

BARGAINING WITH MECHANISMS

MARCIN PEŃSKI

ABSTRACT. We consider an alternating-offer model of bargaining over a heterogeneous pie, with one-sided incomplete information about preferences, and where players can offer arbitrary mechanisms to determine the allocation. When the pie has two parts and offers are frequent, there is a unique limit of Perfect Bayesian Equilibrium outcomes: the uninformed player proposes the optimal screening menu subject to the constraint that each type of the informed player receives at least her payoff under complete information. The optimal menu can be implemented with four allocations. With more than two dimensions, there exist equilibria in which the informed player may receive strictly less than her complete information benchmark.

1. INTRODUCTION

In a standard model of bargaining, one party proposes an allocation of the bargaining surplus and the other party either accepts or rejects it. However, offers made during real-world negotiations are often much more complex. Instead of a single allocation, parties may offer menus of allocations for the other party to choose from.¹ They

Date: April 23, 2021.

Department of Economics, University of Toronto. I am grateful for helpful comments from V Bhaskar, Laura Doval, Thomas Tröger, Thomas Wiseman, and various conference and seminar participants. I gratefully acknowledge financial support from the Insight Grant of the Social Sciences and Humanities Research Council of Canada and the hospitality of HEC Paris, where some of this research was conducted.

¹See Jackson et al. (2018) for real-world and experimental examples. I had an opportunity to observe bargaining over a pension plan reform that took place in 2016-18 between three Ontario universities and the representatives of faculty and staff. Among other issues, the parties negotiated the size of

may offer to settle the dispute with an arbitrator.² They may also offer to alter the bargaining protocol, for example, by dividing the dispute into smaller areas and settling them separately, or by establishing deadlines.³ We teach our students (and our children) that a fair cake division can be found through simple procedures like “I divide and you choose.”⁴ All such offers can be represented as a mechanism, the outcome of which determines the final allocation. The goal of this paper is to study the role of mechanisms as offers in a strategic model of bargaining by addressing the following questions: Does expanding the scope of offers to general mechanisms affect the way in which parties bargain? Which mechanisms are offered in equilibrium? Is the equilibrium efficient?

A natural setting for studying mechanisms as offers is when the object of bargaining is complex (multi-dimensional) and there is incomplete information about player preferences. To stay as close as possible to the existing literature, we consider a version of Rubinstein’s alternating-offer model (Rubinstein (1982)). There are two players, Alice and Bob, who want to divide a heterogeneous pie with $N \geq 2$ parts (for instance, chocolate, strawberry, etc.). Bob’s preferences over different parts are known. Alice’s preferences are linear, but otherwise Bob can have arbitrary beliefs about Alice’s preferences. In alternating periods, each player offers a mechanism, which the other player accepts or rejects. A mechanism is defined as an arbitrary game, where players’ choices determine the final allocations. When the offer is accepted, the mechanism is

the spousal benefit, early retirement options, inflation indexation, etc. While the universities only cared about the total actuarial cost, the preferences of the labor side were uncertain, mostly due to its heterogeneity (for instance, the staff, but not the faculty, valued early retirement more than the spousal benefit). Ultimately, the universities proposed a menu of options, and the labor side chose an option from this menu.

²During the 2019-2020 dispute between the Ontario government and teacher unions, both parties called upon the other to accept mediation but could not agree on the same mediator (Rushowy (2020), Moodie (2019)).

³EU accession negotiations typically take the form of independent bargaining over 30-40 areas.

⁴An example of such mechanism is the Texas shoot-out clause used in the dissolution of a partnership: one partner names a price and the other partner is obliged to either sell her shares or buy the shares of the first partner at the price. I am grateful to T. Tröger for this example.

implemented and the bargaining ends. If the offer is rejected, the other player makes an offer in the next period. Players cannot commit to offers in subsequent periods. We study Perfect Bayesian Equilibria (PBE) with the only restriction that Bob's off-path beliefs about Alice's types do not change after his actions.

Any strategic model of bargaining under incomplete information must deal with two types of problems. Due to a screening problem, a player's offer may be acceptable to some but not all types of the opponent. This may lead to a delay and a new offer for the remaining types, which may change the incentives to accept the original one. Due to a signaling problem, an agent may accept or make an unfavorable offer because off-path deviations are punished with beliefs that lead to a low continuation payoff. The signaling problem typically leads to multiplicity of equilibria that can sometimes be resolved through equilibrium refinements.

