

Nash Bargaining with Endogenous Outside Options*

[most recent version here]

Eduard Talamàs[†]

Abstract

Outside options shape bargaining outcomes, but understanding how they are determined is often challenging, because one's outside options depend on others' outside options, which depend, in turn, on others' outside options, and so on. In this paper, I describe a non-cooperative theory of coalition formation that shows how the classical Nash bargaining solution uniquely pins down *both* the sharing rule and the relevant outside options in each coalition. This provides a tractable framework to investigate how different economic shocks propagate via outside options. In two-sided pairwise matching markets where agents are vertically differentiated by their skills, shocks propagate from the high to the low skill, *but not vice versa*. Positive assortative matching necessarily arises if and only if skills are complementary. In this case, shocks propagate *in blocks*, in the sense that when a shock propagates from one agent to another one, it also propagates to everyone whose skill is in between.

1 Introduction

The Nash bargaining solution is a central concept in economics.¹ It provides a sharing rule in any given coalition as a function of its members' outside options. Its clean axiomatic foun-

*Date printed: November 19, 2018.

[†]University of Pennsylvania and Harvard University. The guidance of Benjamin Golub throughout the process of conducting and presenting this research has been essential. I am grateful to Matthew Elliott, Jerry Green and Rakesh Vohra for extensive discussions and advice, as well as to numerous conference and seminar participants for useful feedback. This work was supported by the Warren Center for Network & Data Sciences, and the Rockefeller Foundation (#2017PRE301). All errors are my own.

¹Nash bargaining is widely used in virtually every branch of economics: See for example Grout (1984), Grossman and Hart (1986), Carraro and Siniscalco (1993), Mortensen and Pissarides (1994), Lundberg and Pollak (1996), Bagwell, Staiger, and Yurukoglu (2018), Manea (2018) and Ho and Lee (forthcoming).

dations (Nash 1950) and close connections to non-cooperative bargaining (e.g., Binmore, Rubinstein, and Wolinsky 1986) make it theoretically appealing, and its simple functional form makes it convenient in applications. In many settings of interest, however, agents simultaneously bargain over *both* which coalitions to form (e.g., which firms employ which workers, which entrepreneurs become partners, which businesses form strategic alliances, etc.) and how to share the resulting gains from trade (e.g., wages, equity shares, etc.), and using the Nash bargaining solution in these settings requires a theory of how the relevant outside options are determined. For example, the outside options of a job candidate when bargaining with a potential employer are often determined by the bargaining outcomes with alternative employers, which depend, in turn, on these alternative employers' outside options, and so on. Hence, understanding the resulting outcomes requires a theory that somehow cuts this outside option Gordian knot.

In this paper, I describe a non-cooperative theory of coalition formation that uniquely pins down both who matches with whom and how the resulting gains from trade are shared in stationary settings. The key observation that cuts the outside option Gordian knot is that there always exists at least one coalition that is sufficiently productive so that—when bargaining to form this coalition—none of its members has a credible outside option. This allows a recursive characterization of the relevant outside options in each potential coalition.

The set of coalitions that form in equilibrium has a nice structure, which makes the resulting theory of coalition formation especially tractable. In particular, the coalitions that form in equilibrium can be organized into tiers, in such a way that the equilibrium sharing rule in each coalition converges—as the bargaining frictions vanish—to the Nash bargaining solution, with the relevant outside options determined by the Nash bargaining solution in higher-tiered coalitions. This implies that (small) changes in market fundamentals propagate—via outside options—from higher to lower tiers, but not vice versa.

More generally, the theory that I describe in this paper overcomes a common indeterminacy problem in standard matching models, and this allows it to provide novel comparative statics in rich matching environments.² The predictions of this theory are broadly consistent with the view that bargaining plays a more prominent role in the determination of high-skill than low-skill wages.³ For example, it predicts that—in two-sided pairwise matching

²As an example of this indeterminacy problem, the classical assignment game of Shapley and Shubik (1971) typically has a “large” core, and the economic properties this game crucially depend upon which among the many possible points in its core is selected (see for example Kranton and Minehart 2001 and Elliott 2015).

³Hall and Krueger (2012) and Brenzel, Gartner, and Schnabel (2014) document a positive correlation between education and wage bargaining in the United States and Germany, respectively.




Figure 1: Illustration of the outside option principle. A graduate and an employer can generate one dollar by matching. The graduate can receive a wage of $w < 1$ elsewhere, and the employer can hire an equally valid candidate at wage $w' > w$. The employer hires the graduate at wage $1/2$ unless $w > 1/2$ or $w' < 1/2$.

markets where agents are vertically differentiated by their skills—bargaining outcomes are determined from the top down. In particular, an increase in the skill of an agent can affect the payoffs of agents whose skills are lower than hers, *but it does not affect the payoff of any agent whose skill is higher than hers*. Intuitively, this is because—in any bargaining encounter—the option of bargaining with a lower-skill agent is not credible and, as a result, it does not affect the equilibrium outcome. Interestingly, however, the analogous reasoning shows that—when the source of heterogeneity is risk aversion instead of skill—bargaining outcomes are determined from the most risk averse agents down. Hence, heterogeneities in skills and heterogeneities in preferences have qualitatively different implications on the way in which shocks propagate in matching markets: When skill is the only source of heterogeneity, shocks propagate *from the agents with the highest payoffs down*. In contrast, when risk aversion is the only source of heterogeneity, shocks propagate *from the agents with the lowest payoffs up*.

This paper is related to [Binmore, Rubinstein, and Wolinsky \(1986\)](#), who describe a non-cooperative bargaining model *in a fixed coalition* to investigate how *exogenous* outside options enter the Nash bargaining solution.⁴ The unique subgame-perfect equilibrium of their game predicts that—as bargaining frictions vanish—the surplus in the coalition of interest is shared according to the Nash bargaining solution, with the *threat points* corresponding to the utilities that the agents get in autarky, and the *outside options entering as lower bounds on the payoffs*. This is the “outside option principle.” (e.g., [Sutton 1986](#)). For example, consider

⁴In many applications, there are different sensible alternatives for both what the relevant outside options are and how they enter the Nash bargaining solution—and different alternatives have qualitatively different implications. For example, the extent to which unemployment is a relevant outside option in wage bargaining determines the effects of unemployment insurance on the labor market—e.g., [Pissarides \(2000\)](#), [Krusell et al. \(2010\)](#), [Hagedorn, Karahan, Manovskii, and Mitman \(2013\)](#) and [Chodorow-Reich, Coglianesi, and Karabarbounis \(2018\)](#)—and the ability of macroeconomic models to generate realistic employment fluctuations—e.g., [Shimer \(2005\)](#), [Hall and Milgrom \(2008\)](#), [Sorkin \(2015\)](#), [Chodorow-Reich and Karabarbounis \(2016\)](#), [Hall \(2017\)](#) and [Ljungqvist and Sargent \(2017\)](#).

the situation described in [Figure 1](#), where a recent graduate and an employer (both risk neutral) can generate 1 dollar by matching. Suppose that (i) the graduate can sell her labor elsewhere at wage $w < 1$, (ii) the employer can hire an equally valuable recent graduate at wage $w' > w$, and (iii) neither the employer nor the graduate in autarky generate any value. In this case, the outside option principle suggests that the employer hires the graduate at wage $1/2$ (as specified by the Nash bargaining solution with the threat point determined by autarky), unless $w > 1/2$ or $w' < 1/2$, in which case it suggests that the employer hires the graduate at wage w or w' , respectively. Intuitively, an agent's outside option only affects her bargaining position if it is credible, in the sense that her outside option is better than what the Nash bargaining solution would otherwise give her.⁵

Crucially, however, the outside option principle is silent about how the relevant outside options in each coalition are determined. For instance, in the example just described, the wages w and w' at which the graduate and the employer, respectively, can match elsewhere are taken as given. But, in many cases, these wages are themselves the result of bargaining with third parties. From this perspective, the contribution of this paper is to describe a non-cooperative theory that shows not only how outside options enter the Nash bargaining solution, but also *how the Nash bargaining solution pins down the relevant outside options in each coalition*. By endogenizing the relevant outside options in each coalition, the resulting theory provides a tractable framework—grounded on classical bargaining theory—to trace out the general equilibrium effects of different economic shocks.

In the model, different types of agents enter a market over time in such a way that there are always agents of each type looking to form a coalition. The model is intended to capture the predominant economic forces in large markets with dynamic entry, where the relevant matching opportunities are roughly constant over time.⁶ Examples include relatively thick labor markets where workers and firms arrive over time in search of profitable (potentially many-to-many) matches, and innovation hubs where startups and entrepreneurs cluster to form (potentially multilateral) strategic alliances. The agents in the market bargain according to a standard protocol (in the spirit of the canonical alternating-offers model of [Rubinstein 1982](#)) over both which coalitions to form and how to share the resulting gains from

⁵[Binmore, Shaked, and Sutton \(1989\)](#) provide experimental evidence that is consistent with the outside option principle. More recently, [Jäger, Schoefer, Young, and Zweimüller \(2018\)](#) find that real-world wages are insensitive to sharp increases in unemployment insurance benefits, which suggests that unemployment is not a credible outside option in wage bargaining.

⁶In particular, I assume that the surplus of each match is independent of which other matches have formed in the past or will form in the future. The approach is similar to the one in [Rubinstein and Wolinsky \(1985\)](#) and the subsequent literature studying non-cooperative bargaining in stationary markets.




Figure 2: An example of the negotiation dynamics between MBA graduates, Wall Street firms and Main Street firms that might lead to the MBA’s outside options being determined in a circular way. When bargaining on Wall Street, MBA’s demand to obtain at least w because they can earn w on Main Street. But the only reason that Main Street pays w is that Wall Street does.

trade. The bargaining friction that incentivizes them to reach agreements is their fear that an exogenous reason will prevent them from matching in the future. As a result, as in the classical bargaining framework of Nash (1950), risk preferences are essential drivers of the bargaining outcomes.⁷

I show that the model admits an essentially unique stationary subgame-perfect equilibrium, and I characterize which coalitions form and how are the resulting gains from trade shared in this equilibrium. In the limit as the bargaining friction vanishes, the equilibrium sharing rule in each coalition is the one prescribed by the Nash bargaining solution, with the relevant outside options entering as prescribed by the outside option principle, and determined as follows:

Each agent’s outside option in any given coalition is her maximum Nash bargaining share—across all the other coalitions—while honoring *the others’* outside options.

The main result of this paper is that there is a unique outside option profile that satisfies this property, and that the non-cooperative bargaining model suggests the resulting outcome as a natural point for the agents to settle on when bargaining in a decentralized way. Roughly speaking, the strategic forces in the non-cooperative model demand that each agent be able to justify her outside option in each coalition as resulting from the Nash bargaining solution in another coalition *without appealing to her own outside option there*. Intuitively, this

⁷Analogous results can be derived if one assumes that no agent is ever exogenously forced out of the market but that—instead—the agents are impatient. The Nash bargaining solution then has to be appropriately constructed from agents’ time preferences (see for example Osborne and Rubinstein 1990).

prevents outside options from being determined in a circular way, and it explains how the equilibrium outcome is uniquely pinned down by the Nash bargaining solution. For example, as illustrated in [Figure 2](#), this prevents MBA graduates from claiming an outside option of w in Wall Street by arguing that this is what they get in Main Street, while the only reason that Main Street pays them w is that Wall Street does.

The theory of coalition formation that I describe in this paper is tractable not only because the coalitions that form in equilibrium have a nice structure that illustrates how shocks propagate via outside options, but also because a simple algorithm identifies which coalitions form and how they share the resulting surplus in equilibrium. This provides a wealth of comparative statics results. For example, as in the canonical marriage market model of [Becker \(1973\)](#), agents necessarily match in a positive assortative way if and only if their skills are complementary. But, in contrast to Becker's theory (as well as much of the subsequent literature), the framework that I describe in this paper pins down prices uniquely, and hence provides testable predictions about how positive assortative matching affects the way in which shocks propagate via outside options. In particular, in two-sided pairwise matching settings where workers and firms, say, match in a positive assortative way, shocks propagate *in blocks*—in the sense that a shock that propagates from one worker to another one also affects every worker whose skill is in between. Hence, this theory suggests a mechanism by which an increase in labor market sorting—as we have observed in many countries over the last decades—can lead to a sharp disconnection between the determinants of high-skill and low-skill wages.⁸

Roadmap

The rest of this paper is organized as follows. I start in [section 2](#) by illustrating the setting and the main result of this paper with a simple example. I then describe the model in [section 3](#) and its essentially-unique stationary subgame-perfect equilibrium in [section 4](#). I illustrate the comparative statics of the resulting theory in [section 5](#), and I further discuss the contribution of this paper to the related literature in [section 6](#). Finally, I conclude in [section 7](#). I defer the formal proofs of most of the results to the appendix.




Figure 3: A chef generates 100 dollars when she matches with a maître (by starting a high-end restaurant, say) and 80 dollars when she matches with a manager (by starting the occasional low-end restaurant with great food, say). Similarly, a cook generates 60 dollars when she matches with a maître (by starting the all-too-common high-end restaurant with unimpressive food, say) and 50 dollars when she matches with a manager (by starting a low-end restaurant, say).

2 Illustration of the setting and the main result

In this section, I illustrate the setting and the main result of this paper using an example. I would like to emphasize that the objective of this example is neither to illustrate the full generality of the setting nor its leading application, but to *illustrate the main ideas of this paper in the simplest possible setting*. In particular, in this example, I assume that only pairs of agents can match, and that productivity is the only source of heterogeneity—but the general model allows coalitions of arbitrary size as well as more varied sources of heterogeneity.

Consider a large city where different agents (in the culinary industry, say) go to in search of business opportunities. For simplicity, assume that there are only four types of agents in this industry: *Managers, maîtres, cooks* and *chefs*, all of them risk neutral.^{9,10} Agents of all types arrive to the city over time (perhaps with the excuse of attending a prestigious culinary school) to find potential partners with whom to start a business venture. For simplicity, assume that each agent can only be part of one such venture (because each feasible venture is a lifelong full-time project, say), and that only bilateral coalitions between one maître/manager and one chef/cook are feasible. Moreover, assume that the surplus of each coalition is independent of which other matches form (because each venture is implemented in a different part of the world, say), and that the surpluses of the four possible coalitions are as illustrated in [Figure 3](#).¹¹

⁸See [Eeckhout \(2017\)](#) for an insightful recent survey of the literature on sorting in labor markets.

⁹For the purposes of this example, “maîtres” and “chefs” are high-end managers and cooks, respectively.

¹⁰I am grateful to Rachel Kranton for encouraging me to illustrate the results of this paper along the lines of [Hart and Moore’s \(1990\)](#) gourmet seafare example.

¹¹Food being the most important part of a culinary experience, I assume that a match between a chef and a




Figure 4: Illustration of the equilibrium when the bargaining friction q is small (.02). The number next to each box is the amount that the corresponding agents are indifferent between accepting and rejecting. An arrow from type i to type j indicates that each agent of type i makes an offer to type j in equilibrium. Every equilibrium offer leaves the receiver indifferent between accepting and rejecting (but is accepted).

