No Holdup in Dynamic Markets*

Matthew Elliott[1] and Eduard Talamàs[2]

[1]*Faculty of Economics, University of Cambridge.*
[2]*Department of Economics, University of Pennsylvania.*

Abstract

Different types of agents make non-contractible investments before bargaining over both who matches to whom and the terms of trade. In thin markets, the *holdup problem*—that is, underinvestment caused by agents receiving only a fraction of the returns from their investments—is ubiquitous. However, we show that holdup is not a problem in markets that attract traders over time—*even when only a few traders are present in the market at any point in time*. In particular, we characterize the type-symmetric Markov perfect equilibria of a non-cooperative investment and bargaining game with sequential entry, and we show that—in every such equilibrium—all the agents receive the full returns from their marginal investments in the limit as they become patient. Intuitively, the option to wait for future market participants creates competition—so even apparently-thin markets can be competitive. This provides non-cooperative foundations for the standard price-taking assumption in the literature investigating investment efficiency in competitive matching markets.

Keywords: Holdup, efficiency, bargaining, competition, dynamic entry, outside options.

***JEL*:** C72, C78, D40, D41.

**Date Printed:* July 28, 2018. We thank Benjamin Golub, George Mailath and Andrew Postlewaite for useful conversations, and various audiences for useful feedback. Financial support from the European Research Council under the grant "EMBED" #757229 (Elliott), and the Warren Center for Network & Data Sciences and the Rockefeller Foundation #2017PRE301 (Talamàs) is acknowledged. All errors are our own.

1 Introduction

In many markets, important investments are sunk by the time agents bargain over prices. For example, in labor markets, workers and firms invest in human and physical capital before bargaining over wages. These investments are crucial determinants of the efficiency of these markets, but their non-contractible nature implies that they can be subject to holdup problems (e.g. Williamson 1975, Grossman and Hart 1986, Hart and Moore 1990): Agents might not have incentives to invest efficiently because they might not expect to fully appropriate the returns of their investments. Indeed, there is a large literature suggesting that these holdup problems are commonplace and can be severe (e.g., Grout 1984, Hosios 1990, Acemoglu 1996, 1997, Cole et al. 2001a, Elliott 2015).

In this paper, we model these markets as a non-cooperative investment and bargaining game featuring inflows and outflows of traders: Different types of agents sequentially enter the market—after making non-contractible investments—to look for trading partners, and each agent exits the market once she agrees to match with another agent. We characterize the type-symmetric Markov-perfect equilibria, and we show that—in the limit as agents become patient—holdup is not a problem in such equilibria. Intuitively, even if the market appears thin at any point in time, the future entry of market participants creates competition—and the market becomes perfectly competitive as agents become perfectly patient.

There is a large literature—dating back at least to Becker 1975—that investigates the holdup problem in matching markets. An important message of this literature is that only in special circumstances can all the agents receive the marginal returns from their investments. To illustrate this, consider a one-to-one matching environment with a finite set of agents (an *assignment game*). Assume that, once investments are sunk, an allocation in the *core* is implemented. Leonard 1983 (see also Demange 1982) shows that an agent receives her marginal contribution to total surplus—and is thus fully rewarded for her investment—if and only if she receives her highest possible core payoff. Hence, the way to guarantee investment efficiency on one side of the market is to select the point in the core that gives each agent on this side her maximum core payoff.[1] But this implies that each agent on the other side of the market necessarily receives her lowest possible payoff, so only when everyone's maximum and minimum core payoffs coincide do all agents have efficient investment incentives. In

[1] Kranton and Minehart 2001 and Felli and Roberts 2016 study bargaining outcomes that select the point of the core that gives each agent on one side of the market her maximum core payoff, guaranteeing that there is no holdup on this side—but holdup is still problem when there are also investment opportunities on the other side.

other words, the holdup problem is eliminated if and only if the core is a singleton.[2] From this perspective, the contribution of this paper is to show that the conditions under which holdup is precluded are significantly more general once we take into account that agents enter the market over time.

This paper complements the literature that studies investment efficiency in competitive environments (e.g. Cole et al. 2001b, Peters and Siow 2002, Nöldeke and Samuelson 2015). The work in this literature typically assumes that agents are price takers—which precludes holdup—in order to investigate other possible sources of investment inefficiencies (like coordination failures and participation constraints).[3] But, as discussed above, in finite markets agents are price takers only under very special conditions. Hence, by describing a dynamic non-cooperative model that features no holdup even in markets that may appear thin at any point in time, we provide non-cooperative foundations for this price-taking assumption.[4] This clarifies the conditions under which we expect a market to be competitive enough to preclude holdup problems. In particular, our results suggest that the dynamic entry of agents into the market and their patience are important determinants of how competitive it is—and hence of its associated investment efficiency.[5]

The rest of this paper is organized as follows. In section 2, we illustrate the main ideas in the context of a simple buyer-seller example. In section 3 we describe the model, in section 4 we present and prove our main result, and in section 5 we discuss it. We relegate the details of section 2 to Appendix A. In Appendix B and Appendix C, we describe relatively standard facts that we use to prove our main result.

[2] Cole et al. 2001a provide conditions under which the core is a singleton. In essence, these require the market to be sufficiently thick.

[3] An important exception is Acemoglu and Shimer 1999, where firms post wages and workers direct their search having sunk their investments. The fact that wages are posted—not bargained—is crucial for their result. In contrast, we show that the holdup problem is overcome in a dynamic bargaining model.

[4] Makowski and Ostroy 1995 identify two conditions under which a version of *The First Theorem of Welfare Economics* holds. The first condition corresponds to the absence of holdup problems, and the second condition to no coordination failures. While we focus on the holdup problem but are silent about coordination problems, the competitive matching literature focuses on coordination failures but is silent about holdup problems.

[5] In a bilateral setting, Che and Sákovics 2004 show that allowing agents to continue to invest while they bargain can lead to the existence of equilibria in which the holdup problem is eliminated. Instead, we illustrate a market-based mechanism—in which agents sink their investments before bargaining starts—that precludes holdup.