We show that, when players are allowed to offer arbitrary mechanisms, both screening and signaling problems have a satisfactory solution. Our main result is that, when $N = 2$ and offers are frequent, there is a unique limit of PBE outcomes. The limit outcome is equivalent to Bob proposing a screening menu Y^* that is optimal for him, subject to the constraint that each of Alice's types receives at least her complete information payoff or, equivalently, her Nash bargaining payoff. To implement the optimal menu, Bob needs no more than 4 allocations. The final outcome is *ex-ante*, but not *ex-post*, efficient. The solution has natural comparative statics with respect to information: Bob is better off when his information improves. When Bob's beliefs converge to certainty, the outcome converges to the complete information Nash solution.

The proof parallels the argument for the uniqueness of the subgame perfect equilibrium payoffs in Rubinstein's alternating-offer game. We develop step-by-step bounds to sandwich the equilibrium payoffs. As offers become frequent, the lower and upper

bounds converge to the same outcome. Two types of mechanisms play a role in the proof. On the one hand, Bob’s ability to offer menus (of allocations, for Alice to choose) allows him to screen among Alice’s types without them worrying about revealing information. On the other hand, Alice’s ability to offer menus (for Bob to choose) of menus (of allocations, for Alice to choose) allows her to protect herself from “punishment with beliefs”. To see a simple intuition for the latter point, suppose that Alice considers an off-path deviation to one of two mechanisms $m \in \{m_1, m_2\}$ with the property that, for each of Bob’s beliefs, one of the two mechanisms would be acceptable to Bob, but none of them is acceptable across all his beliefs. She can be stopped from such a deviation if she is afraid that after off-path offer m , Bob’s beliefs will change to those that find m unacceptable. Such punishment with beliefs would not be possible if she were able to offer a menu $\{m_1, m_2\}$ of mechanisms and let Bob choose whichever mechanism he prefers.

The main result is surprising for at least three different reasons. First, because both the informed and uninformed agents design mechanisms, our model is an example of a dynamic informed principal problem (Myerson (1983)). The uniqueness without any equilibrium refinement is a rare result in the informed principal literature where, typically, multiple equilibria can be supported by belief punishment threats (Mylovanov and Tröger (2012)).

The availability of sophisticated offers plays an important role for uniqueness. If players are only able to offer simple allocations, we show that there may exist multiple equilibria, including an Anti-Coasian one, where each Alice type receives her worst possible payoff across all of Bob’s possible beliefs and Bob receives his best possible payoff. The construction of such an equilibrium involves punishing Alice’s deviations with beliefs that her type is the worst for her (but best for Bob).

Second, although assumptions explicitly disallow commitment across periods, the equilibrium outcome is the same as if Bob could commit himself to any mechanism subject to the constraint that each Alice type receives at least her Nash bargaining payoff. The constraint is clearly a consequence of the connection between the Rubinstein’s model and the Nash solution.

Third, the main result can be contrasted with the Coase conjecture, which predicts that the informed player has all the advantage, the equilibrium is efficient, and the uninformed player receives the worst outcome across all possible types of the informed player. A companion paper, Peski (2019), studies war-of-attrition bargaining in a similar environment, except that players have additional ability to commit to their offers due to reputational types. Interestingly, more commitment leads to a Coasian-type result: in the unique (rational and patient limit) equilibrium, Bob proposes a menu \mathcal{C} of all allocations that give him at least his worst possible complete information payoff. Bob is typically strictly worse off than under the optimal menu Y^* ; Alice types are better off, some of them strictly so. The disparity between alternating-offer and reputational versions of the model is striking to a reader familiar with Abreu and Gul (2000).⁵

Although most of the literature is restricted to the two-dimensional case (Jackson et al. (2018) and Peski (2019) being among the exceptions), we also look at higher dimensions. If $N > 2$, the main result does not hold. We construct an equilibrium, where some Alice types receive a payoff strictly lower than her Nash payoff. This is of interest in itself, as Maskin and Tirole (1990) claims that, in the private value case,