Which business ventures form, and how are the resulting surpluses shared? How does an increase in the productivity of the chef-maître coalition (caused by a global increase in high-end tourism, say), or an improvement in the chefs' bargaining position (caused by a new technology that allows them to directly fly food to their clients' doors, say) affect this market? In order to investigate these types of questions, I study the equilibrium behavior of these agents when they bargain according to an infinite-horizon protocol in the spirit of the alternating-offers model of [Rubinstein \(1982\)](#): In each period, one of the agents that is in the city looking to form a business venture is selected uniformly at random to be the proposer. The selected agent can propose a match as well as how to share its surplus. This captures the fact that starting a business venture requires that someone has an idea: Once an agent has an idea, she can propose to implement it with another agent who, in turn, decides whether to join this venture (at the proposed terms of trade) or to wait for better opportunities to arise. The bargaining friction that incentivizes agreements is that, after each period, each agent has to leave the city (because of personal reasons, say) with some probability q , preventing her from starting any venture.¹² Hence, when an agent is deciding whether to accept or reject an offer, she has to trade off the potential for better opportunities arising in the future (e.g., having a business idea herself), with the risk of having to leave the market before matching.

The unique subgame-perfect equilibrium of this game is illustrated in [Figure 4](#). Chefs propose forming business ventures with maîtres, and vice versa. Hence, even if chefs and maîtres can match with managers and cooks, respectively, they effectively bargain over how

manager generates more value than a match between a cook and a maître.

¹²For simplicity, in this example I normalize to zero the surplus that each agent obtains when she has to leave the market before she has created a venture.




Figure 5: In the limit as the bargaining friction q goes to zero, the cutoffs of chefs, maîtres, managers and cooks converge to 50, 50, 30 and 20, respectively.

to share their gains from trade *as if* they were the only two types in the market. Intuitively, the fact that a chef can always find a maître to bargain with, and vice versa, implies that their surpluses in other matches do not affect their bargaining position. Indeed, since making offers to others is off the equilibrium path, and an agent never benefits from receiving an offer (since equilibrium offers leave the receiver indifferent between accepting and rejecting it), in the limit as the bargaining friction q vanishes, chefs and maîtres share their gains from trade equally. More generally, their terms of trade are as prescribed by the Nash bargaining solution, with the threat points given by their payoffs when they are forced out of the city before they can start a business (which in this example I have normalized to zero), as suggested by the outside option principle.

When a manager is the proposer, in equilibrium she always offers to match with a chef. In this case, the chef has to trade off the gains from accepting such an offer with the expected gains of waiting to be able to make an offer in the future (at the risk that she might be forced to leave the market before this happens). As a result, when the bargaining friction q is small enough, the manager has to offer a chef close to 50 dollars (approximately what she gets when proposing to match with a maître) for her to accept. In particular, in the limit as the bargaining friction q vanishes, they share their surplus $50 - 30$, as suggested by the outside option principle.¹³ Similarly, the cooks propose to match with the managers, and—in the limit as the bargaining friction q vanishes—they share their surplus $20 - 30$, again as suggested by the outside option principle. Figure 5 illustrates this limit equilibrium outcome.

As the bargaining friction q vanishes, the sharing rule in each equilibrium coalition con-

¹³As long as the bargaining friction q is positive, chefs can obtain more from maîtres than from managers when they are the proposers, because they can exploit more their ability to make take-it-or-leave-it offers with the former than with the latter. Intuitively, maîtres have more to lose by rejecting an offer than managers do, because their matching opportunities are better. The difference, however, converges to zero as the bargaining friction q vanishes.




Figure 6: The number associated to an arrow that points from type i to type j is how much type i can justify in the ij coalition using the Nash bargaining solution while honoring the amount that type j can justify—in a similar way—in the other coalition that j is part of. For example, chefs can justify a payoff of 50 as being the result of Nash bargaining with maîtres subject to a lower bound of 30 on the maîtres payoffs (which, in this case does not bind), and managers can justify a payoff of 30 as being the result of Nash bargaining with chefs subject to a lower bound of 50 on the chefs payoffs.

verges to the one prescribed by the Nash bargaining solution, with outside options entering as prescribed by the outside option principle, and determined as follows: Each type’s outside option in any given coalition is her Nash bargaining share in the other coalition that she is part of subject to *the other’s* outside option in that coalition.¹⁴ For example, as Figure 6 illustrates:

1. When bargaining to form a chef-maître coalition (whose surplus is 100), chefs’ and maîtres’ outside options are 40 and 30 respectively, so neither of them binds.¹⁵
2. When bargaining to form a chef-manager coalition (whose surplus is 80), chefs’ and managers’ outside options are 50 and 25 respectively, so only the chefs’ outside options bind.
3. When bargaining to form a cook-manager coalition (whose surplus is 50), cooks’ and managers’ outside options are 10 and 30 respectively, so only the managers’ outside options bind.
4. When bargaining to form a cook-maître coalition (whose surplus is 60), cooks’ and maîtres’ outside options are 20 and 50 respectively, so this coalition is not sufficiently

¹⁴More generally, in the limit as the bargaining friction q vanishes, each type’s share in every equilibrium coalition that she is part of is the maximum amount that she can justify as being the result of the Nash bargaining solution in some coalition subject to the others’ shares. In particular, every binding outside option is determined in a coalition that forms in equilibrium.

¹⁵The fact that none of the outside options in this coalition bind explains why chefs and managers share their surplus equally, even if they are not completely symmetric.

productive to meet its members' outside options—and hence it never forms in equilibrium.

Remarkably, there is a unique outside option profile that satisfies this property, and the strategic forces in the non-cooperative model suggest the resulting outcome as a natural point to settle on when bargaining in a decentralized way. Informally, the strategic forces in the non-cooperative model require that each type is able to justify her outside option in each coalition as resulting from the Nash bargaining solution in some other coalition without appealing to her own outside option there. Intuitively, this prevents outside options from being determined in a circular way, and it explains why the outcome is uniquely pinned down in equilibrium by the Nash bargaining solution. For instance, without this requirement, the chefs would be able to justify that their outside option when bargaining with a maître is getting 55 dollars, say, from a manager, by arguing that, when bargaining with a manager, their outside option is getting 55 dollars from a maître. However, the strategic forces in the model prevent both chefs and maîtres to obtain more than what they can justify by their Nash bargaining shares in their coalition without appealing to their own outside options, and this pins down the equilibrium outcome.

In order to understand how bargaining outcomes are determined, the coalitions that form in equilibrium can be organized into tiers in such a way that the surplus of each coalition is shared—as the bargaining friction q vanishes—according to the Nash bargaining solution, and the binding outside options determined in higher tiers. [Figure 7](#) illustrates that the chef-maître coalition is in the first tier. The Nash bargaining solution in this coalition (without binding outside options) pins down the chefs' binding outside options when bargaining with managers which, in turn, pins down the managers' binding outside options when bargaining with cooks. This illustrates how different economic shocks propagate via outside options from higher to lower tiers, but not vice versa. For instance, an increase in the surplus of the chef-maître coalition propagates downwards—via the chefs' and managers' outside options—to affect everyone. But an increase in the surplus of the cook-manager coalition only affects cooks, who absorb the whole surplus increase.

I defer to [section 4](#) further discussion of the intuition for these results as well as the description of the algorithm that characterizes the equilibrium. I now turn to describing how the results illustrated in this section generalize to settings with arbitrarily many types with (potentially) different risk preferences and proposer probabilities, and where the productive coalitions can be of arbitrary form and size.



Figure 7: The equilibrium tier structure in the culinary example. A type’s payoff is determined (by the Nash bargaining solution subject to the binding outside options determined in higher tiers) in the coalition where her name is in bold.

3 Model: The bargaining game \mathcal{G}

As already emphasized above, the model is intended to capture the predominant economic forces in large markets with dynamic entry—where the relevant matching opportunities are roughly constant over time. For simplicity, and following the approach of [Rubinstein and Wolinsky \(1985\)](#) and the subsequent literature on non-cooperative bargaining in stationary markets, I assume that the agents enter the market over time in such a way that there is always one active agent of each type in the market. In contrast, in [Elliott and Talamàs \(2018\)](#) we model these markets as featuring an exogenous process that determines how agents enter the market over time.¹⁶ While the results of this paper carry over to that more realistic model, I take the traditional modeling approach here both for simplicity and in order to be able to more easily contrast the results of this paper with those in the existing literature.

3.1 Primitives

There is a finite set N of different types of agents, and a sequence of agents of each type. Different types of agents can—by matching—produce different amounts of perfectly divisible surplus (e.g., money). For simplicity, I assume that each match containing at least one

¹⁶In [Elliott and Talamàs \(2018\)](#), we investigate the private incentives to invest in different skills and relationships when these must be sunk before entering the market, and we find that—in dynamically-thick markets—the holdup problem vanishes with the bargaining frictions.

agent of each type in $C \subseteq N$ produces $y(C) \geq 0$ units of surplus when it forms.^{17,18} I refer to $y_i := y(i)$ as *type i 's autarky surplus*, which can be interpreted as *type i 's exogenous outside option*: How much she can obtain without anyone else's consent.

While surplus is perfectly divisible, the utility generated by each match is—in general—imperfectly transferable, because the agents' utility functions need not be linear in money. In particular, as in the canonical bargaining framework of Nash (1950), the preferences of each agent of type i are represented by the von-Neumann Morgenstern utility function u_i , which is a concave, strictly increasing, and twice-continuously differentiable function of the money that she gets.

Finally, I take as given a “bargaining power” profile $\mathbf{p} \in [0, 1]^N$, with $\sum_i p_i \leq 1$, which can be interpreted as an *exogenous measure of the relative bargaining powers* of the different types. In other words, \mathbf{p} can be thought of as capturing primitives other than preferences and productivities (e.g., relative scarcities of different types) that are relevant for bargaining outcomes but that the present framework otherwise abstracts from.

I now turn to describing the bargaining protocol that turns these primitives into a well-defined non-cooperative game of coalition formation.

3.2 Bargaining protocol

Bargaining occurs in discrete periods $t = 1, 2, \dots$. In each period, the first agent in sequence of each type (yet to leave the market) is *active*. At most one active agent is selected at random to be the proposer (the active agent of type i is selected with probability p_i). The proposer, of type i say, chooses one coalition $C \subseteq N$, and proposes a split of the corresponding surplus among its members. The active agents of each type in $C - i$ then decide in (a pre-specified) order whether to accept or reject this proposal.¹⁹ If all of them accept, then they match with the proposer and they, together with the proposer, leave the market with the agreed shares.

¹⁷As long as there is an upper bound on how many agents of a given type are productive in a given coalition, the fact that the surplus of each coalition does not depend on whether it contains one or more agents of a given type is without loss of generality, because types can always be defined so that this property holds. For example, suppose that everyone is identical, and that coalitions of one and two agents produce 1 and 2 units of surplus, respectively. This can be captured by letting there be two types of agents, with coalitions consisting of an agent of any one of these types producing 1 unit of surplus, and coalitions containing both these types producing 2 units of surplus.

¹⁸For expositional clarity, I usually reserve the term “coalition” to refer to a set of *types*, while I use the term “match” to mean a set of *agents* (that match).

¹⁹The order in which the agents respond is not relevant for the results.

Otherwise, they, and the proposer, wait for the next period, as do all the active agents that are neither proposers nor receivers of an offer in this period. At the end of each period, each active agent is independently forced to leave the market with probability $q > 0$, in which case she obtains her autarky surplus.

The game is common knowledge, and it features perfect information. I now turn to describing the notion of equilibrium that I focus on throughout this paper.

3.3 Histories, strategies and equilibrium

For each period t , let h_t be a history of the game up to (but not including) period t , which is a sequence of t pairs of proposers and coalitions proposed—with corresponding proposals and responses. There are two types of histories at which some agent must take an action. First, (h_t, i) consists of h_t followed by the active agent of type i being selected to be the proposer in period t . Second, $(h_t, i \rightarrow C, x, j)$ consists of (h_t, i) followed by the active agent of type i proposing that the surplus of coalition C is shared according to the profile x in \mathbb{R}^C , and all the active agents in C preceding type j in the response order having accepted.

A strategy σ_i for type i specifies, 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 $\sigma_i(h_t, j \rightarrow C, x, i)$ following the history $(h_t, j \rightarrow C, x, i)$.²⁰ The strategy profile $(\sigma_i)_{i \in N}$ is a *subgame-perfect equilibrium* of the game \mathcal{G} if it induces a Nash equilibrium in the subgame following every history. A subgame-perfect equilibrium is *stationary* if no type's strategy conditions on the history of the game except—in the case of a response—on the going proposal and on the identity of the proposer. I often refer to a stationary subgame-perfect equilibrium simply as an *equilibrium*.

I now turn to describing (i) how the bargaining game \mathcal{G} admits an essentially unique equilibrium, and (ii) which coalitions form and how the resulting surplus is shared in this equilibrium.

4 Essentially unique equilibrium

In this section, I show that the bargaining game \mathcal{G} admits an essentially unique equilibrium, and I describe an algorithm that identifies it. This provides the basis for the main result of this paper, which shows how—in the limit as the bargaining friction q vanishes—the

²⁰I allow for mixed strategies, so $\sigma_i(h_t, i)$ and $\sigma_i(h_t, j \rightarrow C, x, i)$ are probability distributions over $2^N \times \mathbb{R}_{\geq 0}^N$ and $\{\text{Yes, No}\}$, respectively.

equilibrium payoff profile is the only one that is such that each type obtains the maximum that she can justify as resulting from the Nash bargaining solution in some coalition *without appealing to her own outside option there*.

I start in [subsection 4.1](#) by deriving the system of equations that determines the equilibrium payoffs. This system formalizes the outside option Gordian knot described in [section 1](#): Each type's payoff (in a period in which she is not the proposer) is determined by the maximum surplus that she can generate (when she is the proposer) net of the others' payoffs. But the others' payoffs depend, in turn, on the others' payoffs, and so on. The objective of most of the rest of this section is to show that this system admits a unique solution. The strategy to prove this is based on the observation that there always exists at least one coalition that is sufficiently productive so that—when bargaining in this coalition—none of its members' has a credible outside option. For instance, in the example of [section 2](#), the chef-maître coalition is sufficiently productive so that neither chefs nor maîtres outside options are credible when bargaining to form this coalition. As a result, they essentially share the surplus of this coalition equally, and this determines their outside options when bargaining in other coalitions.

I divide the equilibrium characterization strategy into three steps. First, in [subsection 4.2](#), I describe an auxiliary non-cooperative game of isolated bargaining within each coalition, and I show how the equilibria of these auxiliary games can be used to derive an upper bound on each type's equilibrium payoff in the game \mathcal{G} . Second, in [subsection 4.3](#), I show that there exists at least one coalition where this bound is tight for all of its members, which implies that this bound is actually the payoff of all of its members *in every equilibrium of the game \mathcal{G}* . Third, in [subsection 4.4](#), I leverage these observations to recursively characterize the unique equilibrium payoffs of all the types.

Finally, I present and discuss the main result of this paper in [subsection 4.5](#): In the limit as the bargaining friction q vanishes, the equilibrium payoff profile converges to the unique profile that gives each type the maximum that she can justify in some coalition using the Nash bargaining solution while honoring *the others'* payoffs.

4.1 Equilibrium cutoff profile

[Proposition 4.1](#) describes the relevant features of the equilibrium of the game \mathcal{G} . This equilibrium is essentially unique—in the sense that (i) each type's payoff is the same in every equilibrium, (ii) each type's on-path responses are the same in every equilibrium, and (iii) the proposals of each type whose expected equilibrium payoff is strictly higher than her

autarky payoff are the same in every equilibrium except in non-generic cases in which one type's maximum surplus net of the others' payoffs is achieved in more than one coalition.

