

Nash bargaining with endogenous outside options*

Eduard Talamàs[†]

Abstract

This paper provides a tractable non-cooperative framework that uniquely pins down which coalitions form and how the resulting gains from trade are shared in complex stationary markets. As bargaining frictions vanish, the equilibrium payoff profile converges to a generalization of the Nash bargaining solution that accommodates endogenous outside options. This generalization prevents outside options from being determined in a circular way and, remarkably, always yields a unique prediction. The resulting theory uncovers an endogenous vertical market structure such that various economic shocks propagate—via outside options—from the top down, but not vice versa.

1 Introduction

In many markets, 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.). Understanding the resulting outcomes is important in order to assess the distributional effects of different types of technological change, for example, but predictions that are sharp, micro-founded and tractable are hard to obtain in general settings. In this paper, I describe a non-cooperative theory of coalition formation that maintains the transparency and amenability to applied exercises provided by the Nash bargaining solution, while providing sharp and easily-computable predictions about how various economic shocks propagate via outside options in rich stationary markets.

*Date printed: August 16, 2019.

[†]IESE Business School. 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, and to Pau Milán, Debraj Ray and numerous conference and seminar participants for useful feedback. This work has been supported by the Warren Center for Network & Data Sciences, and the Rockefeller Foundation (#2017PRE301). All errors are my own.

The highlights of this theory are fivefold. The first is *uniqueness*: The theory provides an essentially-unique map from primitives to outcomes. The second is *tractability*: An intuitive algorithm identifies which coalitions form and how the resulting surplus is shared. The third is *familiarity*: In the limit as the bargaining frictions vanish, the sharing rule in each coalition corresponds to the classical Nash bargaining solution—with the relevant outside options endogenously determined by the Nash bargaining solution in other coalitions. The fourth is a *characterization* of the equilibrium outcome that is independent of the bargaining protocol: The equilibrium payoff profile is the only one that satisfies a natural credibility property that prevents outside options from being determined in a circular way. The fifth is the uncovering of an *endogenous vertical market structure* that provides sharp comparative statics: The coalitions that form in equilibrium can be organized into tiers, in such a way that (small) changes in market fundamentals propagate—via outside options—from higher to lower tiers, but not vice versa.

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.¹ 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 *Nash 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 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 Nash 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

¹In many applications, the analyst faces various 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, Mukoyama, and Şahin \(2010\)](#), [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 \(2017\)](#).




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 $w' < 1/2$.

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.²

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 in the context of rich stationary markets.

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 markets with dynamic entry where the relevant matching opportunities are roughly constant over time. This follows the tradition of the literature on non-cooperative bargaining in stationary markets (e.g., [Rubinstein and Wolinsky 1985](#); [Gale 1987](#); [de Fraja and Sákovics 2001](#); [Manea 2011](#); [Nguyen 2015](#); [Polanski and Vega-Redondo 2018](#)).³ In contrast to previous work in this literature, I allow agents to strategically choose which coalitions to propose as well as how to share the resulting surplus.⁴ In addition to introducing an important bargaining dimension in these settings, this provides a theory of coalition formation that provides non-cooperative foundations for a natural gen-

²[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.

³See [Osborne and Rubinstein \(1990\)](#) for an overview of the origins of this literature, and [Manea \(2016\)](#) for an overview of the more recent advances.

⁴The exception is [Talamàs \(2019\)](#), where I use a version of the framework in the present paper tailored to networked buyer-seller markets. I discuss the connection with this literature in [subsubsection 4.7.3](#).




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.

eralization of the Nash bargaining solution with easy-to-compute unique predictions and sharp comparative statics.

I show that there exists an essentially-unique stationary subgame-perfect equilibrium, and I characterize which coalitions form and how the resulting gains from trade are shared in this equilibrium. The equilibrium payoff profile is uniquely characterized by a natural property that is independent of the bargaining protocol. For ease of exposition, here I state this result in the limiting case as the bargaining frictions vanish: The equilibrium payoff profile converges to the *unique* profile x that satisfies the following credibility property:

Each type i 's payoff x_i is her maximum Nash bargaining share—across all coalitions—subject to the constraint that *every other type j receives at least x_j* .

Roughly speaking, the strategic forces in the non-cooperative model require that each agent gets the maximum that she can justify as the result of Nash bargaining in some coalition *without appealing to her own outside option there*.⁵ Intuitively, this 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, this prevents MBA graduates from claiming an outside option 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 (see Figure 2).

The fact that the above credibility property pins down outcomes uniquely is remarkable. Note that, in general, many different payoff profiles x satisfy the property that each type i 's payoff x_i is her maximum share—across all coalitions—subject to the constraint that every

⁵The Nash threat points in this theory correspond to agents' utilities in autarky, and outside options enter as lower bounds on payoffs, as suggested by the outside option principle.

other type j receives her payoff x_j (see [section 2](#) for further discussion about this point). Surprisingly, however, requiring the above credibility property always pins down one among the many payoff profiles that satisfy this core-like constraint. The particular properties of the endogenous sharing rule (which converges to the Nash bargaining solution as bargaining frictions vanish) are important for this uniqueness result: The analogous credibility condition using other well-known sharing rules—such as the one proposed by [Kalai and Smorodinsky \(1975\)](#), for example—does not necessarily pin down the payoffs uniquely.⁶

I describe an intuitive algorithm that identifies which coalitions form and how they share the resulting surplus in equilibrium. This algorithm transparently characterizes how the primitives shape the outcomes, and provides subtle but also intuitive comparative statics. In particular, it shows that 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 coalitions that are in higher tiers. This implies that (small) changes in market fundamentals propagate—via outside options—from higher to lower tiers, but not vice versa.

The resulting theory of coalition formation is broadly consistent with the view that bargaining plays a more prominent role in the determination of high-skill than low-skill wages (e.g., [Hall and Krueger 2012](#) and [Brenzel et al. 2014](#)) and that the wages in some industries can be entirely determined by the productivity of other industries (e.g., [Baumol et al. 1966](#)). For example, when agents are vertically differentiated by skill, the coalitions of high-skilled agents are in the first tier, where no outside options bind. Hence, an increase in the productivity of one of these coalitions is shared among all its members as specified by the Nash bargaining solution, and it increases the share of surplus that these agents obtain in other coalitions. If, in addition, we restrict attention to two-sided pairwise matching markets, an increase in the skill of an agent can affect the payoffs of agents whose skill is lower than hers, but it does not affect the payoff of anyone with higher skill. This is because the relevant outside options of one’s counterparties are always determined by bargaining with more skilled agents.