⁵In the bargaining literature, the Coase conjecture has been established in the “gap case” of the durable monopoly problem with seller-only offers (Gul et al. (1986)). In the “no-gap case” and seller-only offers, there is a folk theorem of payoffs that reduces to the Coase conjecture under the Markovian assumption (Ausubel and Deneckere (1989b)). Ausubel and Deneckere (1989a) and Ausubel et al. (2002) show that a refinement is also needed for the uniqueness of the equilibrium in the alternating-offer case of the buyer-seller environment.

the informed principal must benefit from incomplete information due to the collapse of agent incentive and individual rationality constraints.⁶ This observation may fail with interdependent values. Although we work with private values, the dynamic setting's continuation payoffs typically depend on the belief of the uninformed agent, which leads to endogenous interdependence.

At this moment, we are not able to give a full characterization of the equilibrium set for general N . Instead, we find the payoffs bounds for both players. In particular, we show that each player and each type can at least ensure his or her worst possible Nash payoff across all possible opponent types. Hence, the equilibrium payoffs are bounded by the worst possible complete information payoffs.

Mechanisms as offers have been considered in the axiomatic theories of bargaining in Harsanyi and Selten (1972) and Myerson (1984). Certain mechanisms, like menus, also appear in some work on strategic bargaining under one-sided incomplete information. With an exception of Jackson et al. (2018), all related papers that rely on sophisticated offers and that we are aware of work solely with two types. Sen (2000) (see also Inderst (2003)) studies a two-type alternating offer game, where players can offer menus but not general mechanisms and demonstrates the existence of a unique outcome in a refinement of PBE (perfect sequential equilibrium due to Grossman and Perry (1986)). The equilibrium behavior depends on whether the high type prefers her own complete information Nash payoff, or the Nash allocation of the low type. In a similar bargaining environment, Wang (1998) studies the Coasian bargaining model with Bob making all the offers. He shows that, in the unique equilibrium, Bob separates Alice's two types with an optimal screening contract. In particular, the Coase conjecture fails as Bob retains all power subject to the incentive compatibility constraints. More recently,

⁶I am grateful to V Bhaskar for this observation.

Strulovici (2017) assumes that, instead of ending the game, any accepted offer becomes the status quo for future bargaining. In this setting, the Coase conjecture holds and the uninformed player is unable to offer an inefficient payoff to type u'_1 in order to screen out the more extreme type u''_1 .

Jackson et al. (2018) considers a general bargaining environment. Although the authors allow for incomplete information on both sides, they make a strong assumption that the total value of bargaining surplus is commonly known. This assumption implies that there are no incentive problems that stop agents from truthfully revealing their information. In the unique equilibrium, the agents use menus to implement information revelation in a single round of bargaining.

An important assumption of our model is that although players cannot commit to any offer in subsequent periods, once the mechanism is offered and accepted, the players are committed to its implementation. Thus, our assumption resembles the recent literature on dynamic mechanism design with limited commitment (Skreta (2006), Doval and Skreta (2018), and others), but with some differences. First, we do not allow for renegotiation of an inefficient outcome, while in Doval and Skreta (2018), if a good is not traded in one period, it can be traded in the future. We discuss this issue in a more detail below. Second, we allow both the uninformed and informed players to offer mechanisms. To the best of our knowledge, ours is the first paper to study the informed principal problem in a dynamic setting either with either limited or no commitment.

2. MODEL

2.1. Environment. Two players, Alice and Bob, bargain over a heterogeneous pie with $N \geq 2$ parts. An allocation is defined as a tuple $x = (x_{i,n}) \in X = \{x \in [0, 1]^{2N} : \forall_n \sum_i x_{i,n} \leq 1\}$,

where $x_{i,n}$ is player i 's share of the n th part of the pie. An allocation x is wasteful if $\sum_i x_{i,n} > 1$ for some n ; otherwise, it is not wasteful. (Dis)allowing for wasteful allocations does not affect our results. An allocation where player i obtains the entire pie is denoted with $\mathbf{1}_i = \mathbf{0}_{-i}$. The main result is about $N = 2$. Section 6.2 discusses what we know about $N > 2$.