2 Example

In subsection 2.1, we describe a simple investment and bargaining game featuring two buyers and two sellers, and we review the classic holdup problem in this setting. In subsection 2.2, we illustrate how the holdup problem vanishes in the homologous market with sequential entry.

2.1 Holdup in a market with no entry

In the first period $t = 0$, two identical buyers, b_1 and b_2, and two identical sellers, s_1 and s_2, simultaneously make (non-contractible) investments. They can choose either to *invest* or to *not invest*. When a buyer and a seller match, they generate

(1)

2	units of surplus	if *both* have invested,
1	unit of surplus	if *only one of them* has invested, and
0	units of surplus	if *none of them* has invested.

Not investing costs zero, and investing costs c, with $1/2 < c < 1$. Hence, when agents are sufficiently patient (so that waiting to obtain the returns from their investments in a subsequent bargaining stage is not too painful), *efficiency requires that everyone invests*.

Once the agents have sunk their investments, they bargain according to the following standard protocol (e.g. Elliott and Nava forthcoming): In each period $t = 1, 2, \ldots$, one agent is selected uniformly at random to be the proposer. If the selected agent has already matched, no trade occurs in this period. Otherwise, the proposer chooses an 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* or *rejects*. If she accepts, the pair match with the agreed shares. Otherwise, no trade occurs in this period. Agents have perfect information and a common discount factor δ, and the game is common knowledge.

This game features the classic *holdup problem*: Each agent pays the full costs of her investment, but she receives only a fraction of the resulting increase in surplus. As a result, agents do not have incentives to invest efficiently. Indeed, we now argue that there does not exist any efficient *type-symmetric Markov-perfect equilibrium*—that is, a subgame-perfect equilibrium in which (i) everyone invests, (ii) agents' proposals depend only on the surpluses that the unmatched agents can generate, and (iii) all the buyers (sellers) follow the same strategy. For brevity, we focus on the case in which agents are patient (i.e. δ is close to 1), and we relegate the details to subsection A.1.

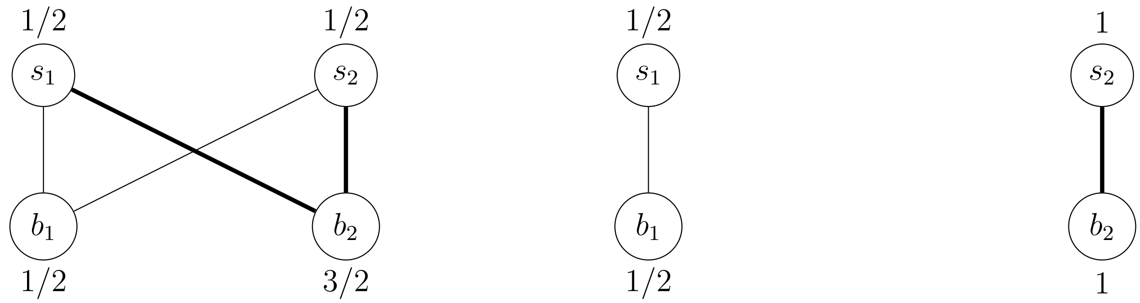

(i) No match has formed. (ii) b_1 is not in first match. (iii) b_1 is in first match.

Figure 1: Surplus profiles that arise in all the relevant subgames (up to isomorphism) of the market with no entry when everyone but b_1 invests. Thin and thick links represent 1 and 2 units of surplus, respectively. The number associated with each agent is her limit (gross) payoff in a type-symmetric Markov-perfect equilibrium conditional on trading in a subgame with these surpluses.

Suppose for contradiction that such an efficient equilibrium exists. On the equilibrium path, the market is fully symmetric, and surplus within each match is split equally in the limit as δ goes to 1. Consider the case in which the buyer b_1 deviates from this equilibrium by not investing. Figure 1 illustrates the available surpluses in all the relevant subgames (up to isomorphism). Intuitively, since b_1 has the option to delay—and share 1 unit of surplus equally with the seller that remains after the first match forms—her payoff is bounded below by $1/2$ in the limit as δ goes to 1. Hence, this deviation is profitable for all sufficiently high discount factors, a contradiction.

2.2 No Holdup in a Market with Entry

Now consider the 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 to either *invest* or to *not invest*, and their investments determine the surplus of each match as specified in (1). Each agent that invests has to pay the investment cost c in the period in which she enters the market.[6] As in the market without entry, when agents are sufficiently patient, *efficiency requires that everyone invests*.

[6]Hence, each agent chooses the investment profile that maximizes her payoffs conditional on entering the market, even if this happens with probability zero.

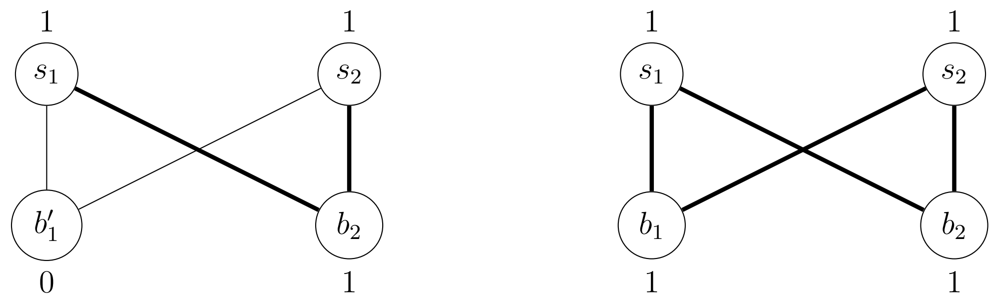

(i) Subgames with b'_1 active. (ii) Subgames with b'_1 inactive.

Figure 2: Surplus profiles that arise in the relevant subgames (up to isomorphism) of the market with entry when all the agents but b'_1 invest. Thin and thick links represent 1 and 2 units of surplus, respectively. The number associated with each agent is her Markov-perfect equilibrium limit (gross) payoff conditional on trading in a subgame with these surpluses.