This paper also suggests a mechanism by which recent increases in labor market sort-

⁶As I discuss in [subsection 4.3](#), the above credibility property in terms of the sharing rule that emerges in equilibrium pins down outcomes uniquely because this sharing rule generates *pairwise-aligned preferences over coalitions* ([Pycia 2012](#)): For any two coalitions C and D containing agents a and b , if a ’s share in C is bigger than a ’s share in D , then the same is true for b . See [Pycia \(2012\)](#) for an example that shows that the Kalai-Smorodinsky solution does not satisfy this property.

ing (e.g., Card, Heining, and Kline 2013; Eeckhout 2018; Song, Price, Guvenen, Bloom, and von Wachter 2019) can divide labor markets into effectively-independent submarkets—potentially leading to a sharp disconnection between the determinants of high-skill and low-skill wages, as well as to widening inequality.⁷ Indeed, as in the canonical marriage market model of Becker (1973), positive assortative matching arises if and only if skills are complementary. But, in contrast to Becker’s framework (and much of the subsequent literature), the present theory pins down prices uniquely, and hence provides testable predictions about how positive assortative matching affects the way in which economic shocks propagate via outside options. For example, in two-sided pairwise matching markets where workers and firms 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. In particular, in this case the market endogenously decomposes into different submarkets that can be ordered by skill, such that (small) economic shocks propagate between them but not across them.

Roadmap

The rest of this paper is organized as follows. I start by illustrating the framework and main result with a simple example in [section 2](#). I describe the general model in [section 3](#). I characterize the essentially-unique stationary subgame-perfect equilibrium and I further discuss the connections with existing work in [section 4](#). Finally, I illustrate the comparative statics of the theory in vertically differentiated markets in [section 5](#), and I conclude in [section 6](#). I defer the formal proofs of most of the results to [Appendix A](#).

2 Example

In this section, I illustrate the setting and the main results of this paper using an example. The objective of this example is not to display the full generality of the framework, but to illustrate the main ideas 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.

⁷See [Acemoglu and Autor \(2011\)](#) for a detailed account of the dramatic rise in U.S. earnings inequality since the 1970s.




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).

Consider a large city where different agents (in the culinary industry, say) go to in search of business opportunities.⁸ 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 Agents of all types arrive to the city over time to find potential partners with whom to start a venture. For simplicity, assume that each agent can only be part of one 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 match is independent of which other matches form (because each venture is implemented in a different geographic market, say), and that the surpluses of the four possible coalitions are as illustrated in Figure 3.¹⁰

Which ventures form, and how is the resulting surplus 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 deliver food to their clients' doors, say) affect this market?

I investigate these questions by studying the equilibrium behavior of these agents when they bargain according to the following infinite-horizon protocol: In each period, one agent of each type is in the city looking to form a venture, and one of them is selected uniformly at random to be the proposer. The selected agent can propose a match as well as how to share

⁸I 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.

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

¹⁰Food being the most important part of a culinary experience, I assume that a match between a chef and a manager generates more value than a match between a cook and a maître.




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).

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 with some probability q , preventing her from starting any venture.¹¹ 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 to form 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 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 coalitions 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 them), in the limit as the bargaining friction q vanishes, chefs and maîtres share their gains from trade equally.¹²

¹¹For 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.

¹²More generally, their terms of trade are as prescribed by the Nash bargaining solution, with the Nash 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.




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.

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 $50 - 30$.¹³ Similarly, the cooks propose to match with the managers, and—in the limit as the bargaining friction q vanishes—they share their surplus $20 - 30$ (Figure 5 illustrates).

As the bargaining friction q vanishes, the equilibrium payoff profile converges to the unique profile x that satisfies the following property: Each type i 's payoff x_i is the maximum that she can justify as the result of equal sharing in some coalition subject to the constraint that *her counterparty j receives at least x_j* . In other words, the strategic forces in the non-cooperative model require that each type is able to justify her payoff as resulting from equal sharing in some coalition while respecting *the others'* payoffs. Remarkably, this requirement alone pins down the equilibrium payoffs uniquely. Figure 6 illustrates four of the many payoff profiles that do not satisfy this requirement despite being such that everyone gets the maximum that she can given others' payoffs.

In order to trace out how different economic shocks propagate via outside options, 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 equal sharing subject to the binding outside options determined in higher tiers. For example, Figure 7

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 to a greater degree 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: Four payoff profiles that are such that everyone gets the maximum that she can given others' payoffs. None of these profiles satisfy the property that each type i 's payoff x_i is the maximum that she can justify as the result of equal sharing in some coalition subject to the constraint that her counterpart j receives at least x_j .

illustrates that the chef-maître coalition is in the first tier. Equal sharing in this coalition 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.

3 Model: The bargaining game \mathcal{G}

I describe the primitives of the model in [subsection 3.1](#), the bargaining protocol that turns these primitives into a well-defined non-cooperative game in [subsection 3.2](#), and the notion of equilibrium that I focus on throughout the paper in [subsection 3.3](#).



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.

3.1 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 agent of each type in $C \subseteq N$ produces $y(C) \geq 0$ units of surplus when it forms.^{14,15} I refer to $y_i := y(i)$ as *type i ’s autarky surplus*. This 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

¹⁴As 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.

¹⁵For 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).

types. In other words, 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. For instance, in the example described in [section 2](#), we can model a situation in which some types are scarcer than others as follows: Suppose that, at each point in time, there are $n_i \geq 1$ active agents of each type i (instead of there being one of each), where $1/n_i$ is a measure of i 's scarcity. In each period, one type is selected uniformly at random, and then one of the corresponding active agents is selected uniformly at random to be the proposer. This leads to a direct relation between the scarcity of a type and the proposer probability of each agent of this type. This relation is the only channel through which type scarcities shape outcomes in this setting and, as I describe next, this framework allows to capture this heterogeneity, by letting the arbitrary profile p determine everyone's proposer probabilities.

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.¹⁶ If all of them accept, then they match with the proposer and they, together with the proposer, leave the market with the agreed shares. 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.

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. While the results of this paper carry over to more realistic models featuring exogenous entry (along the lines of the framework developed in [Elliott and Talamàs 2019](#)), 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.

¹⁶The order in which the agents respond is not relevant for the results.

Also, the results of this paper go through if one assumes that no agent is ever exogenously forced out of the market but that—instead—agents are impatient. The Nash bargaining solution then has to be appropriately constructed from agents' time preferences (see for example [Osborne and Rubinstein 1990](#)).

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 someone 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 $(h_t, j \rightarrow C, x, i)$.¹⁷ 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 Equilibrium characterization

In this section, I characterize the essentially-unique equilibrium of the bargaining game \mathcal{G} . I start, in [subsection 4.1](#), by deriving a system of equations that determines the equilibrium payoffs. This system formalizes the fundamental outside option Gordian knot: Each type's payoff depends on the surplus that she can obtain 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 this section is to show that this system admits a unique solution, and to

¹⁷I 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.

characterize this solution. Informally, the characterization strategy leverages the observation 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. For instance, in the example of [section 2](#), the chef-maître coalition is sufficiently productive so that neither chefs’ nor maîtres’ outside options are relevant when bargaining to form this coalition. As a result, they essentially share the surplus of this coalition equally—and this determines their binding outside options when bargaining to form 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 in any given coalition, and I show how it can be used to derive an upper bound on each type’s equilibrium payoff. 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 its members in every equilibrium of the game \mathcal{G} . Third, in [subsection 4.4](#), I leverage these observations to recursively construct everyone’s unique equilibrium payoffs.

In [subsection 4.5](#), I show that strategic bargaining in the game \mathcal{G} cuts the outside option Gordian knot in a very structured way: The equilibrium payoff profile is *the only one* that satisfies the following credibility property: Each type’s payoff is the maximum that she can justify as the result of isolated bargaining in some coalition subject to the constraint that everyone else in this coalition receives at least her equilibrium payoff. In [subsection 4.6](#), I show that the set of coalitions that form in equilibrium can be organized into tiers, in such a way that the relevant outside options in each coalition are determined exclusively in higher tiers. Finally, in [subsection 4.7](#) I discuss the connection of these results with existing literature.