Both players have linear preferences over allocations. Alice's utility from allocation x is equal to $u(x) = \sum_n u_n x_{A,n}$, where $u \in \mathcal{U} = \{u \in [0, 1]^N : \sum_n u_n = 1\}$. We normalize the preferences so that the coefficients add up to 1. Alice's utility is privately known and Bob's beliefs about Alice's preferences are denoted with $\mu \in \Delta\mathcal{U}$. Bob's utility is publicly known and equal to $v(x) = \sum_n v_n x_{B,n}$ for $v \in \mathcal{U}$.

2.2. Bargaining game. In alternating periods, one player offers to choose an allocation with a mechanism m , and the other player either accepts or rejects. The first offer is made by player j . If the offer is accepted, mechanism m is implemented, the allocation is determined in a continuation equilibrium of the mechanism, and the game ends with players receiving their respective payoffs from the allocation. If the offer is rejected, the game moves onto the next period, with the other player making an offer. The players discount with a common factor $\delta < 1$.

All actions (mechanism choices and acceptance decisions) are perfectly observed. In order to ensure the existence of an equilibrium, we assume that, in each period, (a) players observe a public randomization device, and (b) Alice is able to send cheap talk messages from a sufficiently rich set. Because our result shows that the equilibrium outcome is unique, it is not weakened by allowing for public randomization and cheap talk.

2.3. Mechanisms. A *mechanism* is any normal-form or extensive-form game such that the action choices determine the final allocation in X . Formally, a mechanism is a tuple $m = ((S_i^t)_{i=A,B}^{t \leq T}, \chi)$, where $T \leq \infty$, S_i^t is a set of actions for player i in period t , and $\chi : \prod_i^T S_i^t \rightarrow X$ is an allocation function. Examples include:

- *simple offers*: players do not make any choices and receive a predetermined allocation;
- *(Alice's) menus*: Alice chooses an allocation $x \in Y$ from a closed set of allocations $Y \subseteq X$. Let \mathcal{Y} be the space of all menus;
- *(Bob's) menus of (Alice's) menus*: Bob chooses one of (Alice's) menu $Y \in W$ from a (Hausdorff topology) closed set of menus $W \subseteq \mathcal{Y}$, followed by Alice who chooses an allocation from the menu;
- original bargaining game, or any alteration of the bargaining protocol of the original game. ⁷

Let \mathcal{M} be the space of mechanisms available to players. We assume that \mathcal{M} contains all menus and menus of menus. The statement of the main result refers only to menus. The proof relies heavily on the availability of menus of menus of a particular kind. Whether \mathcal{M} contains any other mechanisms is irrelevant for the results and proofs.

It is important to clearly describe the role of commitment in our model. As in the Coasian bargaining literature, or in the literature on dynamic mechanism design with either limited or no commitment (Skreta (2006), Doval and Skreta (2018), Liu et al. (2019)), we assume that while players *cannot* commit themselves to future offers *unilaterally*, they are committed to implementing it once the mechanism is proposed and

⁷For example, a menu is a mechanism with $T = 1$ period and no non-trivial choices for Bob, $|S_B^1| = 1$. A menu of menus is a mechanism with $T = 2$ periods, no non-trivial choices for Alice in the first period, $|S_A^1| = 1$, and no non-trivial choices for Bob in the second period, $|S_B^2| = 1$.

accepted. However, we also allow for a wider range of mechanisms than this literature typically considers. For example, an agreement on negotiation protocol may force players to restrict their future options, set a deadline, or choose an ex-post inefficient outcome. In other words, we allow players to commit *jointly*. Our model is applicable in situations in which such a commitment is possible, either because the nature of dividing the surplus makes it impossible to divide it again, or because renegotiation is costly, or the agreement is enforced by an arbitrator or a court. Section 6.4 discusses what results stay the same and what changes under renegotiation.

2.4. Equilibrium. An *assessment* is a tuple of strategies for each player and Bob's belief function that specifies Bob's beliefs after any history that ends with Alice's action. A *Perfect Bayesian equilibrium* (or, simply, equilibrium) is an assessment in which (a) the players best respond to the opponent's strategy, and, in Bob's case, given his beliefs, and (b) at each decision point, Bob updates his beliefs through Bayes's formula whenever possible, i.e., after almost all of Alice's decisions, where almost all is with respect to her strategy in the given period. Because our game has perfectly observable actions, there should be no confusion what "Bayes whenever possible" means.