Once the agents have sunk their investments, they bargain according to the following standard protocol (e.g. Talamàs 2018): 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. Hence, in each period, both the bargaining protocol and the set of surpluses that the active agents can create are exactly as in a subgame that starts in period $t = 1$ of the benchmark setting without entry described in subsection 2.1. For brevity, we relegate the details to subsection A.2.

We argue that, in stark contrast to the setting without entry, for all sufficiently high discount factors, *there exists an efficient type-symmetric Markov-perfect equilibrium*—that is, a subgame-perfect equilibrium in which everyone invests, all the buyers (sellers) follow the same strategy, and each agent's proposals condition only on the profile of investments made at $t = 0$ and the investments of the active agents. Moreover, for all high-enough discount factors, *every type-symmetric Markov-perfect equilibrium is efficient*.

First, we discuss why there exists a Markov-perfect equilibrium in which everyone invests. As before, on the equilibrium path, the game is fully symmetric, and surplus within each match is split equally in the limit as δ goes to 1. Let us argue that no agent can have a profitable investment deviation. Indeed, suppose that buyer b'_1 deviates and does not invest. Figure 2 depicts the available surpluses among the active agents in all the relevant subgames (up to isomorphism). The dynamic entry of agents into the market implies that—even when

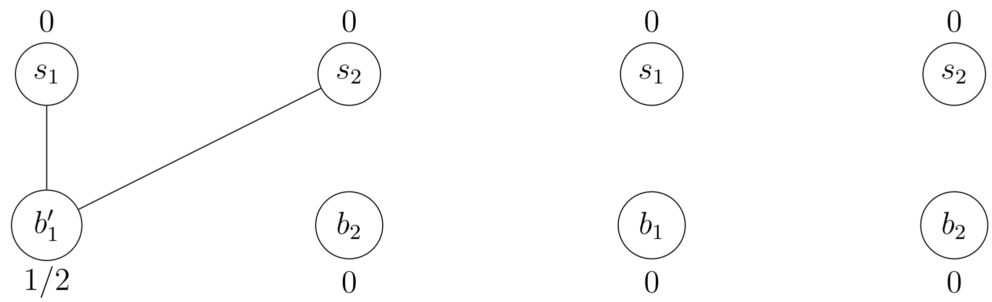

(i) Subgames with b_1' active. (ii) Subgames with b_1' inactive.

Figure 3: Surplus profiles that arise in the relevant subgames (up to isomorphism) of the market with entry when all the sellers invest, and no buyer other than b_1' invests. Thin and thick links represent 1 and 2 units of surplus, respectively. The number associated with each agent is her Markov-perfect equilibrium limit (gross) payoff conditional on trading in a subgame with these surpluses.

the deviator is active—at any point in time each of the sellers has access to a buyer with whom she generates 2 units of surplus, which implies, in turn, that she is able to obtain a payoff arbitrarily close to 1 as δ goes to 1. As a result, the deviator b_1' has to give away almost the whole unit of surplus in order to form a match. Hence, this deviation is unprofitable.

Second, we discuss why there does not exist an inefficient type-symmetric Markov-perfect equilibrium. Indeed, suppose for contradiction that there exists a type-symmetric Markov-perfect equilibrium in which no one invests (a similar argument shows that there exists no type-symmetric Markov-perfect equilibrium in which only the sellers or only the buyers invest). Consider the case in which a buyer b_1' deviates and invests. Figure 3 illustrates all the relevant subgames (up to isomorphism).When b_1' is active, the sellers have the option of trading with her, with whom they generate 1 unit of surplus—and hence might expect to obtain a fraction of these once-in-a-lifetime-high gains from trade. Crucially, however, there is only one deviator b_1', so the sellers are in competition to trade with her. In the limit as δ goes to 1, this competition becomes perfect and, as a result, the deviator b_1' obtains the full surplus created by her investment deviation. Hence, this deviation is profitable.

We now generalize this example to the case in which (i) there is an arbitrary number of types of agents (instead of only two), (ii) agents have access to an arbitrary number of investments, and different types of agents have access to different types of investments (instead of all types making the same binary investment decision), and (iii) agents exogenously enter the market over time (instead of the inflow of traders exactly matching the outflow).

3 Model

There is a finite set I of types agents, and a continuum of agents of each type. The type of an agent determines her investment opportunities and the 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 period $t = 0$

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 $\boldsymbol{x}_i$ and an agent of type $j \neq i$ with investment profile $\boldsymbol{x}_j$ produce $y(\boldsymbol{x}_i, \boldsymbol{x}_j)$ units of surplus when they match, and the costs of their investments are $c(\boldsymbol{x}_i)$ and $c(\boldsymbol{x}_j)$, respectively.

As described below, after deciding their investments, agents enter the market over time. For simplicity, we assume that each agent incurs her investment costs in the period in which she enters the market. We denote the surplus that an agent of type i with investment profile $\boldsymbol{x}_i$ generates in isolation (which can capture, for example, her outside options) by $y(\boldsymbol{x}_i, \boldsymbol{x}_i)$. Assumption 3.1 says that everyone has positive outside options independent of her investments, and that everyone has access to an investment profile whose associated outside option is higher than its cost. This assumption simplifies some of the statements that follow by ensuring that all the agents trade in equilibrium, and that agents' participation constraints do not bind.

Assumption 3.1. *For each type i and each investment profile $\boldsymbol{x}_i \in K_i$, $y(\boldsymbol{x}_i, \boldsymbol{x}_i) > 0$. Moreover, for each type i, there exists $\boldsymbol{x}_i \in K_i$ with $y(\boldsymbol{x}_i, \boldsymbol{x}_i) > c(\boldsymbol{x}_i)$.*

This formulation can capture substantial 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 $y(\boldsymbol{x}_i, \boldsymbol{x}_j)$ associated with any investment profile $(\boldsymbol{x}_i, \boldsymbol{x}_j) \in (K_{i'} \times K_{j'}) \cup (K_{i''} \times K_{j''})$ to be high relative to the associated investment costs, and the surplus $y(\boldsymbol{x}_i, \boldsymbol{x}_j)$ associated with any investment profile $(\boldsymbol{x}_i, \boldsymbol{x}_j) \in (K_{i''} \times K_{j'}) \cup (K_{i'} \times K_{j''})$ to be low relative to the associated investment costs.