4.1 Equilibrium cutoff profile

[Proposition 4.1](#) describes the equilibrium of the game \mathcal{G} . This equilibrium is essentially unique: Each type’s payoff is the same in every equilibrium; each type’s on-path responses are the same in every equilibrium; and 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, when selected to be the proposer, each agent of type i with $w_i > y_i$ 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.

Figure 4 in section 2 illustrates Proposition 4.1 in 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. 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.

Remark 4.1. In pairwise matching settings (in which each feasible match contains at most two agents), every subgame-perfect equilibrium is in stationary strategies; see Talamàs (2019, Proposition A.1).¹⁸ In this case, Proposition 4.1 holds without restricting attention to stationary strategies. When coalitions can contain more than two agents, however, the restriction to stationary strategies is important to pin down payoffs uniquely (see for example Osborne and Rubinstein (1990, section 3.13) for a discussion of this point in the special case in which only one coalition can form).

The rest of this section informally describes the argument behind the proof of Proposition 4.1 (deferring the formal details to Appendix A) and characterizes the equilibrium cutoff profile as well as the set of coalitions that form in equilibrium. The immediate objective is to describe a system of equations 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 $\max_{C \subseteq N} [y(C) - \sum_{j \in C-i} w_j]$. Type i makes acceptable offers in equilibrium if this quantity is strictly bigger than her autarky utility $u_i(y_i)$.¹⁹

By definition, each type i is indifferent between obtaining w_i right away, which gives her

¹⁸Proposition A.1 in Talamàs (2019) 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.

¹⁹To 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.

utility $u_i(w_i)$, and waiting for the next period, which gives her an expected utility of

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

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.

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

In the absence of bargaining frictions (i.e., when $q = 0$) system (1) has a great multiplicity of solutions. For instance, in the example described in [section 2](#), one extreme solution to this system is the profile that gives chefs and cooks 100 and 60, respectively, and 0 to both maîtres and managers. Another extreme solution 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 (1) 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 this system when $q = 0$. In other words, the unique equilibrium of the non-cooperative game picks one of the many plausible bargaining outcomes in the frictionless case.²⁰

Remark 4.2. A version of system (1), with linear utilities, characterizes the no-delay stationary equilibrium in the non-stationary bargaining framework of [Chatterjee et al. \(1993\)](#) and subsequent work (e.g., [Ray and Vohra 1999](#); [Ray 2007](#)). I discuss the connection with this work in [subsubsection 4.7.4](#).

²⁰This is, of course, the comparative advantage of the non-cooperative approach ([Rubinstein 1982](#)).

4.2 A personalized upper bound on the equilibrium cutoffs

The immediate objective is to provide, for each type, an upper bound on her equilibrium cutoff. The basic idea is this: 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. As a result, 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.²¹

To formalize this idea, consider the family $\{\mathcal{G}^C\}_{C \subseteq N}$ of variants of the bargaining game \mathcal{G} in which only one coalition $C \subseteq N$ can form.²² 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, as 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 Nash threat points given by autarky—that is, to the unique profile that solves

$$(2) \quad \operatorname{argmax}_{s \in \mathbb{R}^C} \prod_{j \in C} [u_j(s_j) - u_j(y_j)]^{p_j} \quad \text{subject to the feasibility constraint } y(C) \geq \sum_{j \in C} s_j.$$

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 is part of) is just below half of this coalition's surplus. This reflects the fact that the proposer obtains slightly more than half of the available gains from trade. But, in the limit as 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.

[Proposition 4.2](#) below formalizes the intuition that the equilibrium cutoff of each type in the bargaining game \mathcal{G} cannot be bigger than her *best share*—defined as follows. Informally,

²¹Preventing agents from accepting offers from other coalitions is not necessary for this exercise 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.

²²More 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.

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. [Definition 4.1](#) formalizes these concepts.

Definition 4.1. A type's *best share* is her maximum equilibrium cutoff—across all coalitions $C \subseteq N$ —in the auxiliary game \mathcal{G}^C . 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. A coalition is *perfect* if it is a best coalition of all its members.

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. Let w be a solution to (1). For each type i , w_i is bounded above by i 's best share.

[Figure 8](#) illustrates the exercise 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

The combination of [Proposition 4.3](#) and [Proposition 4.4](#) below shows that the personalized upper bound on each type's equilibrium cutoff (provided by [Proposition 4.2](#) above) is tight for at least one type.

Proposition 4.3. *Let w be a solution to (1). If type i is in a perfect coalition, then w_i is i 's best share.*

Intuitively, the fact that the cutoffs in \mathcal{G}^C are an upper bound on all of its members' equilibrium cutoffs in \mathcal{G} ([Proposition 4.2](#)) implies that no one in a perfect coalition C has a better option—when bargaining in \mathcal{G} —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 . For instance, given that maîtres' cutoff in the example of [section 2](#) is bounded above by 48, in the equilibrium of the game of interest, 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, maîtres' cutoff is 48 (its upper bound in the game of interest). Hence, the upper bound of 48 on chefs' cutoff is tight.

[Proposition 4.3](#) above is especially useful because, 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 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 bigger 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 bigger than j 's cutoff in the auxiliary game \mathcal{G}^D).²³ 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 bigger 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

²³[Pycia \(2012\)](#) labels this property *pairwise-aligned preferences over coalitions*, and shows that the Nash bargaining solution generates pairwise-aligned preferences over coalitions.

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).²⁴

4.4 Recursive characterization of the equilibrium cutoff profile

The immediate objective is to derive an upper bound on the equilibrium cutoff of each type when some of the other types’ cutoffs have already been determined. This is an essential part of the algorithm that pins down everyone’s equilibrium cutoffs in the game \mathcal{G} (Definition 4.3 below). As in subsection 4.2, the basic idea is that no one is hurt when the others’ outside options deteriorate. The difference with the argument above is that, now—in the hypothetical situation in which agents bargain to form one coalition in isolation—the outside options of the agents whose cutoff has already been defined do not deteriorate below their cutoff. To formalize this idea, for every coalition $C \subseteq N$ and every profile x in \mathbb{R}^N with $y(C) \geq \sum_{j \in C} x_j$, 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 leave the market with payoff x_i immediately after rejecting 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 where both the feasible coalition C and the outside option profile x are exogenous, and its members only bargain over how to share the resulting gains from trade subject to their exogenously given outside options (e.g., Binmore, Rubinstein, and Wolinsky 1986). This game has a unique stationary subgame-perfect equilibrium. In particular, the equilibrium cutoff profile of the game \mathcal{G}_x^C converges to the (generalized) Nash bargaining solution in coalition C , with the Nash threat points given by autarky and the outside option profile x imposing only lower bounds on the payoffs—that is, to the profile that solves

$$(3) \quad \operatorname{argmax}_{s \in \mathbb{R}^C} \prod_{j \in C} [u_j(s_j) - u_j(y_j)]^{p_j} \quad \text{subject to} \quad y(C) \geq \sum_{j \in C} s_j \quad \text{and} \quad s_j \geq x_j \quad \text{for all } j \text{ in } C.$$

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 $x_{\text{chefs}}^1 = x_{\text{maitres}}^1 = 48$ and $x_{\text{cooks}}^1 = x_{\text{managers}}^1 = 0$. Naturally, 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

²⁴This is essentially the argument that Pycia (2012) gives to show that the Nash bargaining solution generates pairwise aligned preferences over coalitions.