Proposition 4.1. *Each type i has a cutoff w_i such that, in every stationary subgame-perfect equilibrium of the game \mathcal{G} , she always accepts (rejects) every offer that gives her strictly more (less) than w_i . Moreover, each type whose expected equilibrium payoff is higher than her autarky payoff proposes that one of the coalitions C with the biggest net surplus $y(C) - \sum_{j \in C-i} w_j$ forms, and she offers each of its members $j \neq i$ the amount w_j , all of whom accept.*

Remark 4.1. *In the special case of pairwise matching settings (in which each feasible match contains at most two agents), every subgame-perfect equilibrium is in stationary strategies; see Talamàs (2018, Proposition A.1).²¹ Hence, in this case, Proposition 4.1 holds without restricting attention to stationary strategies.*

Figure 4 in section 2 illustrates Proposition 4.1 in the context of the example described there. The cutoff of chefs and maîtres is 48, the cutoff of managers is 29, and the cutoff of cooks is 19.²² In this case, each proposer i finds the type j that maximizes the net surplus $y(i, j) - w_j$ and offers the active agent of type j her cutoff w_j (and all such offers are always accepted). In particular, chefs and maîtres make offers to each other, managers make offers to chefs, and cooks make offers to managers.

I devote the rest of this subsection together with subsections 4.2, 4.3 and 4.4 to describing the argument behind the proof of Proposition 4.1 (with formal details deferred to the appendix). The immediate objective is to derive the system of equations (2) that determines the equilibrium payoffs in the bargaining game \mathcal{G} . To do this, consider a stationary subgame-perfect equilibrium of this game. For each type i , let w_i be the amount that type i is indifferent between accepting and rejecting in any given period. On the equilibrium path, type j accepts every offer that gives her exactly w_j (otherwise, the proposer would have no best response), so the maximum amount that type i can obtain when she is the proposer is

$$(1) \quad \max_{C \subseteq N} \left[y(C) - \sum_{j \in C-i} w_j \right].$$

Note that type i makes acceptable offers in equilibrium if the quantity in (1) is strictly bigger than her autarky utility $u_i(y_i)$.²³

²¹Proposition A.1 in Talamàs (2018) is stated for the case of linear preferences, but its proof goes through unchanged in the case of possibly heterogeneous and concave utilities of the present paper.

²²For simplicity, I approximate each number to the nearest integer.

²³To see this, let V_i and W_i be the expected utility of an agent of type i when she is and she is not selected to be the proposer, respectively. We have that $W_i = qu_i(y_i) + (1 - q)(p_i V_i + (1 - p_i)W_i)$, so $V_i > u_i(y_i)$ implies that $V_i > W_i$. Hence, every type i with $V_i > u_i(y_i)$ is strictly worse off by delaying.

By definition, each type i is indifferent between obtaining w_i right away, which gives her utility $u_i(w_i)$, and waiting for the next period, which gives her an expected utility of

$$q \underbrace{u_i(y_i)}_{\text{autarky utility}} + (1 - q) \left[p_i u_i \left(\underbrace{\max_{C \subseteq N} \left[y(C) - \sum_{j \in C-i} w_j \right]}_{\text{expected utility when proposing}} \right) + (1 - p_i) \underbrace{u_i(w_i)}_{\text{expected utility when not proposing}} \right].$$

To see this, note that waiting one period involves a risk of being forced to leave the market (which materializes with probability q , and in which case the agent gets her autarky surplus y_i). In the event that she is not forced to leave at the end of the period, she has the opportunity to make a proposal in the next period with probability p_i , in which case she obtains $y(C) - \sum_{j \in C-i} w_j$. Otherwise, she either receives an offer that gives her w_i (which she accepts), or she does not receive any offer; in either case, her expected utility is $u_i(w_i)$. Rearranging terms gives that i 's expected utility when she is not the proposer is a weighted average of her autarky utility and her expected utility when she is the proposer, with the weight $\alpha_i := \frac{q}{(1-q)p_i+q}$ on the former converging to 0 as the bargaining friction q goes to 0.

$$(2) \quad \underbrace{u_i(w_i)}_{\text{exp. utility when not proposing}} = \alpha_i \underbrace{u_i(y_i)}_{\text{autarky utility}} + (1 - \alpha_i) u_i \left(\underbrace{\max_{C \subseteq N} \left[y(C) - \sum_{j \in C-i} w_j \right]}_{\text{exp. utility when proposing}} \right) \quad \forall i \in N,$$

In particular, in the absence of bargaining frictions (i.e., when $q = 0$), system (2) becomes

$$(3) \quad w_i = \max_{C \subseteq N} \left[y(C) - \sum_{j \in C-i} w_j \right] \quad \text{for all types } i \text{ in } N,$$

which, in general, has a great multiplicity of solutions.²⁴ For instance, in the example described in section 2, one extreme solution to system (3) is the profile that gives chefs and cooks 100 and 60, respectively, and 0 to both maîtres and managers. Another extreme solution to system (3) is the profile that gives 0 to both chefs and cooks, and 100 and 50 to maîtres and managers, respectively.

I now turn to showing that—as long as the bargaining friction q is positive—system (2) has a unique solution, which I refer to as the *equilibrium cutoff profile*. The limit of this profile as the bargaining friction q vanishes selects a unique profile out of the many possible solutions to system (3). In other words, the unique equilibrium of the non-cooperative game picks one of the many plausible bargaining outcomes in the frictionless case.²⁵ Most of the rest of this

²⁴Each solution of system (3) can be seen as a profile of *competitive equilibrium prices*, in the sense that every proposer obtains the maximum possible amount of surplus taking as given the others' wages. The dynamic entry of agents into the market implies that the market clearing condition does not discipline the prices in this setting.

²⁵This is, of course, the comparative advantage of the non-cooperative approach (see Rubinstein 1982).

paper is devoted to describing the properties of this solution, and to using this solution to shed light on the determinants of bargaining outcomes in stationary markets.

4.2 A personalized upper bound on the equilibrium cutoffs

The objective of this subsection is to provide a personalized upper bound on the equilibrium cutoff w_i of each type i . The basic idea is the following: When proposing that a coalition forms, its members' outside options have to be met in order to induce them to accept. Hence, roughly speaking, a type cannot be made worse off when the others' outside options deteriorate. This suggests that type i 's equilibrium cutoff in the *hypothetical* situation in which she can (i) choose any coalition, and (ii) prevent its members from making offers to any other coalition is an upper bound on her equilibrium cutoff w_i in the bargaining game \mathcal{G} of interest.²⁶

In order to formalize this idea, I consider a family $\{\mathcal{G}^C\}_{C \subseteq N}$ of variants of the bargaining game \mathcal{G} in which only one coalition $C \subseteq N$ is allowed to form.²⁷ The game \mathcal{G}^C is analogous to a multilateral version of the canonical alternating-offers model of Rubinstein (1982), where the one feasible coalition C is given exogenously, and its members only bargain over how to share the surplus $y(C)$ of this coalition. Hence—as in the canonical multilateral Rubinstein bargaining game—each game \mathcal{G}^C has a unique stationary subgame-perfect equilibrium. Moreover, analogously to the results in Binmore, Rubinstein, and Wolinsky (1986), the equilibrium cutoff profile in this game converges to the (generalized) Nash bargaining solution in coalition C , with the threat points given by autarky. In other words, the equilibrium cutoff profile in the auxiliary game \mathcal{G}^C converges to the unique profile that solves

$$(4) \quad \operatorname{argmax}_{s \in \mathbb{R}^C} \prod_{j \in C} [u_j(s_j) - u_j(y_j)]^{p_j} \quad \text{subject to the feasibility constraint } y(C) \geq \sum_{j \in C} s_j.$$

For instance, Figure 8 illustrates each type's cutoff in each relevant auxiliary game in the example of section 2. Given that the preferences and the proposer probabilities in this example are homogeneous, isolated bargaining between two agents essentially leads to equal sharing. More precisely, each type's cutoff in the auxiliary game associated with any given coalition (that she part of) is just below half of its surplus. This reflects the fact that the proposer obtains slightly more than half of the available gains from trade. But, in the limit as

²⁶In this exercise, preventing agents from accepting offers from other coalitions is not necessary, because—given that the equilibrium offers leave the respondent indifferent between accepting and rejecting—agents' payoffs are determined by the amount that they can obtain when they are given the opportunity to propose.

²⁷More precisely, for each coalition $C \subseteq N$, the game \mathcal{G}^C is defined exactly as the bargaining game \mathcal{G} , with the following modification: The surplus $y(D)$ of each coalition $D \neq C$ is reduced to 0.




Figure 8: The number associated to a link between type i and type j is their equilibrium cutoff in the auxiliary game in which the surplus of all the other coalitions is artificially set to 0. An arrow from type i to type j indicates that the ij coalition is i 's best coalition, and the associated number is her best share.

the bargaining friction q vanishes, these cutoffs converge to exactly half of the surplus of the corresponding coalition—as prescribed by the Nash bargaining solution.

I can now describe more precisely the exercise outlined above: I define each type's *best share* to be her maximum equilibrium cutoff—across all coalitions $C \subseteq N$ —in the auxiliary game \mathcal{G}^C , and I say that a coalition $C \subseteq N$ is one of type i 's *best coalitions* if i 's equilibrium cutoff in \mathcal{G}^C is her best share. Informally, suppose that we ask each type: “Which coalitions would you be happy choosing if you could pick one coalition $C \subseteq N$ and bargain in isolation with its members (according to the auxiliary game \mathcal{G}^C)?” Each type's best coalitions are the ones that she would point to, and her best share is her equilibrium cutoff in this hypothetical situation. Finally, I say that a coalition is a *perfect coalition* if it is a best coalition of all of its members. For instance, in the example of [section 2](#), the best share of both chefs and maîtres is 48, and their best coalition is the chef-maître coalition, so this coalition is a perfect coalition. The best share of managers is 38, and their best coalition is the chef-manager coalition. Finally, the best share of cooks is 29, and their best coalition is the cook-maître coalition.

[Proposition 4.2](#) below formalizes the intuition that the equilibrium cutoff of each type in the bargaining game \mathcal{G} cannot be higher than her best share. This is especially useful because—as I show in [subsection 4.3](#) below—this bound is always tight for at least one type.

Proposition 4.2. *Let the profile w in \mathbb{R}^N be a solution to system (2). For each type i , w_i is bounded above by i 's best share.*

[Figure 8](#) illustrates the hypothetical exercise described above in the context of the example in [section 2](#). For instance, if cooks were able to choose between maîtres and managers, and bargain with them in isolation, their equilibrium cutoff would be 29 and 24, respectively. Intuitively, 29 must then be an upper bound on the cooks' equilibrium cutoff, because, in the

bargaining game \mathcal{G} , both managers and maîtres can in fact choose to make offers to chefs as well, which can only improve their bargaining position. Indeed, as illustrated in Figure 4, their equilibrium cutoff is 19.

4.3 The personalized upper bound is tight for at least one type

I now describe how the personalized upper bound on each type’s equilibrium cutoff provided by Proposition 4.2 above is tight for at least one type, which provides the basis of the recursive characterization of everyone’s equilibrium cutoffs that I describe in subsection 4.4. In particular, the combination of the two results below (Proposition 4.3 and Proposition 4.4) implies that there always exists at least one coalition that is sufficiently productive so that—when bargaining to form this coalition—none of its members’ outside options binds in equilibrium.

On the one hand, Proposition 4.2 above implies that, for each type i in a perfect coalition C , her equilibrium cutoff in the auxiliary game \mathcal{G}^C is an *upper bound* on her equilibrium cutoff in the game \mathcal{G} . On the other hand, as highlighted by Proposition 4.3 below, the fact that this bound holds for all the types in a perfect coalition implies that this upper bound is actually also a *lower bound* on their payoffs. Indeed, the fact that the cutoffs in \mathcal{G}^C are an upper bound on all of its members’ equilibrium cutoffs in the bargaining game \mathcal{G} of interest implies that—when bargaining in the game \mathcal{G} —no one in coalition C has a better outside option than proposing to form coalition C . As a result, in every equilibrium of the game \mathcal{G} , every type in a perfect coalition C can do at least as well as in the auxiliary game \mathcal{G}^C .

Proposition 4.3. *Let the profile w in \mathbb{R}^N be a solution to system (2). If type i is in a perfect coalition, then w_i is her best share.*

For instance, given that the cutoff of the maîtres in the example of section 2 is bounded above by 48, in the equilibrium of the game of interest, the chefs can do as well as they can in the auxiliary game in which they can bargain in isolation with the maîtres because, in this hypothetical case, the cutoff of the maîtres is 48 (its upper bound in the game of interest). Hence, the upper bound of 48 on the cutoff of the chefs is tight, and analogously for the upper bound of 48 on the cutoff of the maîtres.

Proposition 4.3 above is useful mainly for two reasons. First, we can tell whether a coalition is perfect using only the equilibrium cutoff profiles in the auxiliary games $\{\mathcal{G}^C\}_{C \subseteq N}$, which, as I have pointed out above, are familiar and easy to compute. In particular, we do not have to be able to solve the system (2) in order to be able to identify the perfect coalitions.

Second, as highlighted by [Proposition 4.4](#) below, we can always identify at least one perfect coalition. Hence, we can pin down the equilibrium cutoff of a nonempty subset of the types using the auxiliary games $\{\mathcal{G}^C\}_{C \subseteq N}$.

Proposition 4.4. *There exists at least one perfect coalition.*

[Proposition 4.4](#) above follows from the fact that the equilibrium cutoffs in the auxiliary games $\{\mathcal{G}^C\}_{C \subseteq N}$ satisfy the following property: For any two types i and j and any two coalitions C and D containing both these types, if i 's cutoff in the auxiliary game \mathcal{G}^C is higher than i 's cutoff in the auxiliary game \mathcal{G}^D , then the same is true for type j (that is, j 's cutoff in the auxiliary game \mathcal{G}^C is higher than j 's cutoff in the auxiliary game \mathcal{G}^D).²⁸ To see why this implies the existence of a perfect coalition, note that this precludes the existence of cycles in the network whose nodes are coalitions and whose link from C to C' indicates that there is a type in C whose best share is strictly higher in C' than in C . As a result, every path (or sequence of distinct links) in this network must end at a perfect coalition.

Perhaps the best way to gain intuition for [Proposition 4.4](#) is to recall that the Nash bargaining shares s in each coalition C equalize the ratio $\frac{u_i(s_i)}{u'_i(s_i)}$ among all of its members. Since the utility functions are concave, each type's Nash bargaining share is increasing in this ratio, so—given that, as described above, the equilibrium cutoffs in each game \mathcal{G}^C converge to the Nash bargaining shares in C —the coalition with the highest such ratio is a perfect coalition for all small enough bargaining frictions (at least in the generic case in which each type's Nash bargaining share in different coalitions is different).²⁹

4.4 Recursive characterization of the equilibrium cutoff profile

The subsection above describes how to pin down the equilibrium cutoffs of a nonempty subset of types from the equilibrium cutoffs in the family of auxiliary games $\{\mathcal{G}^C\}_{C \subseteq N}$. The goal now is to describe the algorithm that pins down everyone's equilibrium cutoffs in the game \mathcal{G} . For this, the immediate objective is to provide an upper bound on the equilibrium cutoff of each type *when some of the other types' cutoffs have already been pinned down*.

As in [subsection 4.2](#), the basic idea is that no type can be hurt when the others' outside options deteriorate. The difference with the argument above is that, now—in the hypothetical

²⁸[Pycia \(2012\)](#) labels this property *pairwise-aligned preferences over coalitions*, and shows that the Nash bargaining solution generates pairwise-aligned preferences over coalitions.

²⁹This is essentially the argument that [Pycia \(2012\)](#) gives to show that the Nash bargaining solution generates pairwise aligned preferences over coalitions.

situation in which agents bargain in isolation in one coalition—I don't allow the outside options of the agents whose cutoff has already been defined to deteriorate below their cutoff. In order to formalize this idea, for every coalition $C \subseteq N$ and every profile $\mathbf{x} \in \mathbb{R}^N$ with $y(C) \geq \sum_{j \in C} x_j$, I consider the auxiliary game \mathcal{G}_x^C , which is a variant of the bargaining game \mathcal{G}^C in which each agent of type i can choose to immediately leave the market with payoff x_i any time that she rejects any proposal.