3.2 Sequential Bargaining

Once everyone chooses her investment in period $t = 0$, bargaining occurs in discrete periods $t = 1, 2, \ldots$. For simplicity, we assume that, for each type i, there are $n_i \geq 2$ *bargaining slots* available. 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 $n := \sum_{i \in I} n_i$.

In each period $t = 1, 2, \ldots$, 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 output $y(\boldsymbol{x}_a, \boldsymbol{x}_b)$, where $\boldsymbol{x}_a$ and $\boldsymbol{x}_b$ 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 Sequential 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^i_s. Then, s' agents are drawn uniformly at random from those agents of type i that are yet to become active, and they are randomly assigned to different empty slots of type i.

We maintain Assumption 3.2 throughout, which requires that the market never dries up—in the sense that there are always at least two active agents of each type. This holds under fairly mild conditions on the stochastic inflow process $\{q^i_s\}_{i \in I, s \in \mathbb{N}}$. It holds, for example, if there is at least one active agent of each type in period $t = 1$, and $q^i_1(0) = 0$ for each $i \in I$. This assumption simplifies the statements that follow by guaranteeing that when all the agents of the same type make the same investment, the notion of subgame-perfect equilibrium pins down their payoffs (as we show in Appendix B), and that, essentially, a unilateral investment deviation from this strategy profile expands the non-deviators' bargaining opportunities.

Assumption 3.2. *There are always at least two active 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 $(h_t; i \to j; s)$ 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 $\sigma_i(h_t; j \to i; s)$. The strategy profile $\boldsymbol{\sigma}$ is a *type-symmetric Markov-perfect equilibrium* if (i) it induces a Nash equilibrium in every subgame, (ii) all the agents of any given type follow the same strategy, and (iii) each agent a's strategy conditions only on the investment profile chosen at $t = 0$, the set $\{y(\boldsymbol{x}_b, \boldsymbol{x}_c) \mid \text{agents } b, c \text{ active}\}$ of surpluses among the active agents, the set $\{y(\boldsymbol{x}_a, \boldsymbol{x}_b) \mid \text{agent } b \text{ active}\}$ of surpluses that she can create with the active agents, and—in the case of a response—on the going proposal.[7]

4 No Holdup in Equilibrium

Proposition B.2 in Appendix B (see also Talamàs 2018) shows that, for each investment profile $\boldsymbol{x} := (\boldsymbol{x}_i)_{i \in I}$ and each type i, there exists a value $V_i(\boldsymbol{x})$ such that, in every subgame-perfect equilibrium of the subgame that starts at $t = 1$ with the investment profile $\boldsymbol{x}$, $V_i(\boldsymbol{x})$ is the expected equilibrium (gross) payoff at the beginning of each period of each agent of type i.[8] We denote the limit of $V_i(\boldsymbol{x})$ as δ goes to 1 by $V_i^\star(\boldsymbol{x})$.

Theorem 4.1 shows that an investment profile $(\boldsymbol{x}_i)_{i \in I}$ can be implemented in a type-symmetric Markov-perfect equilibrium for all sufficiently high discount factors if and only if for each type i, $\boldsymbol{x}_i$ is the investment profile in K_i that maximizes the maximum—across all types j—of the net surplus that an agent of type i generates when she matches either with herself or with agents of a given type j, *taking as given the amount that agents of type j obtain in any given match.*

Theorem 4.1. *A type-symmetric Markov-perfect equilibrium with investment profile $\boldsymbol{x} := (\boldsymbol{x}_i)_{i \in I}$*

[7] Given that we focus on type-symmetric strategies and that agents are chosen uniformly at random to become active (so that the *distribution* of investment profiles of the agents that fill the bargaining slots of any given type is stationary) these capture all the payoff-relevant states of the world in this game.

[8] In particular, conditional on Assumption 3.2 holding, agents' expected equilibrium payoffs do not depend on which (or how many) bargaining slots are empty.

exists for all sufficiently high discount factors if and only if

$$\text{(2)} \qquad \boldsymbol{x}_i \in \mathop{\mathit{argmax}}_{\boldsymbol{z}_i \in K_i} \left[\max\left(y(\boldsymbol{z}_i, \boldsymbol{z}_i), \max_{j \in I} \left[y(\boldsymbol{z}_i, \boldsymbol{x}_j) - V_j^\star(\boldsymbol{x}) \right] \right) - c(\boldsymbol{z}_i) \right] \textit{ for each } i \in I.$$

Proof. Necessity: Fix a type-symmetric Markov-perfect equilibrium $\boldsymbol{\sigma}$ with investment profile $(\boldsymbol{x}_i)_{i \in I}$. 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 $\boldsymbol{\sigma}$ is subgame perfect and that, by Assumption 3.1, each agent can create a strictly positive surplus by matching to herself, each agent gets—when she is the proposer—the maximum amount that she can obtain while—in case she does not match with herself—leaving the receiver indifferent between accepting and rejecting. Hence,

$$v_i = \max\left(y(\boldsymbol{x}_i, \boldsymbol{x}_i), \max_{j \in I} \left[y(\boldsymbol{x}_i, \boldsymbol{x}_j) - w_j \right] \right).$$

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

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

Hence, it is enough to show that, for any investment deviation from the equilibrium $\boldsymbol{\sigma}$ by any 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 $V_i^\star(\boldsymbol{x})$, when δ is sufficiently close to 1, each agent a must then choose her investment $\boldsymbol{z}_a$ to maximize

$$\max\left(y(\boldsymbol{z}_a, \boldsymbol{z}_a), \max_{j \in I} \left[y(\boldsymbol{z}_a, \boldsymbol{x}_j) - V_j^\star(\boldsymbol{x}) \right] \right) - c(\boldsymbol{z}_a).$$