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 satisfies $x_{\text{chefs}} = x_{\text{maîtres}} = 48$ and $x_{\text{cooks}} = x_{\text{managers}} = 0$. 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.

involving either the chefs or the maîtres (the two types whose outside options have been updated to be their equilibrium cutoffs).

Definition 4.3 describes an algorithm that uses the concepts described in **Definition 4.2** below to pin down the equilibrium cutoff profile w .

Definition 4.2. For each coalition $C \subseteq N$ and each profile x in \mathbb{R}^N , type i 's x -share in C is her cutoff in the auxiliary game \mathcal{G}_x^C if well defined, and 0 otherwise. Type i 's x -best share is her maximum x -share across all coalitions $C \subseteq N$. A coalition C is i 's x -best coalition if i 's x -share in C is her x -best share. A coalition is x -perfect if it is an x -best coalition of all of its members.

Figure 9 illustrates that the chef-manager coalition is the managers' x^1 -best coalition (and is hence x^1 -perfect), and that the managers' x^1 -best share is 29. The argument analogous to the one behind **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 4.3** below formalizes this argument, in the form of an algorithm that recursively pins down everyone's cutoffs.

Definition 4.3 (Algorithm \mathcal{A}). Let $x^0 = \mathbf{0}$. In each step $s = 1, 2, \dots$, for each type i in an x^{s-1} -perfect coalition, let x_i^s be i 's x^{s-1} -best share; for each other type j , let $x_j^s = 0$. Stop in the first step S such that every type is in an x^S -perfect coalition, and let $\chi := x^S$.

Leveraging **Proposition 4.4** and **Proposition 4.3** inductively shows that algorithm \mathcal{A} ends in finitely many steps, and provides the following result.

Proposition 4.5. *The profile χ defined by algorithm \mathcal{A} is the unique equilibrium cutoff profile w .*

Proposition 4.5 makes the theory of coalition formation described in this paper especially tractable. Indeed, it paves the way for the characterization of the equilibrium payoff profile as the only one satisfying a certain credibility property. Also, it allows general comparative statics, and it uncovers the tier structure of the set of equilibrium coalitions that illustrates how shocks propagate via outside options.

4.5 Credibility and equilibrium

Understanding which coalitions form and how the resulting surplus is shared requires understanding how the outside option Gordian knot is cut. I now leverage Proposition 4.5 to show that the equilibrium of the bargaining game \mathcal{G} actually cuts this Gordian in an intuitive way: The equilibrium payoff of any given type in any given coalition is the maximum that she can justify as resulting from isolated bargaining in another coalition subject to the constraint that everyone *else* receives *at least* her equilibrium payoff. Moreover, the equilibrium cutoff profile w is the only profile that satisfies this property.

Notation 4.1. For each profile x in \mathbb{R}^N , x_{-i} denotes x after setting its i^{th} entry to 0.

Definition 4.4. The profile x is *credible* if, for every type i , x_i is i 's x_{-i} -best share.

In words, the profile x is *credible* if, for every type i , x_i is the maximum amount that she can justify as being her equilibrium payoff in the game $\mathcal{G}_{x_{-i}}^C$ for some coalition $C \subseteq N$; that is, the maximum amount that she can justify as being her equilibrium payoff in the game \mathcal{G}^C for some coalition $C \subseteq N$ subject to the constraint that everyone *else* in C receives *at least* what x gives her. At first sight, it might not be clear whether a credible profile always exists, or whether more than one such profile can exist. An inductive argument, however, shows that the profile χ defined by algorithm \mathcal{A} is the unique credible profile (Proposition A.2), giving the following result.

Theorem 1. *The equilibrium cutoff profile w is the unique credible profile.*

Theorem 1 provides intuition for how the strategic forces in the bargaining game \mathcal{G} pin down outcomes uniquely: These forces require each agent to justify her payoff by bargaining in some coalition without appealing to her own outside option, and this rules out multiplicity of consistent payoff profiles driven by outside options losing connection with fundamentals—as it occurs in the example illustrated in Figure 2, for instance. Remarkably,

there is a unique profile that satisfies this property—and this explains why the game \mathcal{G} pins down outcomes uniquely.

The fact that the equilibrium cutoff profile of the game \mathcal{G}_x^C converges to the Nash bargaining solution (3) suggests that the limit equilibrium cutoff profile w of the bargaining game \mathcal{G} satisfies a similar credibility property in terms of the classical Nash bargaining solution. I now turn to formalizing this idea.

Definition 4.5. Fix a coalition $C \subseteq N$ and a profile x in \mathbb{R}^N . The profile of x -Nash shares in C is the solution to (3) if well defined, and 0 otherwise. Type i 's x -Nash best share is her maximum x -Nash share across all coalitions. Coalition C is one of type i 's x -Nash best coalitions if her x -Nash share in C is equal to her maximum x -Nash share across all coalitions. Coalition C is x -Nash perfect if it is an x -Nash best coalition of all of its members.

In words, the profile x is *Nash credible* if, for each type i , x_i is the maximum that i can justify as the result of Nash bargaining in some coalition—with Nash threat points given by autarky—subject to the constraint that everyone else receives at least the payoff that x gives her. An algorithm analogous to \mathcal{A} that uses agents' Nash-best shares as ingredients instead of their best shares (Definition A.2) computes the unique Nash credible payoff profile, providing the following result.

Corollary 4.3. *As the bargaining friction q vanishes, the equilibrium cutoff profile w converges to the unique Nash credible profile.*

For instance, Figure 5 illustrates how each type's limit payoff in the example of section 2 is the maximum (across her two potential coalitions) that she can justify using the Nash bargaining solution (which, in this example, boils down to equal sharing) subject to the constraint that everyone *else* receives *at least* her limit payoff. Corollary 4.3 implies that this is the unique profile that satisfies this property—i.e., the Nash credible profile in this example.

4.6 The equilibrium tier structure

The equilibrium of the game \mathcal{G} cuts the outside option Gordian knot in a structured way: The coalitions that form in equilibrium can be organized into tiers in such a way that the limit sharing rule in each coalition satisfies the outside option principle of Binmore, Rubinstein, and Wolinsky (1986)—with the relevant outside options in each coalition endogenously determined by the Nash bargaining solution in coalitions that are in higher tiers. As a result,

(small) shocks to the primitives can propagate—via outside options—from higher to lower tiers, but not vice versa.