The game \mathcal{G}_x^C is analogous to a multilateral version of the canonical alternating-offers model of Rubinstein (1982) with exogenous outside options (as considered by Binmore, Rubinstein, and Wolinsky 1986, for example) where both the feasible coalition C and the outside option profile \mathbf{x} are exogenous, and its members only bargain over how to share the resulting gains from trade subject to their exogenously given outside options. This game has a unique stationary subgame-perfect equilibrium. In particular, analogously to the results in Binmore, Rubinstein, and Wolinsky (1986), the equilibrium cutoff profile of the game \mathcal{G}_x^C converges to the (generalized) Nash bargaining solution in coalition C , with the threat points given by autarky, and the outside option profile \mathbf{x} imposing only lower bounds on the payoffs. In other words, the equilibrium cutoff in the auxiliary bargaining game \mathcal{G}_x^C converges to the profile that solves

$$(5) \quad \operatorname{argmax}_{\mathbf{s} \in \mathbb{R}^C} \prod_{j \in C} [u_j(s_j) - u_j(y_j)]^{p_j} \quad \text{subject to} \quad y(C) \geq \sum_{j \in C} s_j \quad \text{and} \quad s_j \geq x_j \quad \text{for all } j \text{ in } C.$$

Figure 9 illustrates the equilibrium cutoffs in each of the relevant auxiliary games in the example of section 2, when the outside options are set to

$$(6) \quad x_{\text{chefs}}^1 = x_{\text{maitres}}^1 = 48 \quad \text{and} \quad x_{\text{cooks}}^1 = x_{\text{managers}}^1 = 0.$$

Naturally, in this case, the only cutoffs that are different from the case in which everyone's outside options are zero (the case illustrated in Figure 8) are those associated with coalitions involving either the chefs or the maitres (the two types whose outside options have been updated to be their equilibrium cutoffs).

For each payoff profile \mathbf{x} in \mathbb{R}^N , I define type i 's \mathbf{x} -best share to be i 's maximum cutoff in the auxiliary game \mathcal{G}_x^C across all coalitions $C \subseteq N$. Similarly, I say that a coalition C is type i 's \mathbf{x} -best coalition if i 's equilibrium cutoff in \mathcal{G}_x^C is her \mathbf{x} -best share.³⁰ Finally, I say that a coalition is \mathbf{x} -perfect if it is an \mathbf{x} -best coalition of all of its members. For example, Figure 9 illustrates that the chef-manager coalition is the managers' \mathbf{x}^1 -best coalition (and is hence \mathbf{x}^1 -perfect), and that the managers' \mathbf{x}^1 -best share is 29. The argument analogous to the

³⁰These concepts are well defined as long as there is a coalition $C \subseteq N$ with $i \in C$ and $\sum_{j \in C} x_j \leq y(C)$.




Figure 9: The number that is closest to type i associated with a link between i and j is approximately her equilibrium cutoff in the auxiliary game in which the surplus of all the coalitions other than the one between i and j is artificially set to zero, and the outside option profile x is given by (6). An arrow from type i to type j indicates that the ij coalition is i 's x -best coalition in this case, and the associated number is her x -best share.

one that proves [Proposition 4.2](#) and [Proposition 4.3](#) above shows that these x^1 -best shares are upper bounds on the respective equilibrium cutoffs, and that these bounds are tight in every x^1 -perfect coalition. In this case, this implies that the x^1 -best share of the managers is their equilibrium cutoff. [Definition 4.1](#) below formalizes this argument, in the form of an algorithm that recursively pins down everyone's cutoffs. Informally, this algorithm runs as follows: The initial payoff profile is $x^0 = 0$. In each step $k = 1, 2, \dots$, the payoff profile is x^{k-1} , and each type points to her x^{k-1} -best coalitions. A coalition that is such that all of its members j with $x_j^{k-1} = 0$ point to it "clears," and the outside option x_j^k of each of its members is updated to be her x^{k-1} -best share in this coalition.

Definition 4.1 (Algorithm \mathcal{A}). At each step $s = 1, 2, \dots$,

1. for each type whose cutoff has been determined in a step $s' < s$, let x_i^s be her payoff; for all other types j , let $x_j^s = 0$,
2. let the x^s -best share of each type in an x^s -perfect coalition be her payoff.

Stop in the first step S in which everyone's payoffs have been pinned down, and let $\chi := x^S$.

The reasoning analogous to the one behind [Proposition 4.4](#) shows that, at each step s of algorithm \mathcal{A} , there is always at least one x^s -perfect coalition that is not an $x^{s'}$ -perfect coalition for any $s' < s$. Hence, each step identifies the payoff of at least one type. Since there are finitely many types, this guarantees that algorithm \mathcal{A} ends after finitely many steps. [Proposition 4.5](#) culminates the characterization argument described so far.

Proposition 4.5. *The payoff profile χ defined by algorithm \mathcal{A} is the only one that solves system (2).*

Figure 5 illustrates the limit equilibrium payoffs in the example of section 2. I now turn to showing that—as the bargaining friction q converges to zero—the equilibrium cutoff profile χ converges to the unique payoff profile that is consistent with two principles, roughly stated as follows: The Nash bargaining solution determines the sharing rule in each coalition, and outside options cannot be determined in a circular way.

4.5 Main result

Theorem 1 below highlights that the equilibrium cutoff profile w of the bargaining game \mathcal{G} converges—as the bargaining friction q vanishes—to the unique payoff profile that gives each type the maximum that she can justify using the Nash bargaining solution in some coalition subject to *the others'* payoffs. Definition 4.2 provides the necessary concepts to formalize this idea. For brevity, for each profile $x \in \mathbb{R}^N$, I denote by x_{-i} the profile x after setting its i^{th} entry to zero.

Definition 4.2. Fix a coalition $C \subseteq N$ and a payoff profile x in \mathbb{R}^N .

1. The profile of x -Nash bargaining shares in coalition C is the solution to (8). In words, for each type i in coalition C , i 's x -Nash bargaining share in C is her Nash bargaining share in this coalition (with threat points given by autarky) subject to the constraint that each of its members j obtains at least x_j .
2. Type i 's x -Nash best share is her maximum x -Nash bargaining share across all coalitions.
3. Coalition C is one of type i 's x -Nash best coalitions if her x -Nash bargaining share in it is equal to her maximum x -Nash bargaining share across all coalitions.
4. Coalition C is x -Nash perfect if it is an x -Nash best coalition of all of its members.
5. The payoff profile x is Nash credible if, for each type i , x_i is her x_{-i} -Nash best share. In words, x_i is the maximum that i can justify as the result of Nash bargaining in some coalition subject to the others' outside options—given by x and entering the Nash bargaining solution as bounds on the payoffs, as prescribed by the outside option principle.

Theorem 1 below shows that there always exists a Nash-credible payoff profile, that this profile is always unique, and that the equilibrium cutoff profile converges to it in the limit as the bargaining friction q converges to 0.

Theorem 1. *In the limit as the bargaining friction q vanishes, the equilibrium cutoff profile of the bargaining game \mathcal{G} converges to the unique Nash-credible payoff profile.*

For instance, [Figure 6](#) illustrates how each type's limit payoff in the example of [section 2](#) is the maximum (across her two potential coalitions) that she is able to justify using the Nash bargaining solution while honoring the others' payoffs. [Theorem 1](#) implies that this is the unique payoff profile that satisfies this property—that is, this is the Nash credible payoff profile in this example.

[Proposition 4.6](#) below shows how the algorithm that I now define identifies the unique Nash-credible payoff profile. This algorithm is analogous to algorithm \mathcal{A} using the notion of x -Nash-best shares and x -Nash-best coalitions instead of x -best shares and x -best coalitions.

Definition 4.3 (Algorithm \mathcal{A}^*). At each step $s = 1, 2, \dots$,

1. for each type whose payoff has been determined in any step $s' < s$, let x_i^s be her payoff; for all other types j , let $x_j^s = 0$,
2. let the x^s -Nash best share of each type in an x^s -Nash perfect coalition be her payoff.

Stop in the first step S such that everyone's payoff has been determined, and let $\chi^* := x^S$.

At each step s of algorithm \mathcal{A}^* , there is always at least one x^s -perfect coalition that is not an $x^{s'}$ -perfect coalition for any $s' < s$. Hence, each step identifies the payoff of at least one type. Since there are finitely many types, this guarantees that algorithm \mathcal{A}^* ends after finitely many steps. [Proposition 4.6](#) highlights that the payoff profile defined by this algorithm is the unique Nash-credible payoff profile.

Proposition 4.6. *The payoff profile χ^* defined by algorithm \mathcal{A}^* is the unique Nash-credible payoff profile.*

The reasoning behind [Proposition 4.6](#) is analogous to the one provided in subsections [4.2](#), [4.3](#) and [4.4](#) above: Starting with the profile of outside options x being $\mathbf{0}$, there always exists at least one x -Nash perfect coalition, which defines the payoffs of all of its members in every Nash-credible payoff profile. After updating the payoff profile x accordingly, there is at least an x -Nash perfect coalition that contains a type whose equilibrium cutoff has not yet been defined. Such a type's x -Nash bargaining share in such a coalition defines her payoff in every Nash-credible profile. The process ends after finitely many steps, when the unique Nash credible payoff profile has been identified.

An alternative way to read [Theorem 1](#) is that the equilibrium sharing rule in each coalition C converges to the Nash bargaining solution in that coalition:

$$(7) \quad \operatorname{argmax}_{s \in \mathbb{R}^C} \prod_{j \in C} [u_j(s_j) - u_j(y_j)]^{p_j} \quad \text{subject to} \quad y(C) \geq \sum_{j \in C} s_j \quad \text{and} \quad s_j \geq x_j^C \quad \text{for all } j \in C.$$

with each type i 's outside option x_i^C in C being the maximum, across all coalitions $D \neq C$, of the i^{th} element of

$$(8) \quad \operatorname{argmax}_{s \in \mathbb{R}^D} \prod_{j \in D} [u_j(s_j) - u_j(y_j)]^{p_j} \quad \text{subject to} \quad y(D) \geq \sum_{j \in C} s_j \quad \text{and} \quad s_j \geq x_j^D \quad \text{for all } j \in D - i.$$

In words, each type's outside option in any given coalition is the maximum that she can justify using the Nash bargaining solution *in some other coalition* while honoring to the others' outside options in that coalition.³¹ Indeed, [Theorem 1](#) implies that—when the bargaining friction q is small enough—in equilibrium each type obtains (close to) her Nash bargaining share in one coalition subject to the others' outside options. This determines her binding outside option in all the other coalitions that contain her. For instance, in the example of [section 2](#), the binding outside options of the chefs when bargaining with the managers are determined by the Nash bargaining solution in the chef-maître coalition (without invoking the chefs' outside option there). Similarly, the binding outside options of the managers when bargaining with the cooks are determined by the Nash bargaining solution in the chef-manager coalition (without invoking the managers' outside option there).

5 Comparative statics

In the previous section, I have described how, as the bargaining friction q vanishes, *the unique equilibrium cutoff profile w in the bargaining game \mathcal{G} converges to the unique Nash-credible payoff profile χ^** . I now turn to describing how this payoff profile changes as a function of the primitives of the model.³² I start, in [subsection 5.1](#), by describing the tier structure of the set of coalitions that form in equilibrium—which is behind most of the comparative statics that follow. I then turn to describing general comparative statics in [subsection 5.2](#), comparative statics in the special case of vertically differentiated markets in [subsection 5.3](#), and the implications of positive assortative matching on the nature of shock propagation in [subsection 5.4](#).

³¹In particular, the binding outside options in each coalition are determined in coalitions that form in equilibrium.

³²For simplicity, I focus on the limit as the bargaining friction q vanishes, but analogous comparative statics hold for any bargaining friction $q > 0$.

5.1 The equilibrium tier structure

The coalitions that form in equilibrium can be organized into tiers, in such a way that the surplus in each coalition $C \subseteq N$ is shared as prescribed by the Nash bargaining solution, with the binding outside options determined in higher tiers. As a result, (small) shocks to the primitives propagate—via agents' outside options—from higher to lower tiers, but not vice versa. This structure provides the theory with rich testable predictions.

Formally, the *first-tier coalitions* are those coalitions $C \subseteq N$ that are such that, for each of its members i , type i 's χ_{-i}^* -Nash best share is equal to χ_i^* .³³ In other words, the first-tier coalitions are those coalitions $C \subseteq N$ that are such that, for each of its members $i \in C$,

1. the Nash bargaining solution in the coalition C subject to honoring the payoff χ_j^* of each $j \in C - i$ is equal to χ_i^* , and
2. the Nash bargaining solution in any other coalition $D \neq C$ subject to honoring the payoff χ_j^* for all $j \in D - i$ is smaller than χ_i^* .

The sharing rule in the first tier coalitions converges to the Nash bargaining solution, with threat points given by autarky, and without any binding outside options. Intuitively, no member of a first tier coalition can make a credible threat to propose to a different coalition, so it is *as if* the members of such coalitions where bargaining in isolation. For instance, [Figure 10](#) illustrates an example that has two first-tier coalitions. The surplus in these coalitions is shared according to the Nash bargaining solution, with the threat points given by autarky, and without any binding outside options. The *first tier types* are those that are members of a first tier coalition.

Proceeding inductively, after having identified, for every k in $\{1, 2, \dots, \ell - 1\}$, the k^{th} tier coalitions, a coalition $C \subseteq N$ is in the ℓ^{th} tier if and only if (i) it contains at least one $(\ell - 1)^{\text{th}}$ tier type and (ii) is such that, for each of its members i who is not in the first, second, \dots , or $(\ell - 1)^{\text{th}}$ tier, type i 's χ_{-i}^* -Nash best share is equal to χ_i^* . In other words, the ℓ^{th} tier coalitions are those that are such that each of its members who is not in a higher-tiered coalition (i) can justify her payoff as resulting from the Nash bargaining solution in this coalition subject to the others' payoffs, and (ii) cannot justify in this way more than this payoff in any other coalition. The surplus in each ℓ^{th} tier coalition is shared according to the Nash bargaining solution, with the threat points given by autarky and the binding outside

³³Note that there is always at least one first tier coalition. In particular, every Nash-perfect coalition is in the first tier.

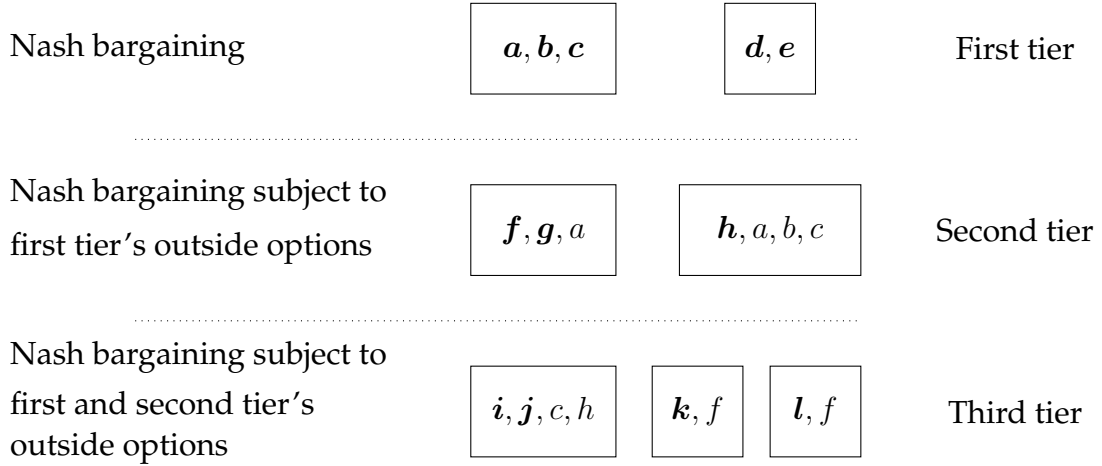


Figure 10: Illustration of the equilibrium tier structure. The types are indexed from a to l . A box containing different types represents an equilibrium coalition containing these types. The index of type i is in bold in the coalition that gives her the highest χ_{-i}^* -Nash best share; this is also the coalition that she proposes to in equilibrium when the bargaining friction q is small enough.

options determined in higher tiers. The ℓ^{th} tier types are those that are in an ℓ^{th} -tier coalition and are not in any k^{th} tier coalition, for any $k < \ell$.