An equilibrium outcome (e_A, e_B) is a (measurable) function $e_A : \mathcal{U} \rightarrow [0, 1]$ and a payoff $e_B \in [0, 1]$, with the interpretation that $e_A(u)$ is the expected payoff of Alice's type u , and e_B is the expected payoff of Bob. Let $E^j(\delta, \mu)$ be the set of expected equilibrium outcomes in a game where player j makes the first offer, the discount factor is equal to δ , and Bob's beliefs are equal to μ .

There is no general result about the existence of PBE in dynamic games with incomplete information and infinitely many actions (and there are well-known problems

with the existence of sequential equilibria - see Myerson and Reny (2020)). However, one can show that, under some compactness condition on \mathcal{M} , a PBE exists.⁸

We are interested in the case of frequent offers, i.e., $\delta \rightarrow 1$. Define

$$E^j(\mu) = \limsup_{\delta \rightarrow 1} E^j(\delta, \mu).$$

2.5. Menus. Because of their importance for the statement of the main result and its proof, we further discuss menus.

The space of all menus is compact under Hausdorff topology. For each menu $Y \subseteq X$, define

$$y(u; Y) = \max_{x \in Y} u(x) \text{ and } \pi(u; Y) = \max_{x \in \arg \max_{x \in Y} u(x)} v(x).$$

Function $y(\cdot; Y)$ describes the payoffs of each type of Alice optimally choosing from menu Y , and $\pi(u; Y)$ is Bob's best possible payoff among Alice's optimal choices. To simplify the subsequent analysis, we assume that whenever a menu is offered, and Alice is indifferent between two optimal allocations in the menu, she will choose the one that is preferred by Bob.⁹ Given this assumption, Bob's expected payoff from a menu Y is equal to

$$\Pi(\mu; Y) = \int \pi(u; Y) d\mu(u).$$

We have the following important observation about menus and equilibrium payoffs.

⁸To state the compactness conditions, let $E(m, \mu)$ be the set of all equilibrium outcomes in mechanism m (i.e., in the continuation game, where the mechanism is proposed and accepted). Because equilibrium payoffs uniformly are equicontinuous in types, $E(m, \mu)$ is a subset of a compact space E^* . We assume that there exists a topology on \mathcal{M} such that the correspondence $E(\cdot, \cdot) : \mathcal{M} \times \Delta\mathcal{U} \rightrightarrows E^*$ is u.h.c., convex valued and non-empty valued. In particular, we can show that the condition is satisfied if \mathcal{M} consists of simple offers, menus, and menus of menus.

⁹The assumption only affects non-generic Alice's types who have multiple optimal choices. This assumption can be easily dropped without affecting the subsequent analysis as we can always perturb menus in a way that aligns indifferent Alice's incentives with Bob's.

Lemma 1. *For any $(e_A, e_B) \in E^j(\delta, \mu)$, there is a menu Y such that $y(\cdot; Y) = e_A$ and $e_B \leq \Pi(\mu; e_A)$.*

The Lemma is a version of the revelation principle for our environment: it says that any equilibrium outcome of the bargaining game can be attained by some menu Y .¹⁰ The proof constructs this menu as the set of expected discounted allocations that each type receives in equilibrium. We refer to a menu associated with a particular equilibrium as the equilibrium menu. We also refer to (Y, μ) as an equilibrium menu-belief pair.

3. COMPLETE INFORMATION

In this section, we discuss the special case of complete information. Here, and in the next two sections, we assume that the pie has $N = 2$ parts and refer to them as chocolate and strawberry, $n = c, s$. W.l.o.g., we assume that $v_c \geq v_s$; i.e., Bob likes chocolate more than strawberry.

3.1. Complete information bargaining. If Alice's preferences u are commonly known, the argument from Rubinstein (1982) implies that our game has unique subgame perfect equilibrium (SPE) payoffs. When $\delta \rightarrow 1$, Alice's equilibrium payoff converges to the axiomatic Nash solution to the bargaining problem in our environment.

We describe the complete information limit outcomes in more detail. When $v_c \geq v_s$, Alice's Nash payoff as a function of her preference for chocolate $u_c = 1 - u_s$ is equal to

$$\mathcal{N}_A(u_c) = \max \left(\frac{1}{2v_c} u_c, \frac{1}{2}, \frac{1}{2} (1 - u_c) \right). \quad (1)$$

¹⁰The claim holds more generally for (e_A, e_B) that are payoff outcomes in arbitrary incentive-compatible mechanisms. The representation of bargaining outcomes as an incentive-compatible mechanism goes back to Myerson (1979) and Ausubel and Deneckere (1989a).