Formally, the *first-tier coalitions* are those coalitions $C \subseteq N$ that are such that, for each of its members i , χ_i is equal to i 's χ_{-i} -share in C .²⁵ The sharing rule in the first-tier coalitions converges to the Nash bargaining solution, with the Nash 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 were bargaining in isolation. 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) it is such that, for each of its members i who is not in the first, second, \dots , or $(\ell - 1)^{\text{th}}$ tier, χ_i is equal to i 's χ_{-i} -share in C . The sharing rule in each ℓ^{th} -tier coalition converges to the Nash bargaining solution, with the Nash threat points given by autarky, and the binding outside options determined in higher tiers. The ℓ^{th} -tier types are those that (i) are in an ℓ^{th} -tier coalition, and (ii) are not in any k^{th} -tier coalition, for any $k < \ell$.

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-maître 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 an ℓ^{th} -tier coalition affects negatively (positively) the $\ell + m^{\text{th}}$ -tier types when m is odd (even).²⁶ 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 small increase in the surplus of one does not hurt (benefit) the members of the other if the path is even (odd).²⁷

²⁵Note that there is always at least one first tier coalition. In particular, every perfect coalition is in the first tier (but other coalitions can be in the first tier as well).

²⁶An 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.

²⁷A path of a network is a sequence of distinct links. A path is even (odd) if it contains an even (odd) number of links.

4.7 Discussion

4.7.1 Relation to the outside option principle

The equilibrium sharing rule in each coalition satisfies the outside option principle of [Binmore, Rubinstein, and Wolinsky \(1986\)](#): Outside options in any given coalition enter as lower bounds on the bargaining shares in this coalition. In contrast to this earlier work, however, in the theory presented in this paper, the relevant outside options in each coalition are not given exogenously, but are themselves determined by the Nash bargaining solution in other coalitions. Indeed, it follows from [Corollary 4.3](#) that—when bargaining frictions are negligible—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* subject to the constraint that, in each coalition, everyone *else* receives *at least* the amount that the Nash credible payoff profile χ gives her. This outside option is binding in all coalitions except her χ -best coalitions.

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-maitre 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).

4.7.2 Holdup in coalition formation

[Pycia \(2012\)](#) describes how a holdup problem can occur in models where agents bargain first over which coalitions to form, and second over their terms of trade. 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. [Corollary 4.4](#) highlights that such a holdup problem does not arise in the present setting.

Corollary 4.4. *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.*²⁸

Intuitively, the maximum payoff that a type can justify using the Nash bargaining solution in some coalition subject to the constraint that others receive at least their equilibrium payoffs can only increase when a coalition that she is part of becomes more productive, when

²⁸I 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 $w = g \circ u$.

her bargaining power increases, or when she becomes less risk averse. This 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.

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-maître 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.

4.7.3 Relation to bargaining in stationary markets

This paper is related to the literature that studies stationary markets (e.g., [Rubinstein and Wolinsky 1985](#), [Gale 1987](#), [Binmore and Herrero 1988](#), [de Fraja and Sákovics 2001](#), [Manea 2011](#), [Lauermann 2013](#), [Nguyen 2015](#), [Polanski and Vega-Redondo 2018](#)). In order to investigate the strategic forces in a steady state of large dynamic markets, this literature typically assumes that the inflow of traders into the market perfectly matches its outflows. In terms of the framework, the innovation with respect to this literature is that, in this paper, the agents 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. For example, [Nguyen \(2015\)](#) uses convex programming techniques to characterize the stationary subgame-perfect equilibrium of a non-cooperative bargaining game where coalitions are not strategically chosen by the proposer, but that is otherwise similar to the one in the present paper.²⁹ The main comparative advantages of the present theory are tractability, sharp comparative statics, and a characterization of the equilibrium that is independent of the bargaining protocol: The sharing rule in each coalition converges to the classical Nash bargaining solution, with the outside options in each coalition determined by the Nash bargaining solution in other coalitions subject to a natural no-circularity condition that, remarkably, always pins down outcomes uniquely.

In [Talamàs \(2019\)](#), I use a version of the framework in the present paper tailored to net-

²⁹In addition to the above-mentioned difference in the protocol, in [Nguyen \(2015\)](#), all the types have linear utilities, while I assume instead that—as in the classical framework of [Nash \(1950\)](#)—preferences can be represented by (possibly heterogeneous) concave vN-M utility functions.

worked buyer-seller markets, 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. In particular, in [Talamàs \(2019\)](#), I provide a necessary and sufficient condition for the law of one price to hold in the limit as bargaining frictions vanish, and I describe an algorithm that decomposes the buyer-seller network into different submarkets, from the submarket with the highest limit price down (or, alternatively, the submarket with the lowest limit price up). The present paper provides a much richer characterization of bargaining outcomes that applies to a substantially more general environment with coalitions of arbitrary size and heterogeneous risk preferences. Moreover, it shows how the limit equilibrium outcome coincides with a natural generalization of the Nash bargaining solution that—by preventing outside options from being determined in a circular way—provides a unique and easy-to-compute prediction in complex stationary markets, and uncovers an endogenous vertical structure such that small shocks propagate from the top down, but not vice versa.

4.7.4 Relation to bargaining in non-stationary markets

This paper is also related to the literature investigating non-cooperative bargaining in non-stationary markets (e.g., [Chatterjee et al. 1993](#); [Ray and Vohra 1999](#); [Ray 2007](#); [Okada 1996, 2011](#); [Abreu and Manea 2012ab](#); [Dilmé 2018](#); [Elliott and Nava 2019](#)). In particular, both the system (1) of equations that determines the unique stationary perfect equilibrium payoffs and the algorithm \mathcal{A} that computes them can be seen as generalizations of the ones that determine the *no-delay* stationary perfect equilibrium in [Chatterjee et al. \(1993\)](#).³⁰ In addition to the stark contrast between the relevant economic environments (this literature focuses on settings without dynamic entry), one notable difference is that—as in the classical bargaining framework of [Nash \(1950\)](#)—I allow heterogeneous vN-M utility functions instead of restricting attention to the case of linear utilities, so the relevant solution concept in the limit as frictions vanish is the Nash bargaining solution instead of the egalitarian solution (e.g., [Dutta and Ray 1989](#)). More importantly, (i) the equilibrium that I characterize always exists and is the unique stationary perfect equilibrium (instead of being one of the many possible stationary perfect equilibria and existing only under certain conditions³¹), (ii) the

³⁰[Ray and Vohra \(1999\)](#) and [Ray \(2007\)](#) go substantially beyond the characterization of the no-delay stationary subgame-perfect equilibrium in [Chatterjee et al. \(1993\)](#) by allowing externalities across coalitions.

³¹Informally, a sufficient condition for a no-delay equilibrium to exist in the setting of [Chatterjee et al. \(1993\)](#) is that no one's expected equilibrium payoff in such equilibrium increases when the set of active players shrinks (see for example condition M in [Ray \(2007, section 7.4\)](#)). This condition is not relevant in the present

equilibrium in the present setting is stationary (instead of evolving as different coalitions form), (iii) the predictions of which coalitions form in equilibrium and everyone's resulting limit payoffs do not depend on the order realization of the proposers, and (iv) I characterize the equilibrium payoff profile as being the only one that satisfies a natural credibility property—which, in the limit as bargaining frictions vanish, provides a natural generalization of the Nash bargaining solution to settings where agents bargain both about which coalitions to form and how to share the resulting surplus.