The tier structure of the set of equilibrium coalitions implies that a type's payoff is not affected by shocks to the primitives that only hit coalitions and types in lower tiers (as long as these shocks are sufficiently small so that they do not change the equilibrium tier structure). Moreover, an increase in the productivity of an ℓ^{th} tier coalition is never beneficial for $\ell + 1^{\text{th}}$ tier types, because it can only increase their partners' outside options. For instance, in the example illustrated in Figure 10, an increase in the surplus of the coalition $\{a, b, c\}$ decreases the payoffs of all the second tier types, because it increases the outside options of the first tier types in the second tier coalitions. This leads, in turn, to an increase in the payoffs of the third tier types k and l . The effect on the payoffs of the other third tier types is ambiguous in this case, because one of the outside options that they have to honor (that of type c) increases, but another one (that of type h) decreases.

As another example, Figure 7 illustrates the tier structure in the example of section 2. In this case, an increase in the productivity of the first tier coalition (the chef-maitre coalition) hurts the second tier type (managers), because it increases the outside options of the first tier type (chefs) that managers have to honor. In contrast, such an increase is beneficial for third tier types (cooks). This suggests that, in certain cases, a positive productivity shock to

a ℓ^{th} tier coalition affects negatively (positively) the $\ell + m^{\text{th}}$ tier types when m is odd (even). Indeed, one can construct the *coalitional overlap network*—whose nodes are all the coalitions that form in equilibrium, and where a link between two coalitions represents the fact that they share at least one type. Then, if there is only one path in the coalitional overlap network from one coalition to another one, a marginal increase in the surplus of one does not hurt (benefit) the members of the other if the path is even (odd).³⁴

5.2 General comparative statics

While the equilibrium tier structure described in the previous subsection is useful to illustrate how shocks propagate via outside options, large shocks will typically affect this structure. Hence, understanding how these shocks propagate requires investigating how they affect the equilibrium tier structure itself. I now turn to discussing several results in this direction, in the form of comparative statics that hold for arbitrarily large shocks to the primitives.

Corollary 5.1 below highlights that each type’s payoff is increasing in her own productivity and bargaining power, and decreasing in her own risk aversion.

Corollary 5.1. *A type’s payoff increases when the surplus of any coalition that she is part of increases, when her bargaining power increases, or when she becomes less risk averse.*³⁵

Intuitively, the maximum payoff that a type can justify using the Nash bargaining solution in some coalition while honoring the others’ payoffs can only increase when a coalition that she is part of becomes more productive, when her bargaining power increases, or when she becomes less risk averse. In the case of small shocks, this can be seen from the equilibrium tier structure, because the outside options that each type has to honor in the coalition where she can justify her payoff are independent of the surplus of any coalition that she is part of, her preferences, and her bargaining power. More generally, **Corollary 5.1** follows from the observation that the outside option x_j^s of any type j at the step s at which type i ’s payoff is determined by algorithm \mathcal{A}^* cannot increase as a result of her becoming more productive, her bargaining power increasing, or her becoming less risk averse.

While **Corollary 5.1** seems intuitive, it contrasts with well-known theories of coalition formation (e.g., **Farrell and Scotchmer 1988** and **Pycia 2012**) where agents bargain first over

³⁴A path of a network is a sequence of distinct links. A path is even (odd) if it contains an even (odd) number of links.

³⁵I say that a type whose utility function changes from u to w has become more risk averse if there exists a concave function g such that $w = g \circ u$.

which coalitions to form, and second over their terms of trade. For example, Pycia (2012, p. 347) describes how a holdup problem can occur in these models. Somewhat paradoxically, this implies that an agent can be worse off when she becomes more productive, her bargaining power increases, or she becomes less risk averse. In his own words,

Inflexible sharing of surplus leads to holdup in coalition formation. For instance, consider the setting in which agents share surplus in Nash bargaining with constant bargaining powers. An agent may be better off with a lower rather than a higher bargaining power—other things held equal—when a low bargaining power allows him or her to form a highly productive coalition, while a high bargaining power makes formation of such a productive coalition impossible by lowering the payoffs of its members below their outside options.

Corollary 5.1 above highlights that such a holdup problem does not arise in the present setting. However, in contrast to the Nash bargaining solution in a fixed coalition (with exogenous outside options), the payoff of an agent can increase when others' bargaining power increases, when others' risk aversion decreases, and when a productivity shock increases others' outside options. For instance, as illustrated by the example of section 2, cooks benefit when the surplus of the chef-maitre coalition increases. Indeed, this improves the chefs' outside options when bargaining with managers which, in turn, deteriorates the managers' outside options when bargaining with cooks.

Corollary 5.2 below highlights that when two types are *perfect complements*—in the sense that one is useless without the other—any shock that affects one also affects the other. Moreover, an increase in the risk aversion or productivity of one has positive spillovers on the other, but an increase in the bargaining power of one has negative spillovers on the other.³⁶

Corollary 5.2. *Consider two types that are perfect complements. An increase in the productivity or risk aversion of one strictly increases the payoff of the other. Similarly, an increase in the bargaining power of one strictly decreases the payoff of the other.*

Intuitively, when two types are perfect complements, their maximum Nash bargaining share—across all coalitions—subject to the others' outside options is determined in the same coalition. As a result, when one of these types becomes more productive, the resulting increase in surplus is shared among these types (and possibly others as well). Similarly, the fact that these two types justify their payoff in the same coalition implies that, when one of

³⁶Formally, two types i and i' are perfect complements if, for any coalition $C \subseteq N - \{i, i'\}$, we have that $y(C) = y(C \cup \{i\}) = y(C \cup \{i'\}) < y(C \cup \{i, i'\})$.

these types becomes more risk averse or her bargaining power decreases, the amount that the other can justify in some coalition using the Nash bargaining solution while honoring the others' payoffs increases.

Corollary 5.3 below highlights that when two types are *perfect substitutes*—in the sense that their marginal contribution to any coalition that does not contain either of them is the same—preference and bargaining power shocks propagate from the most risk averse to the least risk averse one, but not vice versa; and from the least to the most powerful, but not vice versa.³⁷

Corollary 5.3. *Consider two types that are perfect substitutes. If one is more risk averse than the other (and remains so after the change under consideration), an increase in the risk aversion of the most risk averse one can affect the payoff of the other one, but not vice versa. Similarly, if one has a higher bargaining power than the other (both before and after the change under consideration), an increase in the bargaining power of the one with the least bargaining power can affect the payoff of the other one, but not vice versa.*

Indeed, when two types are perfect substitutes, the payoff of the most risk averse one or the one with the least bargaining power is determined by algorithm \mathcal{A}^* first. This means, in particular, that the payoff of the most risk averse type i is not affected by the risk aversion of the other type i' (as long as i' remains less risk averse than i). But an increase in the risk aversion of the most risk averse type can affect the binding outside options of the equilibrium partners of the other (because both of them might match with the same types in equilibrium, for example), and hence affecting, in turn, her payoff.

5.3 Vertically differentiated markets

I now turn to deriving further comparative statics in the context of two-sided pairwise matching markets, in which types are vertically differentiated either by their skill, by their risk aversion, or by their bargaining power. For ease of exposition, I refer to the types on one side as *workers* and the types on the other side as *employers*. Throughout the rest of this section, I assume that each type i is endowed with a *risk aversion parameter* r_i and a *skill (or productivity) parameter* s_i . For each worker-employer pair (i, j) , the surplus $y(i, j)$ is strictly increasing in its members' skills.³⁸ I say that two types i and j *match* if agents of type i match

³⁷Formally, two types i and i' are perfect substitutes if, for any $C \subseteq N - \{i, i'\}$, we have $y(C \cup \{i, i'\}) = y(C \cup \{i\}) = y(C \cup \{i'\}) > y(C)$.

³⁸Formally, let i and i' be any two workers and let j be any employer. We have that $y(i, j) > y(i', j)$ if and only if $s_i > s_{i'}$, and there exists a concave function g such that $u_i = g \circ u_{i'}$ if and only if $r_i > r_{i'}$.




Figure 11: A two-sided market where agents are vertically differentiated by their skills. Each node corresponds to a type. The nodes on the left column represent one side of the market, and the nodes on the right side the other. The surplus that two types i and j in opposite sides of the market generate is the product of i and j . Preferences are homogeneous. A link between node i and node j indicates that i and j match in equilibrium. The number next to a node is her limit equilibrium payoff. An arrow from type i to type j indicates that type i can justify her payoff using the Nash bargaining solution in the ij coalition while honoring j 's payoffs.

with agents of type j in equilibrium.

For example, Figure 11 illustrates a two-sided market whose members are vertically differentiated by their skill. In this example, the set of workers is $\{1, 2, 3, 4, 5, 8, 9, 10\}$ and the set of employers is $\{2, 6, 10\}$. The surplus of each coalition is the product of the labels of the associated types. Figure 11 illustrates the coalitions that form in equilibrium, and Figure 12 illustrates the corresponding tier structure in this example.

Corollary 5.4 below illustrates how—in settings where the agents are vertically differentiated by their productivity but are otherwise identical—the bargaining outcomes are determined from the most productive types down. For instance, in the example illustrated in Figure 11, algorithm \mathcal{A}^* pins down, in the first step, the payoffs of the types with the highest skill on each side, and then works downwards from there to identify the payoffs of types 9




Figure 12: Equilibrium tier structure in the example illustrated in Figure 11. A type’s payoff is determined in the match where her name is in bold. The dotted line indicates how the tier structure can be split into two non-overlapping tier structures, in the sense that each type is either in one or the other.

and 8 in the second step, the payoff of type 6 in the third step, and so on.

Corollary 5.4. *When all the workers have the same risk aversion and bargaining power, an increase in the productivity of worker i from s_i to $s'_i > s_i$ does not affect the payoff of any worker j whose productivity s_j is strictly higher than s'_i .*

For instance, in the example illustrated in Figure 11, an ϵ increase in the skill of worker 8, where $0 < \epsilon < 2$, increases worker 5’s payoff, but it does not affect the payoffs of workers 9 or 10. Intuitively, when all the types have the same risk aversion, algorithm \mathcal{A}^* does not determine a worker’s payoff before determining the payoffs of all the more productive workers. This implies that an increase in a type’s productivity does not affect the payoffs of the more productive types.

Corollary 5.4 implies that when the source of heterogeneity among agents is their productivity, the bargaining outcomes are determined *from those with the highest payoffs down*. In contrast, Corollary 5.5 shows that, when the source of heterogeneity among agents is either their risk aversion or their bargaining power, the bargaining outcomes are determined *from those with the lowest payoffs up*.

Corollary 5.5. *When all the workers have the same skill and bargaining power, an increase in the risk aversion of worker i from r_i to $r'_i > r_i$ does not affect the payoff of any worker whose risk aversion is strictly higher than r'_i . Similarly, when all the workers have the same skill and risk aversion, an increase in the bargaining power of worker i from p_i to $p'_i > p_i$ does not affect the payoff of any worker whose bargaining power is strictly lower than p'_i .*




Figure 13: Equilibrium matching pattern in the example when the surplus of the cook-maître coalition is 80 instead of 60.

Intuitively, when all the types have the same skill and bargaining power, algorithm \mathcal{A}^* only determines a worker’s payoff once it has determined the payoffs of all the more risk averse workers. Similarly, when all the types have the same skill and risk aversion, algorithm \mathcal{A}^* only determines a worker’s payoff when it has determined the payoffs of all the less powerful agents.

5.4 Shock propagation under positive assortative matching

For simplicity, I now focus on settings in which the only source of heterogeneity is skill. I say that the equilibrium features *positive assortative matching* if, for any two worker types i and i' and any two employer types j and j' , with $s_i > s_{i'}$ and $s_j > s_{j'}$, if i matches with j' in equilibrium, then i' does not match with j in equilibrium. For example, the equilibrium matching pattern in the example of [section 2](#) features positive assortative matching. In contrast, when this example is modified to have cooks and maîtres generate 80 units of surplus instead of 60, the equilibrium matching pattern does not feature positive assortative matching, because the most productive cooks (i.e., chefs) match with the lowest productive managers, while the least productive cooks match with the most productive managers (i.e., maîtres). [Figure 13](#) illustrates the matching pattern in this case.

[Pycia \(2012\)](#) shows that when agents bargain over which coalitions to form understanding that the sharing rule within coalitions will be determined by the Nash bargaining solution (with exogenous outside options), the notion of *stability*—in the sense of the *core*—implies that agents match in a positive assortative way with respect to their productivity and their risk aversion. As the example illustrated in [Figure 13](#) illustrates, this is not necessarily the case in the present setting, where agents bargain simultaneously over both which coalitions to form and how to share the resulting surplus.

In the context of his landmark investigation of marriage markets, [Becker \(1973\)](#) showed that—in the context of one-to-one matching environment with a finite set of agents and

transferable utility (the *assignment game*)—agents match in a positive assortative way for all distribution of types if and only if their skills are complementary—in the sense that the match function is supermodular. [Corollary 5.6](#) is the analogous result in the present setting.

Corollary 5.6. *The equilibrium features positive assortative matching for all distribution of types if and only if the production function is strictly supermodular.*³⁹

For example, the surplus function in the example illustrated in [Figure 11](#) is strictly supermodular, so the resulting matching pattern features positive assortative matching. To see the intuition behind [Corollary 5.6](#), consider again the situation illustrated in [Figure 13](#). As argued above, this matching pattern does not feature positive assortative matching, so it must be the case that the production function is not strictly supermodular. Indeed, strict supermodularity in this case boils down to the following condition

$$y(\text{Chefs, Maîtres}) + y(\text{Cooks, Managers}) > y(\text{Chefs, Managers}) + y(\text{Cooks, Maîtres}).$$

When this condition is satisfied, we cannot have the matching pattern illustrated in [Figure 13](#), because—given that cooks and managers do not match in equilibrium—it must be that the sum of their payoffs is strictly bigger than the surplus they generate together. That is,

$$y(\text{Cooks, Managers}) < \underbrace{y(\text{Chefs, Managers}) + y(\text{Cooks, Maîtres})}_{\text{Sum of all four types payoffs}} - \underbrace{y(\text{Chefs, Maîtres})}_{\text{Sum of Chefs and Maîtres payoffs}}.$$

More generally, this argument shows that (i) when the equilibrium does not feature positive assortative matching, the production function is not strictly supermodular, and that (ii) when the production function is not strictly supermodular, a set of four types can be constructed so that their equilibrium matching pattern is not positive assortative.

[Corollary 5.4](#) above highlighted that, in settings where the types are vertically differentiated by their skills, each type is only affected by shocks that hit more productive types. [Corollary 5.7](#) below puts further structure on how shocks propagate under positive assortative matching.