Suppose that an agent d of type k deviates from $\boldsymbol{\sigma}$ by investing $\boldsymbol{x}_d \neq \boldsymbol{x}_k$. Given Assumption 3.2, neither the set $\{y(\boldsymbol{x}_a, \boldsymbol{x}_b) \mid \text{agents } a, b \text{ active}\}$ of surpluses among the active agents nor the set $\{y(\boldsymbol{x}_d, \boldsymbol{x}_b) \mid \text{agent } b \text{ active}\}$ of surpluses that the deviator d can generate with the active agents change while the deviator d is active. Hence, given that $\boldsymbol{\sigma}$ is Markov-perfect, (i) for each agent a we can let $\hat{w}_a$ be her expected equilibrium payoff when rejecting an offer while d is active and, (ii) when d is the proposer, she offers $\hat{w}_a$ to some agent a, who accepts with probability one.[9]

[9] Note that agent d deviates at the investment stage only, so $\boldsymbol{\sigma}$ still governs her bargaining strategy. Given that $\boldsymbol{\sigma}$ 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.

Let $i \in I$ and let $a \neq d$ be an agent of type i with the property that there exists another agent $c \neq a$ with whom the deviator trades with positive probability in equilibrium (Assumption 3.2 ensures that we can find such an agent). We argue that $\hat{w}_a$ converges to $w_i^\star$ in the limit as δ goes to 1. Since $\boldsymbol{\sigma}$ is type symmetric and Markov perfect, this implies that, for each agent b of type i, $\hat{w}_b$ also converges to $w_i^\star$ in the limit as δ goes to 1 (otherwise, Markov perfection would imply that some agents reject offers that other agents of the same type accept).

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 that—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. Suppose for contradiction that there exists a sequence $\mathcal{D}$ of discount factors converging to 1 such that $\hat{w}_a > w_i + \epsilon$ for all sufficiently high discount factors in $\mathcal{D}$. Given that, as we have just argued, for each $j \in I$ 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 in $\mathcal{D}$. 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

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

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

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

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 C.1 (in Appendix C) 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 each investment profile $\boldsymbol{w}$, we can pick an equilibrium $\boldsymbol{\sigma}(\boldsymbol{w})$ of the subgame that starts at $t = 1$. Given any investment profile $(\boldsymbol{x}_i)_{i \in I}$, define a strategy profile as follows: All the agents of type i invest $\boldsymbol{x}_i$, and each agent's bargaining strategy given any investment profile $\boldsymbol{w}$ is as specified by $\boldsymbol{\sigma}(\boldsymbol{w})$. By the *one-stage deviation principle* (e.g. Fudenberg and Tirole 1991, Theorem 4.2), 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 (2). □

5 Discussion

Theorem 4.1 characterizes the investment profiles that can be implemented in a type-symmetric Markov-perfect equilibrium for all sufficiently high discount factors. These are exactly those profiles with the property that each type's investment maximizes the maximum net surplus that she can generate, *taking as given all the other agents' payoffs*.

Let $M_i(\boldsymbol{x})$ be the set of types j with the property that, in a type-symmetric Markov-perfect equilibrium with investment profile $\boldsymbol{x}$, for all sufficiently high discount factors, there are agents of type i who match with agents of type j. In the limit as the discount factor goes to 1, all the agents of a given type receive the same amount independently of which match they form, so we have that $j \in M_i(\boldsymbol{x})$ if and only if $j \in \operatorname{argmax}_{k\in I}\, y(\boldsymbol{x}_i, \boldsymbol{x}_k) - V_k^\star(\boldsymbol{x})$. Hence, for each equilibrium profile $(\boldsymbol{x}_i)_{i\in I}$, we have that

$$\boldsymbol{x}_i \in \operatorname*{argmax}_{\boldsymbol{z}_i\in K_i} \left[y(\boldsymbol{z}_i, \boldsymbol{x}_j) - c(\boldsymbol{z}_i)\right] \text{ for all } j \in M_i(\boldsymbol{x}). \tag{3}$$

When $M_i(\boldsymbol{x})$ contains more than one element, in general there need not exist an investment profile $\boldsymbol{x}_i \in K_i$ that satisfies the system (3). We now describe fairly mild conditions on the investment technology that ensure that *types match one-to-one* in equilibrium—that is, that for each type i and each equilibrium investment profile $(\boldsymbol{x}_i)_{i\in I}$, $M_i(\boldsymbol{x})$ is a singleton. Under these conditions, a non-zero constrained-efficient investment profile (Definition 5.1) exists if and only if there is one pair of types i and j and corresponding investment profiles $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$, with the property that

$$\boldsymbol{x}_i \in \operatorname*{argmax}_{\boldsymbol{z}_i\in K_i} \left[y(\boldsymbol{z}_i, \boldsymbol{x}_j) - c(\boldsymbol{z}_i)\right] \text{ and } \boldsymbol{x}_j \in \operatorname*{argmax}_{\boldsymbol{z}_j\in K_j} \left[y(\boldsymbol{x}_i, \boldsymbol{z}_j) - c(\boldsymbol{z}_j)\right].$$

That is, the investment profile $\boldsymbol{x}_i$ maximizes the joint surplus of a match between an agent of type i and an agent of type j conditional on j's investment profile $\boldsymbol{x}_j$, and vice versa. This is satisfied, for instance, in the example discussed in section 2, where invest maximizes the joint surplus of a match independently of whether the other party in the match invests or not.

Let the investment set K_i of each agent of type i be the union of $\{K_{ij}\}_{j\in I}$, where, for every $j \in I$, $K_{ij} \subset \mathbb{R}_{\geq 0}$ is a finite set, with $0 \in K_{ij}$. We refer to an agent's investment in K_{ij} as her *type-j-specific investment*. Without loss of generality, we assume that the cost function c is strictly increasing, we normalize $c(\mathbf{0}) = 0$, and we assume that the sets in $\{K_i\}_{i\in I}$ only have $\mathbf{0}$ in common.

Assumption 5.1 requires that agents' investments possibilities be sufficiently rich—in the sense that making a type-i-specific investment is never the most cost-effective way to increase one's gains from trade with a different type j. This rules out situations in which agents make certain type-specific investments only to increase their gains from trade with other types.