The connection between the strategic forces in [Chatterjee et al. \(1993\)](#) 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, this earlier work focuses on a *rejector-proposes protocol* (in which the rejector of a proposal becomes the proposer in the next period) instead of the *random-proposer protocol* (with possibly heterogeneous proposal probabilities) of the present paper. [Ray](#) (section 7.7; 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 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 with the bargaining frictions. Not surprisingly given the qualitative differences between the settings, the predictions of the resulting theories are qualitatively distinct. For example, the endogenous evolution of the market in [Chatterjee et al. \(1993\)](#) 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).

4.7.5 Relation to other Nash bargaining approaches to coalition formation

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

setting because the dynamic entry of agents ensures that the relevant matching opportunities are constant over time. As a result, in contrast to this previous work, no conditions need to be put on the production function to guarantee equilibrium existence.

agent's disagreement point being the maximum that she can achieve in any other match (keeping the others' payoffs fixed).^{32,33} [Burguet and Caminal \(2018\)](#) show that a modification and extension of this idea (in a context in which only one coalition can form) uniquely pins down the agents' payoffs, and provide strategic foundations for the resulting coalition formation solution concept.³⁴ 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.

4.7.6 Relation to Nash-in-Nash

[Collard-Wexler et al. \(2019\)](#) 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. In particular, being a surplus division rule *for a given network*, the Nash-in-Nash solution does not generally pin down the equilibrium network.

³²[Rochford \(1984\)](#) 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).

³³The idea of endogenizing the Nash 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](#), [Peleg 1986](#), [Hart and Mas-Colell 1989](#), [Serrano and Shimomura 1998](#)). See also [Moldovanu \(1993\)](#) and [de Fontenay and Gans \(2014\)](#) for alternative approaches to endogenous outside options in bilateral settings.

³⁴Relatedly, [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.

5 Vertically differentiated markets

I now turn to describing comparative statics in the context of two-sided pairwise matching markets, in which types are vertically differentiated either by skill, risk aversion, or bargaining power. For ease of exposition, I refer to the types on one side as *workers* and the types on the other side as *firms*. Also, 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-firm pair (i, j) , the surplus $y(i, j)$ is strictly increasing in its members' skills.³⁵ I say that two types i and j *match* if agents of type i match with agents of type j in equilibrium.

5.1 Shock propagation from the top down

Corollary 5.1 below illustrates how—in settings where the agents are vertically differentiated by skill but are otherwise identical—the bargaining outcomes are determined from the most skilled types down.

Corollary 5.1. *When all the workers have the same risk aversion and bargaining power, an increase in the skill of worker i from s_i to $s'_i > s_i$ does not affect the payoff of any worker j whose skill s_j is strictly higher than s'_i .*

Indeed, 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 skilled workers. This implies that an increase in a type's productivity does not affect the payoffs of the more productive types. In other words, the relevant outside options of a firm that matches with a worker i are never determined by bargaining with a less skilled worker and—as a result—the productivity of workers whose skills are lower than i 's do not affect i 's payoff.

Corollary 5.1 implies that when the source of heterogeneity among agents is their productivity, the bargaining outcomes are determined *from the types with the highest payoffs down*. In contrast, **Corollary 5.2** shows that, when the source of heterogeneity among agents is either their risk aversion or their bargaining power, the bargaining outcomes are determined *from the types with the lowest payoffs up*.

Corollary 5.2. *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*

³⁵Formally, let i and i' be any two workers and let j be any firm. 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 10: Equilibrium matching pattern in the example when the surplus of the cook-maître coalition is 80 instead of 60.

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 .

Indeed, 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. In other words, the relevant outside options of a firm that matches with a worker i are never determined by bargaining with a less risk averse worker and—as a result—the risk aversion of workers whose risk aversion is higher than i 's do not affect i 's payoff. 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.

5.2 Shock propagation under positive assortative matching

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 firm 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 the cook-maître match 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 10](#) 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 10 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 two-sided 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 profiles of types if and only if their skills are complementary—in the sense that the match function is supermodular. Corollary 5.3 is the analogous result in the present setting.³⁶

Corollary 5.3. *The equilibrium features positive assortative matching for all profiles of types if and only if the production function is strictly supermodular.*³⁷

To see the intuition behind Corollary 5.3, consider again the situation illustrated in Figure 10. 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 10, 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 5.1 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.

³⁶In the present setting, a profile of types consists of a finite set W of workers, a finite set F of firms, and a profile $\{s_i\}_{i \in W, F}$ of corresponding skill parameters.

³⁷The 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 $y(i, j) + y(i', j') > y(i', j) + y(i, j')$.

Corollary 5.4 below puts further structure on how shocks propagate under positive assortative matching.

Corollary 5.4. *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.*

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. Indeed, under positive assortative matching, algorithm \mathcal{A} pins down the workers' (firms') 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 5.4 contrasts with the case in which the production function is strictly submodular.³⁸ Indeed, in this case, all the types that match do so with the most productive type on the other side of the market.³⁹ 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

³⁸The 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 $y(i, j) + y(i', j') < y(i', j) + y(i, j')$.

³⁹To see this in the context of the example, suppose for contradiction that the surplus function is submodular and that chefs and maitres 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-maitre coalition. In particular, at least one of these coalitions is more profitable than the sum of its members' payoffs, a contradiction.

as in [Figure 10](#), an increase in the surplus of the chef-manager coalition from 80 to 90 is fully absorbed by the managers.

6 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.

The objective of this paper has been to investigate how prices and allocations are determined in stationary 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 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 the constrain that *everyone else gets at least her equilibrium payoff*. This theory provides a tractable framework to investigate the determinants of bargaining power in stationary markets. In particular, an intuitive algorithm finds which coalitions form and how they share the resulting surplus and uncovers an endogenous vertical structure such that different economic shocks can propagate via outside options from the top down, but not vice versa. An interesting avenue for future work is to investigate how Nash's axioms can be extended to the present setting to characterize this generalization of the Nash bargaining solution.

A Appendix: Details omitted from section 4

In subsection A.1, I rewrite the system of equations (1) that any equilibrium cutoff profile w in the bargaining game \mathcal{G} satisfies. In subsection A.2, I characterize the stationary subgame-perfect equilibrium of the auxiliary games \mathcal{G}_x^C . In subsection A.3, I provide the proofs of Propositions 4.2, 4.3 and 4.4. Finally, in subsection A.4, I provide the proofs of Proposition 4.5, Theorem 1 and Corollary 4.3.

A.1 Rewriting the system of equations (1)

I start by defining the function that I use to compactly rewrite system (1) as (4).

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 $\alpha_i := \frac{q}{(1-q)p_i+q}$. Figure 11 illustrates.

In words, $f_i(x)$ is the amount that an agent of type i can obtain when she is the proposer in a stationary equilibrium conditional on her being indifferent between accepting and rejecting the amount x .⁴⁰ 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 $f_i(x) > x$ for all $x > y_i$. 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 11). System (1) can be written as

$$(4) \quad f_i(w_i) = \max_{C \subseteq N} \left(y(C) - \sum_{j \in C-i} w_j \right) \text{ for all } i \text{ in } N.$$

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

I now turn to characterizing the equilibrium of the auxiliary game \mathcal{G}_x^C . Naturally, the fact that there is only one relevant coalition in this game substantially simplifies the system that pins down its equilibrium cutoff profile. Indeed, I now use a relatively simple argument to show that this cutoff profile exists and is unique (Lemma A.1).⁴¹

⁴⁰This interpretation relies on the argument used in section 4 to derive the system (1), which, for brevity, I do not reproduce here.

