of period t . First, consider a period- t subgame where i is the proposer. Note that i cannot make an offer that generates an expected payoff greater than

$$\max \left(\max_{j \in N} (s_{ij} - \delta m_j^k), \delta M_i^k \right).$$

To see this note that, under the surviving strategies, all agents of type j reject all offers lower than δm_j^k , and when an agent of type j rejects an offer, every agent of type i can expect a period- $(t+1)$ payoff of at most M_i^k . Second, consider a period- t subgame where an agent of type i is not the proposer; under the surviving strategies, this agent can expect a period- t payoff of at most M_i^k . Therefore, *agent of type i has an expected payoff (at the beginning of each period) of at most M_i^{k+1} .* These are all the actions that are eliminated in round (k+1)b; after this round, *no agent ever offers strictly more than δM_j^{k+1} to any agent of type j .*

Round (k+1)c. Given the surviving strategies, every strategy κ of agent of type i that ever makes an offer that gives $y > \delta M_j^{k+1}$ to agent of type j is strictly dominated by a similar strategy κ' that specifies *offer $y - \epsilon > \delta M_j^{k+1}$ to agent of type j with probability π* in every instance in which κ specifies *offer y to agent of type j with probability π* , since every agent of type j must accept both y and $y - \epsilon$. These are all the actions that are eliminated in round (k+1)c; after this round *no agent ever makes an offer giving $y > \delta M_j^{k+1}$ to any agent of type j .*

(ii) The sequences $(m_i^k)_{i \in N}$ and $(M_i^k)_{i \in N}$ converge to the same limit.

First, we prove by induction on k that for all $i \in N$, the sequence $(m_i^k)_{k \geq 0}$ is increasing in k , the sequence $(M_i^k)_{k \geq 0}$ is decreasing in k , and $\max_{j \in N} (s_{ij}) \geq M_i^k \geq m_i^k \geq 0$ for all $k \geq 0$. This implies that both sequences $(m_i^k)_{i \in N}$ and $(M_i^k)_{i \in N}$ converge.

Note that $m_i^0 = 0$ and $M_i^0 := \max_j \{s_{i,j}\}$, and that Equation 11 and Equation 13 imply that $m_i^1 \geq 0$ and $M_i^1 \leq \max_j \{s_{i,j}\}$, so $m_i^1 \geq m_i^0$ and $M_i^1 \leq M_i^0$. Now suppose that for some $l \in \mathbb{N}$:

$$m_i^l \geq m_i^{l-1} \text{ and } M_i^l \leq M_i^{l-1}.$$

We show that

$$m_i^{l+1} \geq m_i^l \text{ and } M_i^{l+1} \leq M_i^l.$$

Note that, by the induction hypothesis, every summand in Equation 11 when $k = l + 1$ is smaller than when $k = l$, which implies that $m_i^{l+1} \leq m_i^l$. Similarly, every summand in Equation 13 when $k = l + 1$ is bigger than when $k = l$, which implies that $M_i^{l+1} \geq M_i^l$. Hence, the sequence $(m_i^k)_{k \geq 0}$ is increasing in k and the sequence $(M_i^k)_{k \geq 0}$ is decreasing in k , which, implies that

$$\max_{j \in N} (s_{ij}) \geq M_i^k \geq m_i^k \geq 0 \text{ for all } k \geq 0.$$

since $\max_{j \in N} (s_{ij}) = M_i^0 > m_i^0 = 0$.

Second, we show that the sequences $(m_i^k)_{i \in N}$ and $(M_i^k)_{i \in N}$ converge to the same limit. Let D^k be $\max_{i \in N} (M_i^k - m_i^k)$. We show that

$$D^k \leq \left(\max_{j \in N} \delta \right)^k D^0 = \left(\max_{j \in N} \delta \right)^k \max_{j, j' \in N} (s_{jj'})$$

for all $k \geq 0$; that is, that D^k converges to 0 as k grows large. Indeed,

$$\begin{aligned}
D^{k+1} &= \max_{i \in N} [M_i^{k+1} - m_i^{k+1}] \\
&= \max_{i \in N} \left[\frac{1}{n} \max \left(\max_{j \in N} (s_{ij} - \delta m_j^k), \delta M_i^k \right) + \left(1 - \frac{1}{n}\right) \delta M_i^k \right. \\
&\quad \left. - \frac{1}{n} \max \left(\max_{j \in N} (s_{ij} - \delta M_j^k), \delta m_i^k \right) + \left(1 - \frac{1}{n}\right) \delta m_i^k \right] \\
&= \max_{i \in N} \left[\frac{1}{n} \left[\max \left(\max_{j \in N} (s_{ij} - \delta m_j^k), \delta M_i^k \right) - \max \left(\max_{j \in N} (s_{ij} - \delta M_j^k), \delta m_i^k \right) \right] \right. \\
&\quad \left. + \left(1 - \frac{1}{n}\right) [\delta M_i^k - \delta m_i^k] \right] \\
&\leq \max_{i \in N} \left[\frac{1}{n} \left[\max \left(s_{ij'} - \delta m_{j'}^k, \delta M_i^k \right) - \max \left(s_{ij'} - \delta M_{j'}^k, \delta m_i^k \right) \right] \right. \\
&\quad \left. + \left(1 - \frac{1}{n}\right) [\delta M_i^k - \delta m_i^k] \right] \\
&\leq \max_{i \in N} \left[\frac{1}{n} \max \left(\delta (M_{j'}^k - m_{j'}^k), \delta (M_i^k - m_i^k) \right) + \frac{n-1}{n} \delta D^k \right] \\
&\leq \max_{j \in N} \delta D^k
\end{aligned}$$

where j' in the first inequality is any element of $\operatorname{argmax}_{j \in N} (s_{i,j} - \delta M_j^k)$, and the second inequality is a consequence of Lemma C.4 below. \square

Lemma C.4 (Manea 2017). *For all $w_1, w_2, w_3, w_4 \in \mathbb{R}$,*

$$|\max(w_1, w_2) - \max(w_3, w_4)| \leq \max(|w_1 - w_3|, |w_2 - w_4|).$$