Assumption 5.1 (Rich investment opportunities). *Let $i \neq j \in I$, $\boldsymbol{x}_i \in K_i$ and $\boldsymbol{x}_j \in K_j$ with $x_{ji} > 0$. If $x_{ik} > 0$ for $k \neq j$, then there exists $\boldsymbol{z}_i \in K_i$ with $c(\boldsymbol{z}_i) \leq c(\boldsymbol{x}_i)$, $z_{ik} = 0$ and $y(\boldsymbol{z}_i, \boldsymbol{x}_j) > y(\boldsymbol{x}_i, \boldsymbol{x}_j)$. Similarly, if $x_{ij} > 0$, then there exists $\boldsymbol{z}_i \in K_i$ with $c(\boldsymbol{z}_i) \leq c(\boldsymbol{x}_i)$, $z_{ij} = 0$ and $y(\boldsymbol{z}_i, \boldsymbol{z}_i) > y(\boldsymbol{x}_i, \boldsymbol{x}_i)$.*

Assumption 5.2 requires gains from trade to only emerge when the parties involved have made type-specific investments: If an agent a of type i and an agent b of type j are to have gains from trade, it is necessary that a makes a non-zero type-j-specific investment and that b makes a non-zero type-i-specific investment.

Assumption 5.2 (Gains from trade require type-specific investments). *Let $i \neq j \in I$, $\boldsymbol{x}_i \in K_i$ and $\boldsymbol{x}_j \in I_j$. We have $y(\boldsymbol{x}_i, \boldsymbol{x}_j) > 0$ if and only if $x_{ij}x_{ji} > 0$.*

Intuitively, Assumption 5.1 implies that each agent chooses at most one type-specific investment. Hence, given that, by Assumption 5.2, gains from trade require type-specific investments, under these two assumptions, types necessarily match one-to-one. As a result, under these two conditions, Theorem 4.1 implies that an investment profile can be implemented in a type-symmetric Markov-perfect equilibrium for all sufficiently high discount factors if and only if it is constrained-efficient and individually rational, defined as follows.

Definition 5.1. Under Assumption 5.1 and Assumption 5.2, we say that the investment profile $(\boldsymbol{x}_i)_{i \in I}$ is *constrained efficient* if, for each type i, $x_{ij} > 0$ for at most one $j \neq i$ and, for each pair (i, j) with $x_{ij} > 0$, we have that $x_{ji} > 0$ and that

$$\boldsymbol{x}_i \in \underset{\boldsymbol{z}_i \in I_i}{\operatorname{argmax}}\, y(\boldsymbol{z}_i, \boldsymbol{x}_j) - c(\boldsymbol{z}_i),$$

Moreover, we say that the constrained-efficient investment profile $(\boldsymbol{x}_i)_{i \in I}$ is *individually rational* if $x_{ij} > 0$ implies that

$$y(\boldsymbol{x}_i, \boldsymbol{x}_j) - c(\boldsymbol{x}_i) \geq \max_{\boldsymbol{z}_i \in K_i} y(\boldsymbol{z}_i, \boldsymbol{z}_i) - c(\boldsymbol{z}_i).$$

Appendices

A Details of section 2

A.1 Market with no entry

A.1.1 Subgame-perfect equilibrium after one match has formed

When there are only two unmatched agents, who can generate y units of surplus, there is a unique subgame-perfect-equilibrium (e.g. Ståhl 1972, Rubinstein 1982, Binmore 1987): Each agent accepts an offer if and only it gives her the amount $w(y)$ that satisfies

$$w(y) = \delta \left(\frac{1}{4}(y - w(y)) + \frac{3}{4} w(y) \right), \text{ or, rearranging, } w(y) = \frac{\delta}{4 - 2\delta} y \to \frac{y}{2}, \tag{4}$$

and the proposer offers $w(y)$ to the other unmatched agent, who accepts with probability one.

A.1.2 Non-existence of efficient equilibrium

Let us show that, for all sufficiently high discount factors, there does not exist an efficient type-symmetric Markov-perfect equilibrium. For this, it is enough to show that, when everyone but b_1 invests, b_1's type-symmetric Markov-perfect equilibrium limit (gross) payoff is bounded below by $1/2$. For each agent a, let w_a be her expected equilibrium (gross) payoff in a period in which she is active but she is not the proposer, and no match has formed.

Let $\epsilon > 0$. For contradiction, suppose that there exists a sequence $\mathcal{D}$ of discount factors converging to 1 such that, for all sufficiently high discount factors in $\mathcal{D}$, $w_{b_1} \leq 1/2 - \epsilon$. This implies that, for all sufficiently high discount factors in $\mathcal{D}$, before any match has formed, the sellers trade with b_2 with a vanishingly small probability (otherwise, b_1 could just wait for them to match, which would give her a limit payoff of $1/2$). Hence, letting $w_s := w_{s_1} = w_{s_2}$ (which holds because strategies are type symmetric), when a seller is the proposer, she weakly prefers to delay than to make an acceptable offer to b_2. In particular,

$$2 \leq w_{b_2} + w_s. \tag{5}$$

Letting $\pi \geq 0$ be the probability that b_1 forms a match when b_2 is not the proposer,

$$w_{b_2} \leq \delta \left[\frac{1}{4} w_{b_2} + \frac{3}{4} \left(\pi w(2) + (1 - \pi) w_{b_2} \right) \right],$$

which implies that $w_{b_2} < w(2)$, which combined with Equation 5 gives $w_s > w(2)$. That is, each seller strictly prefers to be in a subgame where no match has formed, which implies that

(6) $$w_s < \delta\left[\frac{1}{4}(1 - w_{b_1}) + \frac{3}{4}w_s\right].$$

But, by our assumption that, for all sufficiently high discount factors in $\mathcal{D}$, $w_{b_1} \leq 1/2 - \epsilon$, and the fact that $w(1)$ converges to $1/2$ as δ goes to 1, we have that $w_{b_1} < w(1)$ for all sufficiently high discount factors in $\mathcal{D}$, so b_1 strictly prefers to be in a subgame where a match has formed, which implies that

(7) $$w_{b_1} > \delta\left[\frac{1}{4}(1 - w_s) + \frac{3}{4}w_{b_1}\right].$$

Equation 6 and Equation 7 together imply that $w_{b_1} > w(1)$ for all sufficiently high discount factors, a contradiction.