FIGURE 1. Nash payoffs and allocations.

The left panel of Figure 1 illustrates Alice's Nash payoffs as a function of her type. The corresponding allocations are drawn on the right panel, which depicts the Edgeworth box of all non-wasteful allocations. Alice's utility increases towards the north-east. For example, under allocation s in the top-left corner, Alice gets all strawberry and Bob gets all chocolate. There are four distinct cases:

- If $u_c = v_c$, i.e. Alice's preferences are the same as Bob's, then the Nash solution awards a payoff of $\frac{1}{2}$ to each player; any allocation on Bob's payoff- $\frac{1}{2}$ indifference line (the dashed line between allocations $p = (p_{A,c}, p_{A,s}) = (\frac{1}{2v_c}, 0)$ ¹¹ and $q = (1 - \frac{1}{2v_c}, 1)$) is a solution to the Nash bargaining problem.
- If $u_c > v_c$, i.e. Alice likes chocolate more than Bob does, she is going to get her favorite allocation subject to the constraint that Bob's payoff is at least $\frac{1}{2}$, i.e., allocation p .¹² In such a case, Bob's payoff is $\frac{1}{2}$ and Alice's payoff is strictly larger than $\frac{1}{2}$.

¹¹Here, and in the rest of the paper, non-wasteful allocations are described by Alice's shares.

¹²Notice that the Nash allocation must be efficient, which means that, if $u_c \geq v_c$, the Nash allocation lies along the bottom-right part of the Edgeworth box. Then, $\Pi(y) =$

- If $\frac{1}{2} \leq u_c \leq v_c$, i.e. Alice prefers chocolate to strawberry, but the intensity of her preference for chocolate is less than Bob's, Bob receives his favorite allocation subject to the constraint that Alice's payoff is at least $\frac{1}{2}$, i.e., allocation $r = \left(1 - \frac{1}{2u_c}, 1\right)$. The allocation and Bob's payoff depends on Alice's preference; Alice's payoff is $\frac{1}{2}$.
- Finally, if $u_c \leq \frac{1}{2}$, i.e. Alice prefers strawberry to chocolate, each player receives his or her favorite part of the pie (allocation s).

3.2. Nash menu. An attempt to implement the Nash solution under incomplete information about Alice's payoffs must address Alice's incentive problem. If she likes strawberry more than Bob does, her Nash allocation is either s or one of the allocations in the interval $r \in R = \left\{(r_c, 1) : 0 \leq r_c \leq 1 - \frac{1}{2v_c}\right\}$. However, she is better off if Bob thinks that her preferences are as close to his as possible, in which case she gets an allocation close to q .

It is possible to ensure each type of Alice gets exactly her Nash payoff in an incentive-compatible manner by offering, what we call, the Nash menu:

$$\mathcal{N} = \{p, h, s\}, \quad (2)$$

where $h = \left(\frac{1}{2}, \frac{1}{2}\right)$ is the allocation of splitting each part of the pie equally, as illustrated on the left panel of Figure 2. Although the Nash menu ensures complete information payoffs for Alice, it does not implement the complete information allocations, and it does not ensure Nash payoffs for Bob. For instance, if Alice likes chocolate more than strawberry, but she likes chocolate less than Bob does, she will choose the ex-post

$\overline{\begin{cases} v_s + v_c \left(1 - \frac{y}{u_c}\right) & \text{if } y \leq u_c, \\ v_s \left(\frac{1}{u_s} - \frac{y}{u_s}\right) & \text{otherwise.} \end{cases}}$ A simple calculation shows that $y\Pi(y)$ is maximized by $y = \frac{u_c}{2v_c}$, which corresponds to allocation p . We omit the details of finding Nash allocations in other cases.

inefficient allocation h , with a payoff of $\frac{1}{2}$ for her and for Bob. However, given such a preference type, Bob's Nash payoff is strictly above $\frac{1}{2}$.