⁴¹This result is close existing results in the literature, but 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




Figure 11: Illustration of the definition of the function f_i (Definition A.1).

Note A.1. The game \mathcal{G}^C defined in subsection 4.2 corresponds to the game \mathcal{G}_x^C defined in subsection 4.3, for $x = 0$.

Lemma A.1. *For any coalition $C \subseteq N$ and any profile x in \mathbb{R}^N satisfying $\sum_{j \in C} x_j \leq y(C)$, the game \mathcal{G}_x^C has unique stationary subgame-perfect equilibrium. As the bargaining friction q goes to zero, the equilibrium cutoff profile converges to the profile that solves*

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

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 (1) shows that the equilibrium cutoff profile v satisfies

$$(6) \quad f_i(v_i) = \max \left[f_i(x_i), y(C) - \sum_{j \in C-i} v_j \right] \text{ for all } i \text{ in } C.$$

Existence of a solution to system (6) follows from Brouwer's fixed point theorem. To prove uniqueness, suppose for contradiction that there are two profiles v, v' that solve system (6). 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 $\{f_j(v'_j) - v'_j\}_{j \in S}$. Since $f_i(v_i) - v_i$ is increasing in v_i , we also have that $v_i > v'_i$. Moreover we have that

$$(7) \quad f_i(v_i) = y(C) - \sum_{j \in C-i} v_j,$$

transferable utility), so I provide the proof of Lemma A.1 below. This argument is a slight generalization of standard arguments in the literature (e.g., Ray, 2007, Chapter 7) to the setting without perfectly transferable utility considered in this paper.

since, otherwise, $v_i = x_i$, which contradicts the fact that $v_i > v'_i \geq x_i$. In particular,

$$f_j(v_j) - v_j \geq y(C) - \sum_{j \in C} v_j = f_i(v_i) - v_i \geq f_j(v'_j) - v'_j \text{ for all } j \text{ in } S,$$

or, using again that $f_j(v_j) - v_j$ is increasing in v_j , $v_j \geq v'_j$ for all j in S . But then, Equation 7 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. For brevity, I omit the proof of the characterization of the limit equilibrium payoffs in terms of the Nash bargaining solution; it is the analog of the main result in Binmore, Rubinstein, and Wolinsky (1986) in the setting of the present paper. \square

A.3 Proof of Proposition 4.2, Proposition 4.3 and Proposition 4.4

A.3.1 Proof of Proposition 4.2

Let w be a solution of system (1), and fix an arbitrary type i . Let $C \subseteq N$ be such that $f_i(w_i) - w_i = y(C) - \sum_{j \in C} w_j$, and let v be the equilibrium cutoff profile in \mathcal{G}_0^C . By definition, we have that $f_i(v_i) - v_i = y(C) - \sum_{j \in C} v_j$. Suppose for contradiction that $v_i < w_i$. This implies that there exists $k \in C$ such that $v_k > w_k$. Since $f_i(v_i) - v_i$ is increasing in v_i , this implies that $f_i(v_i) - v_i < f_i(w_i) - w_i$ and $f_k(v_k) - v_k > f_k(w_k) - w_k$, which contradicts the fact that $f_i(v_i) - v_i = y(C) - \sum_{j \in C} v_j = f_k(v_k) - v_k$ and $f_k(w_k) - w_k \geq y(C) - \sum_{j \in C} w_j = f_i(w_i) - w_i$. \square

A.3.2 Proof of Proposition 4.3

Let w be a solution of system (1), let C be a perfect coalition, and let v be the equilibrium cutoff profile in \mathcal{G}_0^C . Suppose for contradiction that there exists $i \in C$ such that $w_i < v_i$. This implies that there exists $k \in C$ with $w_k > v_k$, which contradicts Proposition 4.2. \square

A.3.3 Proof of Proposition 4.4

Given the discussion of Proposition 4.4 in the main text, it only remains to show that 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 bigger than i 's cutoff in the auxiliary game \mathcal{G}^D , then j 's cutoff in the auxiliary game \mathcal{G}^C is bigger than j 's cutoff in the auxiliary game \mathcal{G}^D . To see that this is the case, let v^C and v^D be the equilibrium cutoff profile in \mathcal{G}^C and \mathcal{G}^D , respectively. We have

that

$$f_i(v_i^C) - v_i^C = y(C) - \sum_{j \in C} v_j^C = f_j(v_j^C) - v_j^C \text{ and } f_i(v_i^D) - v_i^D = y(D) - \sum_{j \in D} v_j^D = f_j(v_j^D) - v_j^D.$$

Hence, given that $f_i(x) - x$ is increasing in x , $v_i^C \geq v_i^D$ implies that $v_j^C \geq v_j^D$. \square

A.4 Proofs of Proposition 4.5, Theorem 1 and Corollary 4.3

Proposition A.1 and Proposition A.2 together prove Proposition 4.5 and Theorem 1. Corollary 4.3 follows from the combination of Lemma A.1 above and Proposition A.3 below.

Proposition A.1. *A profile \mathbf{x} in \mathbb{R}^N solves system (4) if and only if it is credible.*

Proof. Let the profile \mathbf{x} in \mathbb{R}^N be credible, and fix an arbitrary type i . Given that $f_i(x_i) \leq \max_{C \subseteq N} (y(C) - \sum_{j \in C-i} x_j)$, we only need to show that there exists a coalition $C \subseteq N$ with $f_i(x_i) \geq y(C) - \sum_{j \in C-i} x_j$, but this is satisfied by i 's \mathbf{x}_{-i} -best coalition.

In the other direction, suppose that the profile \mathbf{x} in \mathbb{R}^N solves system (4), and fix an arbitrary type i in N . Since, $x_i = \max_{C \subseteq N} (y(C) - \sum_{j \in C-i} x_j)$, i 's \mathbf{x}_{-i} -best share is bounded above by x_i , so letting $C \subseteq N$ be such that

$$(8) \quad f_i(x_i) = y(C) - \sum_{j \in C-i} x_j,$$

it is enough to show that i 's \mathbf{x}_{-i} -share in C is bounded below by x_i . Let the profile \mathbf{v} in \mathbb{R}^C solve

$$\begin{aligned} f_i(v_i) &= y(C) - \sum_{j \in C-i} v_j \\ f_k(v_k) &= \max \left[f_k(x_k), y(C) - \sum_{j \in C-k} v_j \right] \text{ for all } k \text{ in } C - i. \end{aligned}$$

Suppose for contradiction that $v_i < x_i$. Since $v_j \geq x_j$ for all $j \in C - i$, this implies that $v_j > x_j$ for some type j in C . Using that $f_i(v_i) - v_i$ is increasing in v_i , Equation 8, and that \mathbf{x} solves system (6), we get

$$f_j(v_j) - v_j = f_i(v_i) - v_i < f_i(x_i) - x_i = y(C) - \sum_{j \in C} x_j \leq f_j(x_j) - x_j$$

which implies that $v_j < x_j$, a contradiction. \square

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

Proof. First, I prove by induction in the step number s that algorithm \mathcal{A} only updates payoffs that are 0 (i.e., if $x_i^s \neq 0$, then $x_i^{s+1} = x_i^s$). The base step is vacuous. For the induction step, consider a step s of algorithm \mathcal{A} , and suppose that, if $x_i^{s-1} \neq 0$, then $x_i^s = x_i^{s-1}$. Given that the outside options that have to be met only increase from step s to step $s+1$, every x^{s-1} -perfect coalition is an x^s -perfect coalition, so $x_i^{s-1} \neq 0$ implies that $x_i^{s+1} = x_i^s$.