Corollary 5.7. *Consider an increase in the skill of worker i from s_i to s'_i , and assume that the equilibrium features positive assortative matching both before and after this shock. If this shock does not affect the payoff of a worker j with $s_j < s_i$, then it does not affect the payoff of any worker j' with $s_{j'} \leq s_j$ either.*

³⁹The production function is strictly supermodular if, for any two buyer types i and i' and two seller types j and j' , with $s_i > s_{i'}$ and $s_j > s_{j'}$, we have that $y(i, j) + y(i', j') > y(i', j) + y(i, j')$.

In other words, under positive assortative matching, shocks propagate *in blocks*, in the sense that if a shock that affects worker i propagates to worker i' , it also affects every worker whose skill is in between. For instance, in the example illustrated in [Figure 11](#), an increase in the skill of worker 10 affects the payoff of worker 4. Given that the matching is positive assortative, this implies that such an increase affects every worker whose skill is between 4 and 10 as well.

Intuitively, under positive assortative matching, algorithm \mathcal{A}^* pins down the workers' (employers') equilibrium payoffs in sequence—from the most skilled down. As a result, in the scenario in which worker i becomes more productive and worker j , with $s_j < s_i$, is not affected, the output of the algorithm \mathcal{A}^* from the step at which it pins down j 's payoff on is not affected by this shock. In order to gain more intuition for this observation, consider the case in which the shock is sufficiently small so that it does not affect the equilibrium tier structure. If an increase in the skill of a worker i from s_i to s'_i does not affect the payoff of another worker j with $s_j < s_i$, it means that there is no path from i to j in the matching network (which has a link from one type to another if they match in equilibrium). When the equilibrium matching is positive assortative, this implies that there is no path to i from any other worker k less productive than j either, which implies, in turn, that such an increase cannot affect k 's payoff either.

The situation described by [Corollary 5.7](#) contrasts with the case in which the production function is strictly submodular.⁴⁰ Indeed, in this case, all the types that match do so with the most productive type on the other side of the market.⁴¹ As a result, everyone's payoffs depend on the most productive types. In particular, in this case, an increase in the surplus of the match among the two most productive types makes everyone other than these types worse off. However, an increase in the surplus of any other match that includes the highest type is fully absorbed by the low type of the match. For example, when the equilibrium is as in [Figure 13](#), an increase in the surplus of the chef-manager coalition from 80 to 90 is fully absorbed by the managers.

⁴⁰The production function is strictly submodular if, for any two buyers i and i' and any two sellers j and j' , with $s_i > s_{i'}$ and $s_j > s_{j'}$, we have that $y(i, j) + y(i', j') < y(i', j) + y(i, j')$.

⁴¹To see this in the context of the example, suppose for contradiction that the surplus function is submodular and that chefs and maîtres match with each other, and so do cooks and managers. Then, the sum of the payoffs of these four types is equal to the sum of the surpluses of these two types of coalitions, which, by submodularity, is smaller than the sum of the surpluses generated by a chef-manager coalition and a cook-maître coalition. In particular, at least one of these coalitions is more profitable than the sum of its members' payoffs, a contradiction.

6 Related literature

This paper is related to two complementary branches of the literature on non-cooperative coalition formation—one that studies *stationary markets* (e.g., Rubinstein and Wolinsky 1985, Gale 1987, Binmore and Herrero 1988, Fraja and Sákovics 2001, Manea 2011, Lauer mann 2013, Nguyen 2015, Polanski and Vega-Redondo 2018) and the other one that studies *non-stationary markets* (e.g., Chatterjee, Dutta, Ray, and Sengupta 1993, Moldovanu and Winter 1995, Ray and Vohra 1999, Corominas-Bosch 2004, Gale and Sabourian 2006, Ray 2007, Okada 2011, Abreu and Manea 2012ab, Elliott and Nava 2018).⁴²

These two lines of research focus on two opposite extremes. On the one hand, in order to investigate the strategic forces in a steady state of large dynamic markets, the former typically assumes that *the inflow of traders into the market perfectly matches its outflows*. On the other hand, in order to investigate the consequences of the endogenous evolution of the set of active traders over time, the latter typically assumes that *there are no inflows of traders into the market*. Hence, roughly speaking, these two lines of research focus on the likely predominant strategic forces in *thick* and *thin* markets, respectively.

This paper's main contribution is to the literature on bargaining in stationary markets. The main innovation of this paper with respect to this literature is that I allow the agents to strategically choose which coalitions to propose, which is an essential feature behind the connection between the predictions of the non-cooperative model and the Nash bargaining solution.⁴³ For example, Nguyen (2015) uses convex programming techniques to characterize the stationary subgame-perfect equilibrium of a non-cooperative bargaining game similar to the one in the present paper, but the equilibrium payoff profile in his framework cannot be understood in terms of classical bargaining theory.⁴⁴ This contrasts with the present

⁴²Also related, albeit with a somewhat different spirit, is the large literature that builds on legislative bargaining model of Baron and Ferejohn (1989) (e.g., Eraslan 2002, Eraslan and McLennan 2013 and Eraslan 2016).

⁴³In Talamàs (2018), I study networked buyer-seller markets using a framework similar to the one in the present paper, and I discuss how allowing the agents to strategically choose whom to make offers to fundamentally alters the determinants of price dispersion in these markets. The characterization in Talamàs (2018) is related to the one of the present paper. There, I provide a simple necessary and sufficient condition for the law of one price to hold in equilibrium (in the limit as bargaining frictions vanish), and I describe an algorithm that decomposes the buyer-seller network into submarkets, from the submarket with the highest limit price down (or, alternatively, the submarket with the lowest limit price up).

⁴⁴The main difference between the model in Nguyen (2015) and the one that I study in the present paper is that, in the former, coalitions are proposed at random (instead of being selected strategically by the proposer) and, in each period, its members bargain only over whether to form the proposed coalition and, if so, their terms of trade. Another difference is that, in Nguyen (2015), all the types have linear utilities, while I assume

paper, where I show that the (limit) equilibrium payoff profile is the unique profile that satisfies the following property: Each type obtains the maximum that she can justify using the Nash bargaining solution in some coalition while honoring the others' equilibrium payoffs. As I illustrate in this paper, this provides a tractable theory of coalition formation that allows general comparative statics.

Despite the contrast between the settings under consideration, the strategic forces in the present paper are most related to those that arise in the literature investigating non-cooperative bargaining in non-stationary markets. Indeed, the structure of the equilibrium is similar to that of the *no-delay* perfect equilibrium characterized in Chatterjee, Dutta, Ray, and Sengupta (1993). Other than the setting, the main difference is that—as in the classical bargaining framework of Nash (1950)—I allow the agents to have arbitrary vN-M utility functions instead of restricting attention to the case of linear utilities. As a consequence, in the limit as the bargaining friction vanishes, in the present setting the coalitional surpluses are shared according to the Nash bargaining solution—instead of according to equal sharing. Moreover, the dynamic entry of the agents into the market implies that the structure of the equilibrium in the present setting is fixed throughout—instead of evolving as different coalitions form—and that the equilibrium that I characterize always exists and is the unique stationary perfect equilibrium—instead of being one of the possible stationary perfect equilibria and existing only under certain conditions. Not surprisingly given the qualitative differences between the settings, however, the predictions of the resulting theories are qualitatively distinct. For example, the endogenous evolution of the market in Chatterjee, Dutta, Ray, and Sengupta (1993) implies that—unlike in the present setting—an agent does not necessarily benefit when she becomes more productive (because her improved outside option can lead others to avoid making offers to her, which can in turn make it more likely that the market will evolve against her).

The similarity between the strategic forces in Chatterjee, Dutta, Ray, and Sengupta (1993) and the present paper might seem surprising given that these two papers differ not only in the setting under study but also in the bargaining protocol. Indeed, they use a *rejector-proposes protocol* (in which the rejector of a proposal becomes the proposer in the next period) instead of the *random-proposer protocol* of the present paper. Ray (2007) (see also Compte and Jehiel 2010) discusses how the former protocol gives considerably more bargaining power to the receiver of the offer than the latter, and how this explains the contrasting predictions often obtained under these two protocols. Intuitively, however, the dynamic entry of agents

instead that—as in the classical framework of Nash (1950)—each type's preferences can be represented by a concave vN-M utility function.

into the market featured in the present paper implies that agents do not have to consider how the market might evolve after they reject an offer, which implies that the difference between these protocols is much less pronounced—and it actually vanishes as the bargaining frictions go to zero.

In the context of *convex games*,⁴⁵ Chatterjee, Dutta, Ray, and Sengupta (1993) show that the prediction of the no-delay stationary perfect equilibrium of their coalition formation game converges—as the bargaining frictions vanish—to the *egalitarian solution* of Dutta and Ray (1989). The similarity between the structure of the equilibrium in their game and the one in the present paper suggests that their results can be generalized beyond the case of linear utilities and can more generally be understood in terms of the Nash bargaining solution. Relatedly, Compte and Jehiel (2010) focus on environments in which the grand coalition generates the highest surplus, and in which only one coalition may form. They show that, if an (asymptotically) efficient stationary equilibrium exists, the corresponding profile of payoffs is the one that maximizes the product of agents' payoffs among those in the core.⁴⁶ The analysis of the present paper suggests that this result generalizes to the case in which agents' have vN-M utility functions, as follows: If an (asymptotically) efficient stationary perfect equilibrium exists, the corresponding profile of payoffs is the unique payoff profile that gives each agent her maximum Nash bargaining share among all the possible coalitions subject to the others' outside options.

This paper is further related to two other lines of research. First, the idea of building a theory of coalition formation from the Nash bargaining solution goes back at least to Rochford (1984), who defines a *symmetrically-pairwise-bargained* payoff profile of an assignment game with transferable utility as one that satisfies the following property: Each matched pair shares output according to the Nash bargaining solution—with each agent's disagreement point being the maximum that she can achieve in any other match (keeping the others' payoffs fixed).^{47,48} Burguet and Caminal (2018) show that a modification and extension of this

⁴⁵A game is *convex* if, for any two coalitions C_1 and C_2 , $y(C_1 \cup C_2) \geq y(C_1) + y(C_2) - y(C_1 \cap C_2)$.

⁴⁶They also provide a necessary and sufficient condition for such an equilibrium to exist. This study contrasts with Krishna and Serrano (1996), who establish a connection between the Nash bargaining solution and a strategic bargaining game in which the grand coalition is the only one that can form.

⁴⁷Rochford (1984) shows that the set of symmetrically-pairwise-bargained payoff profiles is the intersection of the *kernel* and the *core*. Kleinberg and Tardos (2008) refer to such profiles as “balanced outcomes.” Alternative approaches to select a point from the core of the assignment game include Kranton and Minehart (2001) (who focus on an extreme point of the core) and Elliott (2015) (who focuses on different convex combinations of the extreme points of the core).

⁴⁸The idea of endogenizing agents' “threat points” was pursued by Nash (1953) himself (see also Binmore 1987 and Abreu and Pearce 2015) and it is the essence of well-known consistency notions (e.g. Sobolev 1975,

idea (in a context in which only one coalition can form) uniquely pins down the agents' pay-offs, and provide strategic foundations for the resulting coalition formation solution concept. While these concepts are similar in spirit to the one described in the present paper, the non-cooperative approach described here suggests that—in the setting of this paper—the outside option principle holds, so outside options do not enter through disagreement points, but act instead as bounds on the range of validity of the Nash bargaining solution.⁴⁹

Second, [Collard-Wexler, Gowrisankaran, and Lee \(forthcoming\)](#) provide strategic foundations for the *Nash equilibrium in Nash bargains* ([Horn and Wolinsky 1988](#)), which is a widely used bargaining solution concept for bilateral oligopoly settings. In contrast to the coalition formation approach of the present paper—in which each agent can be part of at most one coalition—the *Nash-in-Nash* solution assumes that all the parties trade with each other (i.e., that all possible coalitions form) and derives prices for each bilateral contract as a function of the fundamentals. [Ho and Lee \(forthcoming\)](#) provide a modification of the Nash-in-Nash solution to investigate agents' incentives to restrict the set of agents with whom they trade.

7 Conclusion

The *Nash bargaining solution* is a central solution concept in economics. Nash proposed this solution concept using an axiomatic approach. In his own words ([Nash, 1953](#), p. 129),

One states as axioms several properties that it would seem natural for the solution to have and then one discovers that the axioms actually determine the solution uniquely.

My objective in this paper has been to investigate how prices and allocations are determined in dynamically-thick matching markets, in which agents bargain both about which coalitions to form and how the resulting surplus is shared within them. One possible approach to fulfill this objective is to extend Nash's axioms to this setting, and then to discover what solution comes out of these axioms.

In this paper, I have taken an alternative approach, which leverages the celebrated connection between non-cooperative bargaining and the Nash bargaining solution: I have extended the canonical non-cooperative bargaining model that connects with the Nash bargaining solution (e.g., [Binmore, Rubinstein, and Wolinsky 1986](#)) to the setting of interest, and I have let

[Peleg 1986](#), [Hart and Mas-Colell 1989](#), [Serrano and Shimomura 1998](#)).

⁴⁹In [Agranov, Elliott, and Talamàs \(in preparation\)](#), we experimentally investigate, among other things, the empirical validity of this prediction.

this model suggest how the Nash bargaining solution generalizes to this setting. The payoff profile under the resulting theory of coalition formation is the unique profile that is such that each agent's payoff is her maximum Nash bargaining share among all the coalitions subject to honoring the others' payoffs. An interesting avenue for future research is to investigate whether Nash's axioms have natural analogs in the framework of this paper that pin down this solution.

This paper suggests a handful of exciting directions for future research. First, as discussed in [section 1](#), the contribution of this paper can be seen as enriching the outside option principle of [Binmore, Rubinstein, and Wolinsky \(1986\)](#) to determine, not only how outside options enter the Nash bargaining solution, but also *how they are determined by the Nash bargaining solution*. [Binmore, Shaked, and Sutton \(1989\)](#) and [Jäger, Schoefer, Young, and Zweimüller \(2018\)](#) provide empirical evidence that is credible with the outside option principle. In [Agranol, Elliott, and Talamàs \(in preparation\)](#), we design and implement a new experimental approach that replicates the relevant strategic forces in stationary markets, and we use it to investigate the empirical relevance of the qualitative predictions that emerge from this theory.

Second, the theory that emerges from this paper opens a door to investigate the agents' incentives to invest in different skills and relationships—when these investments must be sunk before bargaining takes place. This is the case, for example, in labor markets, in which both employers and employees must make substantial investments before they know who is going to end up matching with whom. In [Elliott and Talamàs \(2018\)](#), we study the extent to which the classical holdup problem (e.g., [Williamson 1975](#), [Grossman and Hart 1986](#), [Hart and Moore 1990](#), [Hosios 1990](#), [Acemoglu 1996, 1997](#), [Cole, Mailath, and Postlewaite 2001](#), [Elliott 2015](#)) is actually a problem in markets that—as in the present paper—attract traders over time. Leveraging the characterization of the present paper to study how agents' use their investments to obtain more favorable outcomes in decentralized markets is an interesting avenue of future research.

Third, following most of the related literature, this paper focuses on settings in which each coalition's productivity is independent of which other coalitions form. In many settings of interest, however, this is counterfactual (e.g., because different coalitions are in competition with each other, or because they provide products that complement each other). The analysis of this paper can be extended to characterize the structure of the coalitions that form in equilibrium in settings with externalities but, in general, these externalities can imply that such equilibria do not exist, or that there is more than one such equilibrium. [Ray and Vohra \(1999\)](#) and [Ray \(2007\)](#) extend the construction of the (no-delay stationary subgame-perfect)

equilibria in [Chatterjee, Dutta, Ray, and Sengupta \(1993\)](#) to settings with externalities across coalitions.⁵⁰ A similar generalization of the construction of perfect equilibria that I describe in this paper to settings with general externalities across coalitions seems attainable—even if substantially more involved because of the more general preferences considered in this paper.