4. MAIN RESULT

4.1. Unique limit equilibrium payoffs. Notice that if menu Y contains the Nash menu, $Y \supseteq \mathcal{N}$, then the Y -payoff for each of Alice's types is not smaller than her Nash, i.e., complete information payoff:

$$y(\cdot; Y) = \max_{x \in Y} u(x) \geq \max_{x \in \mathcal{N}} u(x) = y(\cdot; \mathcal{N}) = \mathcal{N}_{\mathcal{A}}(\cdot).$$

We are ready to state the main result of this paper.

Theorem 1. *Suppose that $N = 2$ and that \mathcal{M} contains all menus and menus of menus.*

Then,

$$E^j(\mu) \subseteq \left\{ (y(\cdot; Y^*), \Pi(\mu; Y^*)) : Y^* \in \arg \max_{Y \supseteq \mathcal{N}} \Pi(\mu; Y) \right\}.$$

Bob's limit equilibrium payoff

$$\max_{Y \supseteq \mathcal{N}} \Pi(\mu; Y) \tag{3}$$

is convex and continuous (in the weak sense) in μ .*

As offers become increasingly frequent, Bob's equilibrium payoff converges to a unique value: the expected payoff from his optimal screening menu among all menus that ensure each of Alice's types receive her complete information payoff. The same payoff would be obtained if Bob were able to commit to an optimal mechanism subject to the complete information constraint. If the optimal screening menu is unique, the payoff of each of Alice's types would be also unique.

Bob's optimal payoff is continuous in his beliefs. In particular, when $\mu \rightarrow \delta_u$ for some Alice type u , Bob's payoff converges to $\mathcal{N}_B(u)$, his Nash payoff against type u . This

stands in contrast to the Coase conjecture literature, where the durable monopolist payoff in the limit $\delta \rightarrow 1$ typically depends on the support of its beliefs, and may change discontinuously with beliefs.

Additionally, the optimal payoff is convex in μ . An implication is that it has a natural comparative statics with respect to information: Bob is better off if his information improves in the sense of Blackwell's ordering. In particular, Bob is worse off due to his incomplete information about Alice's preferences. Each of Alice's types is either the same or better off under incomplete information.

The result is silent about equilibrium behavior. In fact, we are unable to construct any equilibrium in this game. However, the equilibrium behavior cannot be too different from if Bob offered the optimal screening menu, and Alice accepted it. In particular, the mechanisms must be accepted without too much delay, and, because the optimal screening menu depends significantly on Bob's beliefs, there cannot be any substantial revelation of information prior to that moment.

The proof requires that all menus and menus of menus are available. Section 6.1 argues that the thesis of Theorem 1 fails if \mathcal{M} contains only simple offers. We do not know if the theorem also holds if \mathcal{M} contains only menus but no menus of menus. However, in such a case, our proof shows that the optimal payoff (3) is a lower bound on Bob's equilibrium payoffs.

4.2. Optimal menu. We can provide additional information about the shape of the optimal menu (3).

Theorem 2. *For each Bob's belief μ , there exists allocation $x_\mu \in X$ such that*

$$\mathcal{N} \cup \{x_\mu\} \in \arg \max_{Y \supseteq \mathcal{N}} \Pi(\mu; Y).$$




FIGURE 2. Left panel: an example of optimal menu. Right panel: Anti-Coasian equilibrium from Section 6.1.

The theorem says the choice of the optimal menu can be restricted to four-element menus that contain the three allocations p, h, s of the Nash menu and a single additional allocation x_μ . An example of (the completion of) an optimal menu is illustrated on the left side of Figure 2.

Notice that the optimal screening menu is ex-ante efficient, but it typically does not ensure ex-post efficient allocations. Unless preferences of Alice are the same as preferences of Bob, all Pareto-optimal allocations lie at the edges of the Edgeworth box of non-wasteful allocations. However, in a typical optimal screening menu, some of Alice's types may choose an interior allocation x or h . The ex-post inefficiency may create a demand for renegotiation. In our main model, renegotiation is impossible because players are committed to implementing the jointly accepted mechanism. We discuss the effect of adding renegotiation in Section 6.4.

In the proof, we show that the space of all menus can be equipped with a natural convex structure such that Bob's expected payoffs $\Pi(\cdot; \mu)$ are affine in menus. Then,

we show that menus $\mathcal{N} \cup \{x\}$ are extreme points of the structure. The details are in Appendix B.