Second, note that in every step s of algorithm \mathcal{A} , $x^{s-1} < x^s$, so algorithm \mathcal{A} ends in finitely many steps. Indeed, the argument analogous to the one behind [Proposition 4.4](#) restricted to the set $K_s = \{i \in N \mid x_i^s = 0\}$ shows that there exists at least one coalition C that is an x^s -perfect coalition and contains at least one of the types in K_s .

Finally, I prove by induction in the step number s that, for each $s \leq S$, every credible profile gives x_i^s to each type i in $N - K_s$ (so that χ is the only possible credible profile). Let z be a credible payoff profile. Let s be such that z_i is equal to x_i^{s-1} for each type i in $N - K_{s-1}$ (this induction hypothesis is vacuously true when $s = 1$, so there is no need to prove the base step separately). Suppose for contradiction that, for some i in K_{s-1} , x_i^s is strictly bigger than z_i (the induction hypothesis together with fact that algorithm \mathcal{A} only updates outside options upwards implies that x_i^s cannot be strictly smaller than z_i). In other words, there exists a coalition C such that i 's z_{-i} -share in C is strictly smaller than i 's x^{s-1} -share in C . This implies that for some j in K_{s-1} , z_j is strictly bigger than j 's x^{s-1} -share in C (that is, j 's x^{s-1} -best share), which, as just argued, contradicts the induction hypothesis. \square

The algorithm \mathcal{A}^* below 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 A.2 (Algorithm \mathcal{A}^*). Let $x^0 = 0$. In each step $s = 1, 2, \dots$, for each type i in an x^{s-1} -Nash perfect coalition, let x_i^s be i 's x^{s-1} -Nash best share; for each other type j , let $x_j^s = 0$. Stop in the first step S such that every type is in an x^S -Nash perfect coalition, and let $\chi^* := x^S$.

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 s with $K_s \neq \emptyset$, there is at least one x^s -perfect coalition containing at least one type in K_s . The argument in this case is analogous to that in [Pycia \(2012, pages 330-331\)](#).⁴² Denoting, for each coalition C , type i 's x^{s-1} -Nash share in C by x_i^C , and letting u_i' denote the derivative of the utility function u_i , we have that

⁴²[Pycia \(2012\)](#) 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).

$u_i(x_i^C)/u'_i(x_i^C)$ is the same for every type i in $C \cap K^{s-1}$; denote by μ_C this common value. A coalition C with maximum μ_C is i 's \mathbf{x}^s -Nash best coalition for each $i \in K_s$, since each type i '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. AND D. AUTOR (2011): "Skills, tasks and technologies: Implications for employment and earnings," in *Handbook of Labor Economics*, Elsevier, vol. 4, 1043–1171.
- BAUMOL, W. J., W. G. BOWEN, ET AL. (1966): *Performing arts-the economic dilemma: a study of problems common to theater, opera, music and dance.*, MIT Press.
- 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*.

- CARD, D., J. HEINING, AND P. KLINE (2013): "Workplace heterogeneity and the rise of West German wage inequality," *The Quarterly Journal of Economics*, 128, 967–1015.
- 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 cyclical cost of the opportunity cost of employment," *Journal of Political Economy*, 124, 1563–1618.
- COLLARD-WEXLER, A., G. GOWRISANKARAN, AND R. LEE (2019): "'Nash-in-Nash' Bargaining: A Microfoundation for Applied Work," *Journal of Political Economy*, 127, 163–195.
- COMPTE, O. AND P. JEHL (2010): "The coalitional Nash bargaining solution," *Econometrica*, 78, 1593–1623.
- DE FONTENAY, C. C. AND J. S. GANS (2014): "Bilateral bargaining with externalities," *The Journal of Industrial Economics*, 62, 756–788.
- DE FRAJA, G. AND J. SÁKOVICS (2001): "Walras retrouvé: Decentralized trading mechanisms and the competitive price," *Journal of Political Economy*, 109, 842–863.
- DILMÉ, F. (2018): "Bargaining and delay in thin markets," *Mimeo*.
- DUTTA, B. AND D. RAY (1989): "A concept of egalitarianism under participation constraints," *Econometrica*, 615–635.
- EECKHOUT, J. (2018): "Sorting in the Labor Market," *Annual Review of Economics*, 10, 1–29.
- ELLIOTT, M. (2015): "Inefficiencies in networked markets," *American Economic Journal: Microeconomics*, 7, 43–82.
- ELLIOTT, M. AND F. NAVA (2019): "Decentralized bargaining in matching markets: Efficient stationary equilibria and the core," *Theoretical Economics*, 14, 211–251.
- ELLIOTT, M. AND E. TALAMÀS (2019): "No holdup in dynamic markets," *Mimeo*.
- GALE, D. (1987): "Limit theorems for markets with sequential bargaining," *Journal of Economic Theory*, 43, 20–54.

- 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.
- HORN, H. AND A. WOLINSKY (1988): “Bilateral monopolies and incentives for merger,” *The RAND Journal of Economics*, 408–419.
- JÄGER, S., B. SCHOEFER, S. YOUNG, AND J. ZWEIMÜLLER (2018): “Wages and the value of nonemployment,” *Mimeo*.
- KALAI, E. AND M. SMORODINSKY (1975): “Other solutions to Nash’s bargaining problem,” *Econometrica*, 43, 513–518.
- 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.
- 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.
- MANEA, M. (2011): "Bargaining in stationary networks," *American Economic Review*, 101, 2042–2080.
- (2016): "Models of Bilateral Trade in Networks," *The Oxford Handbook of the Economics of Networks*, (Yann Bramoullé, Andrea Galeotti, and Brian Rogers, eds.).
- MOLDOVANU, B. (1993): "Price indeterminacy and bargaining in a market with indivisibilities," *Journal of Mathematical Economics*, 22, 581–597.
- NASH, J. (1950): "The bargaining problem," *Econometrica*, 155–162.
- (1953): "Two-person cooperative games," *Econometrica*, 128–140.
- NGUYEN, T. (2015): "Coalitional bargaining in networks," *Operations Research*, 63, 501–511.
- OKADA, A. (1996): "A noncooperative coalitional bargaining game with random proposers," *Games and Economic Behavior*, 16, 97–108.
- (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.
- 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.
- SONG, J., D. J. PRICE, F. GUVENEN, N. BLOOM, AND T. VON WACHTER (2019): "Firming up inequality," *The Quarterly Journal of Economics*, 134, 1–50.
- 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. (2019): "Price dispersion in stationary networked markets," *Games and Economic Behavior*, 115, 247–264.