Finally, an important direction for future research is to investigate the conditions under which the outcome of decentralized bargaining in the thick markets considered in this paper is efficient. If one considers the economy consisting of the set of all the agents that have matched in equilibrium before any given period, the resulting allocation is in the core of this economy, and hence efficient. As a result, the only source of inefficiency in these markets can be the frequency with which agents of different types match in equilibrium. The modeling approach that we develop in [Elliott and Talamàs \(2018\)](#) to study investment incentives is better suited to investigate this question than the one of the present paper (mainly because, in contrast to the present paper, it features exogenous entry). Using a framework along the lines of the one in [Elliott and Talamàs \(2018\)](#), we are currently investigating the potential sources of inefficiencies in dynamically-thick markets, thus providing the natural counterpart to the large literature investigating the efficiency of decentralized bargaining outcomes in thin markets (e.g., [Compte and Jehiel 2010](#), [Burguet and Caminal 2018](#) and [Elliott and Nava 2018](#)).

⁵⁰See also [de Clippel and Serrano \(2008\)](#) and [Maskin \(2016\)](#).

A Appendix: Details omitted from section 4

The objective of this section is to prove the main result of this paper: [Theorem 1](#). It follows from [Lemma A.2](#) below that the payoff profile χ defined by algorithm \mathcal{A} converges to the payoff profile χ^* defined by algorithm \mathcal{A}^* . As a result, [Theorem 1](#) follows from combining [Proposition A.1](#), [Proposition A.2](#) and [Proposition A.3](#) below.⁵¹ I start in [subsection A.1](#) by rewriting—in a more tractable way—the system of equations (2) that defines the equilibrium cutoff profile w in the bargaining game \mathcal{G} . In [subsection A.2](#), I define the family of auxiliary games $\{\mathcal{G}_x^C\}$, and I characterize the stationary subgame-perfect equilibrium of each of these games. In [subsection A.3](#), I use these equilibria to characterize the solution to the system of equations (2). In [subsection A.4](#), I show how the algorithm \mathcal{A} finds the unique solution to the system of equations (2). Finally, in [subsection A.5](#), I show that the algorithm \mathcal{A}^* finds the unique Nash-credible payoff profile.

A.1 Rewriting the system of equations (2)

I start by defining a set of functions $\{f_i\}_{i \in N}$ that allows me to rewrite the system (2) of equations that characterizes the equilibrium cutoffs in the game \mathcal{G} in a more tractable way. [Figure 14](#) illustrates this definition.

Definition A.1. For each type i , let the function $f_i : [y_i, \infty] \rightarrow \mathbb{R}$ be (implicitly) defined by

$$u_i(x) = \alpha_i u_i(y_i) + (1 - \alpha_i) u_i(f_i(x))$$

where (recall that) $\alpha_i := \frac{q}{(1-q)p_i + q}$.

In words, $f_i(x)$ is the amount that an agent of type i needs to be able to obtain when she is the proposer in order for her to be indifferent between accepting and rejecting the amount x (when she is not the proposer).⁵² The fact that the utility function u_i is strictly increasing ensures that f_i is well defined and, since $\alpha_i \in (0, 1)$, that $f_i(x) > x$. Moreover, it follows from the concavity of the utility function u_i that the difference between $f_i(x)$ and x is increasing in x (see [Figure 14](#)). [Lemma A.1](#) highlights this fact, which plays an important role in the analysis that follows.

⁵¹For compactness, I incorporate the proofs of [Proposition 4.2](#), [Proposition 4.3](#), and [Proposition 4.4](#) highlighted in [section 4](#) within the proofs of [Proposition A.1](#), [Proposition A.2](#) and [Proposition A.3](#) below.

⁵²The validity of this interpretation follows from the argument used in [section 4](#) to derive the system (2), which, for brevity, I do not reproduce here.




Figure 14: Illustration of the definition of the function f_i (Definition A.1).

Lemma A.1. $f_i(x) - x$ is increasing in x .

With the definition of the function f_i at hand, system (2) can be written as

$$(9) \quad f_i(w_i) = \max_{C \subseteq N} \left(y(C) - \sum_{j \in C-i} w_j \right) \text{ for all } i \text{ in } N.$$

A.2 Auxiliary game: Bargaining in a fixed coalition with exogenous outside options

In this subsection, I prove Lemma A.2 below, which characterizes the equilibrium of the auxiliary game \mathcal{G}_x^C in which (i) only one coalition $C \subseteq N$ produces positive surplus when it forms, and (ii) agents have exogenous outside options given by $x \in \mathbb{R}^N$. As I described in section 4, the equilibria of the set of auxiliary bargaining games $\{\mathcal{G}_x^C\}$ provide the basis for building the equilibrium of the bargaining game \mathcal{G} of interest.

Definition A.2 (Game \mathcal{G}_x^C). For any coalition $C \subseteq N$ and any payoff profile x in \mathbb{R}^N satisfying the feasibility constraint $\sum_{j \in C} x_j \leq y(C)$, the bargaining game \mathcal{G}_x^C is defined exactly as the bargaining game \mathcal{G} with the following two modifications. First, the surplus $y(D)$ of each coalition $D \neq C$ is reduced to 0. Second, at the end of each period, every active agent of type $j \in C$ can opt to leave the market and obtain x_j units of surplus.

The fact that, in the game \mathcal{G}_x^C , there is only one relevant coalition substantially simplifies the system that pins down the equilibrium cutoffs in this auxiliary game. Indeed, I now use a relatively direct argument to show that this cutoff profile exists and is unique (Lemma A.2). While this result is not surprising given the existing results in the literature, I am not aware of the existence of a proof of this result in the general framework of the present

paper (featuring coalitions of arbitrary size and imperfectly transferable utility), so I provide the proof of [Lemma A.2](#) below. The first part of this argument (existence and uniqueness of the equilibrium cutoff profile) is a generalization of standard arguments in the literature (see for example [Ray, 2007](#), Chapter 7) to the setting without perfectly transferable utility considered in this paper. The second part of this argument (convergence to the Nash bargaining solution) is a generalization of the analogous argument in [Binmore, Rubinstein, and Wolinsky \(1986\)](#).

Lemma A.2. *For any coalition $C \subseteq N$ and any payoff profile x in \mathbb{R}^N satisfying the feasibility constraint $\sum_{j \in C} x_j \leq y(C)$, the game \mathcal{G}_x^C has unique stationary subgame-perfect equilibrium. Moreover, as the bargaining friction q goes to zero, the stationary subgame-perfect equilibrium payoff profile of \mathcal{G}_x^C converges to the profile that solves*

$$(10) \quad \operatorname{argmax}_{s \in \mathbb{R}^C} \prod_{j \in C} [u_j(s_j) - u_j(y_j)]^{p_j} \quad \text{subject to} \quad y(C) \geq \sum_{j \in C} s_j \quad \text{and} \quad s_j \geq x_j \quad \text{for all } j \text{ in } C.$$

Proof. Consider a stationary subgame-perfect equilibrium of the game \mathcal{G}_x^C . For each type i in C , let v_i be the amount that an agent of type i is indifferent between accepting and rejecting in any given period. The argument analogous to the one used to derive system (2) in [section 4](#) implies that the profile v must satisfy

$$(11) \quad f_i(v_i) = \max \left[f_i(x_i), y(C) - \sum_{j \in C-i} v_j \right] \quad \text{for all } i \text{ in } C.$$

Let the profile v in \mathbb{R}^C be a solution to system (11). The fact that each function f_i is increasing implies that $v_j \geq x_j$ for each type j different than ℓ , and hence that

$$(12) \quad f_\ell(v_\ell) = y(C) - \sum_{j \in C-\ell} v_j \leq y(C) - \sum_{j \in C-\ell} x_j.$$

Intuitively, [Equation 12](#) says that the amount that an agent of type ℓ gets when she is the proposer in the auxiliary game \mathcal{G}_x^C is bounded above by the available net surplus in the coalition C after giving to each of its members his outside option. Moreover, the fact that, by definition, $f_i(v_i) \geq y(C) - \sum_{j \in C-i} v_j$ with equality for all types i with $v_i > x_i$, implies that

$$(13) \quad f_\ell(v_\ell) - v_\ell \leq f_i(v_i) - v_i \quad \text{with equality for all } i \text{ with } v_i > x_i.$$

In words, [Equation 13](#) says that the difference between the amount that an agent of type ℓ gets when she is the proposer and when she is the receiver in the auxiliary game \mathcal{G}_x^C is the same as the difference between the amount that an agent of any other type i gets when she

is the proposer and when she is the receiver, unless type i 's outside option binds, in which case the former is smaller than the latter.

Existence of a solution to system (11) follows from Brouwer's fixed point theorem. To prove uniqueness, suppose for contradiction that there are two profiles v, v' that solve system (11). Define S to be the set of all types for which these two solutions differ; that is, $S := \{i \in N \mid v_i \neq v'_i\}$. Let i be one of the elements of the set S for which $f_i(v_i) - v_i$ is highest, and suppose without loss of generality that $f_i(v_i) - v_i$ is an upper bound on $\{f_j(v'_j) - v'_j\}_{j \in S}$. By Lemma A.1, $f_i(v_i) - v_i$ is increasing in v_i , so we also have that $v_i > v'_i$. Moreover we have that

$$(14) \quad f_i(v_i) = y(C) - \sum_{j \in C-i} v_j,$$

since otherwise, $i \neq \ell$ and $v_i = z_i$ (see Equation 13), which contradicts the fact that $v'_i \geq z_i$. In particular, $f_i(v_i) - v_i = f_\ell(v_\ell) - v_\ell$, so Equation 13 implies

$$(15) \quad f_j(v_j) - v_j \geq f_i(v_i) - v_i \text{ for all } j \text{ in } C.$$

Given the choice of type i , Equation 15 implies that

$$f_j(v_j) - v_j \geq f_j(v'_j) - v'_j \text{ for all } j \text{ in } C,$$

or, using again that $f_j(v_j) - v_j$ is increasing in v_j , that $v_j \geq v'_j$. But then, Equation 14 combined with the fact that the function f_i is increasing and, by definition,

$$f_i(v'_i) \geq y(C) - \sum_{j \in C-i} v'_j$$

implies that $v'_i \geq v_i$, a contradiction.

It only remains to show that the payoff profile that solves system (11) converges, as q goes to zero, to the payoff profile that solves (10). Using Definition A.1, system (11) can be rewritten as

$$u_i(v_i) = \max \left[u_i(x_i), \alpha_i u_i(y_i) + (1 - \alpha_i) u_i \left(y(C) - \sum_{j \in C-i} v_j \right) \right] \text{ for all } i \text{ in } C.$$

Fix two types a and b in C be such that

$$\begin{aligned} u_a(v_a) &= \alpha_a u_a(y_a) + (1 - \alpha_a) u_a \left(y(C) - \sum_{j \in C-a} v_j \right) \\ u_b(v_b) &= \alpha_b u_b(y_b) + (1 - \alpha_b) u_b \left(y(C) - \sum_{j \in C-b} v_j \right). \end{aligned}$$




Figure 15: Illustration of the fact that the payoff profile v that solves system (11) is such that, for any two types a and b with $v_a > x_a$ and $v_b > x_b$, (v_a, v_b) converges, as q goes to zero, to the payoff profile that maximizes (16). The graph U corresponds to the set of utility pairs $\{(u_a(s_a), u_b(s_b)) \mid s \in S\}$. This figure is analogous to Figure 4.2 in Osborne and Rubinstein (1990).

It is enough to show that (v_a, v_b) converges to the maximizer of

$$(16) \quad [u_a(s_a) - u_a(y_a)]^{p_a} [u_b(s_b) - u_b(y_b)]^{p_b}$$

over the set

$$S = \left\{ (s_a, s_b) \in \mathbb{R}^2 \mid s_a + s_b = y(C) - \sum_{j \in C-a-b} v_j \right\}.$$

We have that $u_a(f_a(v_a)) - u_a(y_a) = \frac{1}{1-\alpha_a} [u_a(v_a) - u_a(y_a)]$. Hence, using the fact that the first order approximation around $\hat{q} = 0$ of $\left[\frac{1}{1-\alpha_a}\right]^{p_a}$ is e^q , we get that the difference between

$$[u_a(f_a(v_a)) - u_a(y_a)]^{p_a} [u_b(v_b) - u_b(y_b)]^{p_b}$$

and

$$[u_a(v_a) - u_a(y_a)]^{p_a} [u_b(f_b(v_b)) - u_b(y_b)]^{p_b}$$

is (to a first order) 0. Moreover, given that, by definition, for each type j and each number x , we have that $f_j(x)$ converges to x as q goes to zero, this implies that (v_a, v_b) converges to the maximizer of (16) over the set S . Figure 15 illustrates. \square

A.3 Credibility and equilibrium

I say that the payoff profile x in \mathbb{R}^N is *credible* if, for every type i , x_i is i 's x_{-i} -best share. I now turn to proving Proposition A.1, which implies that a profile w in \mathbb{R}^N is an equilibrium

cutoff profile in the bargaining game \mathcal{G} if and only if it is *credible*.

Proposition A.1. *A profile \mathbf{x} in \mathbb{R}^N solves system (9) if and only if it is credible.*

Proof. Let the profile \mathbf{x} in \mathbb{R}^N be such that, for every type i , x_i is i 's \mathbf{x} -best share, and let ℓ be in N . Recalling Equation 12, we have that

$$f_\ell(x_\ell) \leq \max_{C \subseteq N} \left(y(C) - \sum_{j \in C-\ell} x_j \right),$$

so we only need to show that there exists a coalition $C \subseteq N$ with $f_\ell(x_\ell) \geq y(C) - \sum_{j \in C-\ell} x_j$, but this is satisfied by ℓ 's \mathbf{x} -best coalition. Indeed, let C be ℓ 's \mathbf{x} -best coalition, and suppose for contradiction that $f_\ell(x_\ell) < y(C) - \sum_{j \in C-\ell} x_j$. This implies that the profile \mathbf{w} in \mathbb{R}^C that solves system (11) is such that $w_\ell = x_\ell$ (by the assumption that x_ℓ is ℓ 's \mathbf{x} -best share), and that $w_i > x_i$ for some type i in $C - \ell$ (otherwise, $y(C) - \sum_{j \in C-\ell} x_j$ is equal to $y(C) - \sum_{j \in C-\ell} w_j$, which by definition is itself equal to $f_\ell(x_\ell)$, a contradiction). Hence, the same profile \mathbf{w} solves system (11) after interchanging the roles of i and ℓ in this system, a contradiction of the assumption that x_i is i 's \mathbf{x} -best share.