A.1.3 Equilibrium when everyone but one agent invests

Consider the case in which everyone but buyer b_1 invests. Let us check that, for all sufficiently high discount factors, there exists a unique Markov-perfect equilibrium (of the subgame that starts at $t = 1$) in which (i) b_1 delays (with some probability $\pi > 0$) and (ii) b_2 matches with probability one unless b_1 is the proposer, and that, in this equilibrium, b_1's limit payoff is $1/2$. Buyer b_2 must be making offers to the seller with the lowest continuation value, so she must offer to each seller with equal probability (otherwise, the seller who is more likely to receive offers has a strictly higher continuation value). This implies, in turn, that s_1's continuation value, w_{s_1}, is equal to s_2's continuation value, w_{s_2}, so we can let $w_s := w_{s_1} = w_{s_2}$. Moreover, b_1 must be indifferent between *delaying* and *not delaying*. Hence, we have that $w_{b_1} = \delta\left(\frac{1}{4}w_{b_1} + \frac{3}{4}w(1)\right)$, or, $w_{b_1} = w(1)\frac{3\delta}{4-\delta}$ and that

$$\begin{aligned}
w_{b_2} &= \delta\left(\frac{1}{4}\underbrace{(2 - w_s)}_{\text{proposer is } b_2} + \frac{1}{4}\underbrace{(\pi w_{b_2} + (1-\pi)w(2))}_{\text{proposer is } b_1'} + \frac{1}{2}\underbrace{w_{b_2}}_{\text{proposer is } s_1 \text{ or } s_2}\right) \\
w_s &= \delta\left(\frac{1}{4}\underbrace{(2 - w_{b_2})}_{\text{proposer is } s_1} + \frac{1}{4}\underbrace{\left(\frac{1}{2}w_s + \frac{1}{2}w(1)\right)}_{\text{proposer is } b_2} + \frac{1}{4}\underbrace{\left(\pi w_s + \frac{1-\pi}{2}w_s + \frac{1-\pi}{2}w(2)\right)}_{\text{proposer is } b_1'} + \frac{1}{4}\underbrace{w(1)}_{\text{proposer is } s_2}\right) \\
1 &= w_{b_1'} + w_s.
\end{aligned}$$

It can be easily checked that this system has a unique solution for all δ sufficiently large, which satisfies

$$w_s \to \frac{1}{2},\ w_{b_1'} \to \frac{1}{2},\ w_{b_2} \to \frac{3}{2},\ \text{and } \pi \to 1.$$

Moreover, for all δ sufficiently large, $1 - w_{b_1'} < 2 - w_{b_2}$ and $2 < w_{b_2} + w_s$, so that each seller and b_2 indeed want to make acceptable offers to each other.

A.2 Market with Sequential Entry

A.2.1 Existence of an efficient type-symmetric Markov-perfect equilibrium

In Appendix C, 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, by the *one-stage deviation principle*, it is enough to describe strategies for the subgames that start with (i) everyone that is yet to trade (active or inactive) having invested, and (ii) everyone but one active agent having invested, and to show that, under any equilibrium featuring these strategies in these subgames, and specifying that everyone invests in $t = 0$, any unilateral investment deviation from this strategy profile is unprofitable.

First, let us describe Markov-perfect strategies for the subgames where everyone that is yet to trade 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,

$$w = \delta \left(\frac{1}{4} \underbrace{(2 - w)}_{\text{proposer's payoff}} + \frac{3}{4} \underbrace{w}_{\text{non-proposer's payoff}} \right), \text{ that is, } w = \frac{\delta}{2 - \delta} \to 1. \tag{8}$$

Second, let us describe Markov-perfect strategies for every subgame in which all but one agent (buyer b_1' say), who is active, has invested (assuming that the strategy profile will be as described above once the deviator leaves): Each non-deviator (i) offers w (as defined in Equation 8) 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 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 b_1' 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,

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

Hence, this is indeed an equilibrium, since $2 - w > 1 - w'$ implies that the best the non-deviators can do is to obtain $2 - w$ when they are the proposers. We conclude that the deviator saves c but loses 1 in the limit as δ goes to 1, so, for all sufficiently high discount factors, her deviation is not profitable.

A.2.2 Non-Existence of a non-efficient type-symmetric Markov-perfect equilibrium

Suppose for contradiction that there exists a type-symmetric Markov-perfect equilibrium $\boldsymbol{\sigma}$ in which only the sellers invest. We show that a buyer b_1' can profitably deviate by investing. The same argument shows that there exists no type-symmetric Markov-perfect equilibrium in which only buyers invest, or in which neither buyers nor sellers invest.

On the equilibrium path, each agent's expected payoff when she rejects an offer satisfies

$$(9) \qquad w = \delta \left(\frac{1}{4} \underbrace{(1-w)}_{\text{proposer's payoff}} + \frac{3}{4} \underbrace{w}_{\text{non-proposer's payoff}} \right), \text{ that is, } w = \frac{\delta}{4 - 2\delta} \to \frac{1}{2}.$$

Suppose that buyer b_1' deviates and invests, and consider a subgame in which b_1' is active. From the point of view of the deviator, the environment is stationary, so when she is the proposer, she makes offers that leave the receiver indifferent between accepting and rejecting, and which are accepted with probability one.[10]

Let s be a seller such that there exists another 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 $\boldsymbol{\sigma}$ 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 this deviation is profitable.