In the other direction, suppose that the profile \mathbf{x} in \mathbb{R}^N solves system (9). Let ℓ in N and $C \subseteq N$ be such that $f_\ell(x_\ell) = y(C) - \sum_{j \in C-\ell} x_j$. First, note that ℓ 's \mathbf{x} -share z_ℓ in any coalition $D \neq C$ is bounded above by x_ℓ , since, using Equation 12,

$$f_\ell(z_\ell) \leq y(D) - \sum_{j \in D-\ell} x_j \leq y(C) - \sum_{j \in C-\ell} x_j = f_\ell(x_\ell).$$

Hence, it is enough to show that ℓ 's \mathbf{x} -share in C is bounded below by x_ℓ . Let the profile \mathbf{w} in \mathbb{R}^C solve system (11). Suppose for contradiction that $w_\ell < x_\ell$. Then, $f_\ell(w_\ell) < y(C) - \sum_{j \in C-\ell} x_j$, which implies that $w_j > x_j$ for some type j in C . Using Equation 13, the fact that $f_\ell(w_\ell) - w_\ell$ is increasing in w_ℓ , and \mathbf{x} solves system (9), we get

$$f_j(w_j) - w_j = f_\ell(w_\ell) - w_\ell < f_\ell(x_\ell) - x_\ell = y(C) - \sum_{j \in C} x_j \leq f_j(x_j) - x_j$$

which implies that $w_j < x_j$, a contradiction. □

A.4 Construction of the unique credible payoff profile

Proposition A.2. *The payoff profile χ defined by algorithm A is the unique credible profile.*

Proof. The fact that algorithm \mathcal{A} updates each type's outside option at most once, and such updates only increase types' outside options, implies that both the χ -best share and the χ -best coalitions of each type in S_k , for each $k \leq \kappa$, are exactly as her x^k -best share and her x^k -best coalitions. This implies, in turn, that χ is credible as long as $X_\kappa = N$. The rest of this proof consists of showing that X_κ is equal to N and that every credible profile gives χ_i to each type i .

First, I prove by induction in k that S_k is empty only if X^{k-1} is equal to N (so that X_κ is equal to N). Let k be such that X_{k-1} is a strict subset of N (this induction hypothesis is vacuously true when $k = 1$, so there is no need to prove the base step separately). Denoting, for each coalition C , agent i 's x^{k-1} -share in C by x_i^C , by definition, there exists a number μ_C such that $f_i(x_i^C) - x_i^C$ is equal to μ_C for every agent i in $C - X_{k-1}$. A coalition D with maximum μ_D (among those coalitions C such that $C - X_{k-1}$ is nonempty) is an x^{k-1} -best coalition of all types in $C - X_{k-1}$, since the fact that u_i is concave implies that $f_i(x_i) - x_i$ is increasing in x_i (see Figure 14).

Second, I prove by induction in k that, for each $k \leq \kappa$, every credible profile gives x_i^k to each type i in S_k (so that χ is the only possible credible profile). Let x be a credible payoff profile. Let k be such that x_i is equal to x_i^{k-1} for each agent i in X_{k-1} (again, this induction hypothesis is vacuously true when $k = 1$, so there is no need to prove the base step separately). Let C be a coalition that is an x^{k-1} -best coalition of all its members in $N - X_{k-1}$. Suppose for contradiction that, for some i in $C - X_{k-1}$, x_i is strictly smaller than i 's x^{k-1} -share in C (the induction hypothesis together with fact that algorithm \mathcal{A} only updates outside options upwards implies that x_i cannot be strictly bigger than i 's x^{k-1} -share in C). This implies that i 's x -share in C is strictly smaller than i 's x^{k-1} -share in C which implies, in turn, that for some j in $C - X_{k-1}$, x_j is strictly bigger than j 's x^{k-1} -share in C (that is, j 's x^{k-1} -best share), which, as just argued, contradicts the induction hypothesis. \square

A.5 Construction of the unique Nash credible payoff profile

Proposition A.3. *The payoff profile χ^* defined by algorithm \mathcal{A}^* is the unique Nash-credible profile.*

The only part of the proof of Proposition A.3 that differs from the proof of Proposition A.2 is the reasoning behind the fact that, for each step k with $X_k \neq N$, S_k is not empty. The argument in this case is analogous to that in Pycia (2012, pages 330-331).⁵³ Denoting, for

⁵³Pycia (2012) uses this argument to illustrate how there exists a stable coalitional structure when coalitional output is shared according to the Nash bargaining solution (with exogenous outside options).

each coalition C , type i 's \mathbf{x}^{k-1} -Nash share by x_i^C , and letting u_i' denote the derivative of the utility function u_i , we have that $u_i(x_i^C)/u_i'(x_i^C)$ is the same for every type i in $C - X^{k-1}$; denote by μ_C this common value. A coalition C with maximum μ_C is the \mathbf{x}^{k-1} -best coalition of all its members outside of X_{k-1} , since each type's \mathbf{x}^{k-1} -Nash share in C is increasing in μ_C .

References

- ABREU, D. AND M. MANEA (2012a): "Markov equilibria in a model of bargaining in networks," *Games and Economic Behavior*, 75, 1–16.
- (2012b): "Bargaining and efficiency in networks," *Journal of Economic Theory*, 147, 43–70.
- ABREU, D. AND D. PEARCE (2015): "A dynamic reinterpretation of Nash bargaining with endogenous threats," *Econometrica*, 83, 1641–1655.
- 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.
- AGRANOV, M., M. ELLIOTT, AND E. TALAMÀS (in preparation): "A thick market bargaining experiment," *Mimeo*.
- BAGWELL, K., R. W. STAIGER, AND A. YURUKOGLU (2018): "'Nash-in-Nash' Tariff Bargaining with and without MFN," *Mimeo*.
- BARON, D. P. AND J. A. FERREJOHN (1989): "Bargaining in legislatures," *American Political Science Review*, 83, 1181–1206.
- BECKER, G. S. (1973): "A theory of marriage: Part I," *Journal of Political Economy*, 81, 813–846.
- BINMORE, K., A. RUBINSTEIN, AND A. WOLINSKY (1986): "The Nash bargaining solution in economic modelling," *The RAND Journal of Economics*, 176–188.
- BINMORE, K., A. SHAKED, AND J. SUTTON (1989): "An outside option experiment," *The Quarterly Journal of Economics*, 104, 753–770.
- BINMORE, K. G. (1987): "Nash bargaining theory (II)," in *The economics of bargaining*, ed. by K. Binmore and P. Dasgupta, Cambridge: Basil Blackwell.
- BINMORE, K. G. AND M. J. HERRERO (1988): "Matching and bargaining in dynamic markets," *The Review of Economic Studies*, 55, 17–31.
- BRENZEL, H., H. GARTNER, AND C. SCHNABEL (2014): "Wage bargaining or wage posting? Evidence from the employers' side," *Labour Economics*, 29, 41–48.

- BURGUET, R. AND R. CAMINAL (2018): "Coalitional Bargaining with Consistent Counterfactuals," *Mimeo*.
- CARRARO, C. AND D. SINISCALCO (1993): "Strategies for the international protection of the environment," *Journal of public Economics*, 52, 309–328.
- CHATTERJEE, K., B. DUTTA, D. RAY, AND K. SENGUPTA (1993): "A noncooperative theory of coalitional bargaining," *The Review of Economic Studies*, 60, 463–477.
- CHODOROW-REICH, G., J. COGLIANESE, AND L. KARABARBOUNIS (2018): "The Macro Effects of Unemployment Benefit Extensions: a Measurement Error Approach," *The Quarterly Journal of Economics*.
- CHODOROW-REICH, G. AND L. KARABARBOUNIS (2016): "The cyclicity of the opportunity cost of employment," *Journal of Political Economy*, 124, 1563–1618.
- COLE, H., G. MAILATH, AND A. POSTLEWAITE (2001): "Efficient Non-Contractible Investments in Finite Economies," *The BE Journal of Theoretical Economics*, 1, 1–34.
- COLLARD-WEXLER, A., G. GOWRISANKARAN, AND R. LEE (forthcoming): "'Nash-in-Nash' Bargaining: A Microfoundation for Applied Work," *Journal of Political Economy*.
- COMPTE, O. AND P. JEHIEL (2010): "The coalitional Nash bargaining solution," *Econometrica*, 78, 1593–1623.
- COROMINAS-BOSCH, M. (2004): "Bargaining in a network of buyers and sellers," *Journal of Economic Theory*, 115, 35–77.
- DE CLIPPEL, G. AND R. SERRANO (2008): "Bargaining, coalitions and externalities: A comment on Maskin," *Mimeo*.
- DUTTA, B. AND D. RAY (1989): "A concept of egalitarianism under participation constraints," *Econometrica*, 615–635.
- ECKHOUT, J. (2017): "Sorting in the Labor Market," *Annual Review of Economics*.
- ELLIOTT, M. (2015): "Inefficiencies in networked markets," *American Economic Journal: Microeconomics*, 7, 43–82.
- ELLIOTT, M. AND F. NAVA (2018): "Decentralized bargaining: Efficiency and the core," *Theoretical Economics*.

- ELLIOTT, M. AND E. TALAMÀS (2018): “No holdup in dynamic markets,” *Mimeo*.
- ERASLAN, H. (2002): “Uniqueness of stationary equilibrium payoffs in the Baron–Ferejohn model,” *Journal of Economic Theory*, 103, 11–30.
- (2016): “Uniqueness of stationary equilibrium payoffs in the Baron–Ferejohn model with risk-averse players,” *International Journal of Economic Theory*, 12, 29–40.
- ERASLAN, H. AND A. MCLENNAN (2013): “Uniqueness of stationary equilibrium payoffs in coalitional bargaining,” *Journal of Economic Theory*, 148, 2195–2222.
- FARRELL, J. AND S. SCOTCHMER (1988): “Partnerships,” *The Quarterly Journal of Economics*, 103, 279–297.
- FRAJA, G. D. AND J. SÁKOVICS (2001): “Walras retrouvé: Decentralized trading mechanisms and the competitive price,” *Journal of Political Economy*, 109, 842–863.
- GALE, D. (1987): “Limit theorems for markets with sequential bargaining,” *Journal of Economic Theory*, 43, 20–54.
- GALE, D. AND H. SABOURIAN (2006): “Markov equilibria in dynamic matching and bargaining games,” *Games and Economic Behavior*, 54, 336–352.
- 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.
- HAGEDORN, M., F. KARAHAN, I. MANOVSKII, AND K. MITMAN (2013): “Unemployment benefits and unemployment in the great recession: the role of macro effects,” Tech. rep., National Bureau of Economic Research.
- HALL, R. E. (2017): “High discounts and high unemployment,” *American Economic Review*, 107, 305–30.
- HALL, R. E. AND A. B. KRUEGER (2012): “Evidence on the incidence of wage posting, wage bargaining, and on-the-job search,” *American Economic Journal: Macroeconomics*, 4, 56–67.
- HALL, R. E. AND P. R. MILGROM (2008): “The limited influence of unemployment on the wage bargain,” *American Economic Review*, 98, 1653–74.

- HART, O. AND J. MOORE (1990): "Property rights and the nature of the firm," *Journal of Political Economy*, 98, 1119–1158.
- HART, S. AND A. MAS-COLELL (1989): "Potential, value, and consistency," *Econometrica*, 589–614.
- HO, K. AND R. S. LEE (forthcoming): "Equilibrium provider networks: bargaining and exclusion in health care markets," *American Economic Review*.
- HORN, H. AND A. WOLINSKY (1988): "Bilateral monopolies and incentives for merger," *The RAND Journal of Economics*, 408–419.
- HOSIOS, A. J. (1990): "On the efficiency of matching and related models of search and unemployment," *The Review of Economic Studies*, 57, 279–298.
- JÄGER, S., B. SCHOEFER, S. YOUNG, AND J. ZWEIMÜLLER (2018): "Wages and the value of nonemployment," *Mimeo*.
- KLEINBERG, J. AND É. TARDOS (2008): "Balanced outcomes in social exchange networks," in *Proceedings of the fortieth annual ACM symposium on Theory of computing*, ACM, 295–304.
- KRANTON, R. E. AND D. F. MINEHART (2001): "A theory of buyer-seller networks," *American Economic Review*, 91, 485–508.
- KRISHNA, V. AND R. SERRANO (1996): "Multilateral bargaining," *The Review of Economic Studies*, 63, 61–80.
- KRUSELL, P., T. MUKOYAMA, AND A. ŞAHİN (2010): "Labour-market matching with precautionary savings and aggregate fluctuations," *The Review of Economic Studies*, 77, 1477–1507.
- LAUERMANN, S. (2013): "Dynamic matching and bargaining games: A general approach," *American Economic Review*, 103, 663–89.
- LJUNGQVIST, L. AND T. J. SARGENT (2017): "The fundamental surplus," *American Economic Review*, 107, 2630–65.
- LUNDBERG, S. AND R. A. POLLAK (1996): "Bargaining and distribution in marriage," *Journal of Economic Perspectives*, 10, 139–158.
- MANEA, M. (2011): "Bargaining in stationary networks," *American Economic Review*, 101, 2042–2080.

- (2018): “Bottleneck links, essential intermediaries, and competing paths of diffusion in networks,” *Mimeo*.
- MASKIN, E. (2016): “How can cooperative game theory be made more relevant to economics? An open problem,” in *Open Problems in Mathematics*, Springer, 347–350.
- MCLENNAN, A. (1982): “A general noncooperative theory of bargaining,” *Mimeo*.
- MOLDOVANU, B. AND E. WINTER (1995): “Order independent equilibria,” *Games and Economic Behavior*, 9, 21–34.
- MORTENSEN, D. T. AND C. A. PISSARIDES (1994): “Job creation and job destruction in the theory of unemployment,” *The Review of Economic Studies*, 61, 397–415.
- MOULIN, H. (1982): *Bargaining and non-cooperative implementation*, Laboratoire d'économétrie de l'École polytechnique.
- NASH, J. (1953): “Two-person cooperative games,” *Econometrica*, 128–140.
- NASH, J. F. (1950): “The bargaining problem,” *Econometrica*, 155–162.
- NGUYEN, T. (2015): “Coalitional bargaining in networks,” *Operations Research*, 63, 501–511.
- OKADA, A. (2011): “Coalitional bargaining games with random proposers: Theory and application,” *Games and Economic Behavior*, 73, 227–235.
- OSBORNE, M. J. AND A. RUBINSTEIN (1990): *Bargaining and markets*, Academic Press, San Diego.
- PELEG, B. (1986): “On the reduced game property and its converse,” *International Journal of Game Theory*, 15, 187–200.
- PISSARIDES, C. A. (2000): *Equilibrium unemployment theory*, MIT press.
- POLANSKI, A. AND F. VEGA-REDONDO (2018): “Coalitional bargaining in repeated games,” *International Economic Review*.
- PYCIA, M. (2012): “Stability and preference alignment in matching and coalition formation,” *Econometrica*, 80, 323–362.
- RAY, D. (2007): *A game-theoretic perspective on coalition formation*, Oxford University Press.

- RAY, D. AND R. VOHRA (1999): "A theory of endogenous coalition structures," *Games and Economic Behavior*, 26, 286–336.
- ROCHFORD, S. C. (1984): "Symmetrically pairwise-bargained allocations in an assignment market," *Journal of Economic Theory*, 34, 262–281.
- 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.
- SERRANO, R. AND K.-I. SHIMOMURA (1998): "Beyond Nash bargaining theory: the Nash set," *Journal of Economic Theory*, 83, 286–307.
- SHAPLEY, L. S. AND M. SHUBIK (1971): "The assignment game I: The core," *International Journal of game theory*, 1, 111–130.
- SHIMER, R. (2005): "The cyclical behavior of equilibrium unemployment and vacancies," *American Economic Review*, 95, 25–49.
- SOBOLEV, A. I. (1975): "The characterization of optimality principles in cooperative games by functional equations," *Mathematical methods in the social sciences*, 6, 94–151.
- SORKIN, I. (2015): "Are there long-run effects of the minimum wage?" *Review of Economic Dynamics*, 18, 306–333.
- SUTTON, J. (1986): "Non-cooperative bargaining theory: An introduction," *The Review of Economic Studies*, 53, 709–724.
- TALAMÀS, E. (2018): "Price dispersion in stationary networked markets," *Mimeo*.
- WILLIAMSON, O. E. (1975): *The Economic Institutions of Capitalism*, New York Free Press.