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 δ. The same 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 $\epsilon > 0$ and a sequence $\mathcal{D}$ of discount factors converging to 1

[10]Note that we are considering an *investment* deviation from $\boldsymbol{\sigma}$, so it still governs the deviator's bargaining strategy.

such that $\hat{w}_s \geq 1/2 + \epsilon$ for all sufficiently high discount factors in $\mathcal{D}$. 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 discount factors 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

$$\hat{w}_s = \delta \left[\frac{1}{4}(2 - \hat{w}_{b_1'}) + \frac{3}{4}\left(\pi w + (1-\pi)\hat{w}_s\right) \right] \text{ and } \hat{w}_{b_1'} = \delta \left[\frac{1}{4}(2 - \hat{w}_s) + \frac{3}{4}\hat{w}_{b_1'} \right],$$

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

B Uniqueness of perfect equilibrium payoffs

Proposition B.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 2018, where it is assumed that exactly one agent of each type is active in the market at each point in time. Proposition B.2 holds under the following assumption, which is weaker than Assumption 3.2.

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

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

The proof of Proposition B.2 is identical to the corresponding result in Talamàs 2018, 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). For completeness, we provide this proof here. Proposition B.2 follows from Proposition B.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 B.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 B.3 using the techniques developed in Manea 2017.

Proposition B.3. *Fix an investment profile $(\boldsymbol{x}_i)_{i\in I}$. For every type i, there exists $w_i \geq 0$ such that, in every game in which Assumption B.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 $y(\boldsymbol{x}_i, \boldsymbol{x}_j)$ 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 $0, 1, \ldots$, 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

$$m_i^0 := 0. \tag{10}$$

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

$$M_i^0 := \max_j \{s_{ij}\}. \tag{11}$$

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 $s = k + 1$.

<u>Step $k + 1$.</u>

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:

$$(12) \qquad m_i^{k+1} := \frac{1}{n} \max \left(\max_{j \in N} (s_{ij} - \delta M_j^k), \delta m_i^k \right) + \frac{n-1}{n} \delta m_i^k$$

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 12 can be written as

$$(13) \qquad m_i^{k+1} = \begin{cases} \delta m_i^k & \text{if } \max_{j \in N} (s_{ij} - \delta M_j^k) \leq \delta m_i^{k+t+1-t'} \\ \frac{1}{n} \max_{j \in N} (s_{ij} - \delta M_j^k) + \frac{n-1}{n} \delta m_i^k & \text{otherwise} \end{cases}$$

and an analogous equation can be used to expand the term m_i^k in Equation 13, 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 $\epsilon \to 0$. 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

$$(14) \qquad M_i^{k+1} := \frac{1}{n} \max \left(\max_{j \in N} (s_{ij} - \delta m_j^k), \delta M_i^k \right) + \frac{n-1}{n} \delta M_i^k$$

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 o 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 $\max\limits_{j\in N}(s_{ij}) \geq M_i^k \geq m_i^k \geq 0$ 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 $M_i^0 := \max\limits_j \{s_{i,j}\}$, and that Equation 12 and Equation 14 imply that $m_i^1 \geq 0$ and $M_i^1 \leq \max\limits_j \{s_{i,j}\}$, so $m_i^1 \geq m_i^0$ and $M_i^1 \leq M_i^1$. 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 12 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 14 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 $\max\limits_{j\in N}(s_{ij}) = M_i^0 > m_i^0 = 0$.

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 $\max\limits_{i\in N}(M_i^k - m_i^k)$. 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,

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

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

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

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

C Existence of a type-symmetric Markov-perfect equilibrium

This is analogous to the Markov-perfect equilibrium existence proof in Elliott and Nava forthcoming.

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

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 $V : \mathcal{K} \to \mathbb{R}^m$, 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 $k \in K \backslash \{i\}$ is

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

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

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

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

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

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 $V : \mathcal{K} \to \mathbb{R}^m$, consider the correspondence $F : \mathcal{V} \to \mathcal{V}$ defined by

$$\text{(15)} \qquad F(V)(K) = \frac{1}{n} \sum_{i \in K} f^{i,K}(V), \text{ for all value functions } V \text{ and all } K \in \mathcal{K}.$$

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. $\square$

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.

BECKER, G. S. (1975): "Front matter, human capital: a theoretical and empirical analysis, with special reference to education," in *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education, Second Edition*, NBER, 22–0.

BINMORE, K. G. (1987): *Nash bargaining theory II*, In Binmore and Dasgupta (1987).

BINMORE, K. G. AND P. DASGUPTA (1987): *The Economics of Bargaining*, Oxford: Blackwell.

CHE, Y.-K. AND J. SÁKOVICS (2004): "A dynamic theory of holdup," *Econometrica*, 72, 1063–1103.

COLE, H., G. MAILATH, AND A. POSTLEWAITE (2001a): "Efficient Non-Contractible Investments in Finite Economies," *The BE Journal of Theoretical Economics*, 1, 1–34.

COLE, H. L., G. J. MAILATH, AND A. POSTLEWAITE (2001b): "Efficient non-contractible investments in large economies," *Journal of Economic Theory*, 101, 333–373.

DEMANGE, G. (1982): "Strategyproofness in the assignment market game," *Laboratoire d'Économétrie de l'École Polytechnique*.

ELLIOTT, M. (2015): "Inefficiencies in networked markets," *American Economic Journal: Microeconomics*, 7, 43–82.

ELLIOTT, M. AND F. NAVA (forthcoming): "Decentralized bargaining in matching markets: Efficient stationary equilibria and the core," *Theoretical Economics*.

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," *Cambridge, Massachusetts*, 393, 80.

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.

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.

KRANTON, R. E. AND D. F. MINEHART (2001): "A theory of buyer-seller networks," *American Economic Review*, 91, 485–508.

LEONARD, H. B. (1983): "Elicitation of honest preferences for the assignment of individuals to positions," *Journal of Political Economy*, 91, 461–479.

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. (2017): "Bargaining in dynamic markets," *Games and Economic Behavior*, 104, 59–77.

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.

RUBINSTEIN, A. (1982): "Perfect equilibrium in a bargaining model," *Econometrica*, 97–109.

STÅHL, I. (1972): "Bargaining theory," *Economic Research Institute at the Stockholm School of Economics (EFI)*.

TALAMÀS, E. (2018): "Price dispersion in stationary networked markets," *Mimeo*.

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