4.3. Comparison to Coasian bargaining. We contrast the results of Theorems 1 and 2 with the famous result from the Coasian bargaining literature on the durable good monopolist without commitment (Gul et al. (1986)). The solution to the latter, in the gap case, exhibits three features: (a) it is ex-post efficient, (b) the uninformed agent's payoff is as if he faces an informed player type that is worst for him, and (c) each type of the informed agent is able to mimic the behavior of the type that would maximize her payoffs.

Our setting differs from the durable-good monopolist case in multiple ways. Most prominently, it has a different procedure, as we allow both parties to make offers. The bargaining environment is slightly different as well, as we have 2-dimensional good, but no transfers (see Section 6.3 for a more direct comparison). However, it is easy to identify an outcome that shares the three aforementioned features. If we ignore the incentive problems described in Section 3.2, and allow Alice to mimic an arbitrary type, Alice shall choose her optimal allocation among all Nash allocations. A generic type will choose q if she likes chocolate less than Bob and p otherwise, or, equivalently, will choose from menu $\mathcal{C} = \{p, q\}$. For this reason, we refer to \mathcal{C} as the Coasian menu. In a companion paper (PeŃski (2019)), we show that the Coasian menu is an equilibrium outcome in a closely related model of bargaining: a war-of-attrition bargaining but with small-probability reputational types who insist on the opponent accepting their offer.

In the Coasian menu, Bob's payoff is equal to $\frac{1}{2}$ regardless of Alice's type. Theorem 2 shows that Bob's equilibrium expected payoff is typically higher, as allocation x_μ gives Bob a payoff that is strictly higher than $\frac{1}{2}$ (otherwise Bob would gain by making

x_μ equal to h and re-directing some of the types from x_μ to a profitable allocation s), and all other allocations p, h , or s lead to a payoff either equal or strictly higher than $\frac{1}{2}$.

In order to explain why the Coase conjecture fails in our paper, recall the basic logic of the Coasian bargaining literature. First, the uninformed player is not able to commit to not offering a trade to a low type in the future. He may want to postpone the transaction with the low type in order to reach a better deal with a higher type before that. Because such a deal would be unacceptable to the low type, a rejection would convince the uninformed player that he is facing the low type, rendering him more inclined to offer a trade that is acceptable to such a type in the next period. Because any offer that is acceptable to the low type is highly attractive to the high type, if the delay between offers is not costly enough, the high type to imitate the low type, reject the initial offer, which in turn destroys the equilibrium.

In our setting, notice that the “low” type is the worst possible type from Bob’s point of view, i.e., the type who likes chocolate as much as Bob himself. The last step of the above logic is not valid here.¹³ With menus, Bob is able to make an offer that is both acceptable to the “low” type and unattractive to “higher” types of Alice. In particular, because menu $\mathcal{N} \cup \{x_\mu\}$ contains allocation h , the payoff of the “low” type is the same as the payoff under Coasian menu \mathcal{C} , even if the payoffs of the higher types are lower. A rejection of such an offer does not have to be interpreted as evidence that Bob is facing the “low” type. The screening of types using menus with offers that are

¹³To be clear, because the players in our model are able to offer and accept ex-post inefficient allocations, the first step of the above logic does not hold as well. However, that does not seem to be of first importance. As we explain below in Section 6.4, giving the players the ability to renegotiate does not lower Bob’s payoff relative to what is described in Theorems 1 and 2.

acceptable to some types plays an important role in our argument (Lemma 3 below).

14

5. PROOF OF THEOREM 1

The purpose of this section is to present the proof of Theorem 1. The section focuses on the parts of the proof that do not depend on the details of the environment described in Section 2.1 and can be, possibly, generalized to other settings (see Section 6.3 for an example). We indicate the parts that are environment-specific and leave the discussion of their details to the Appendix.

5.1. Complete information bargaining. Our proof of Theorem 1 develops in parallel to the argument for the uniqueness of the equilibrium payoffs in the complete information version of the alternating-offer game. We divide this argument into few steps. Some of these steps are quite obvious, but they have more complex counterparts in the incomplete information case.

Let $\Pi(y) = \max_{x: u(x) \geq y} v(x)$ denote the best payoff available to Bob when Alice's payoff is at least y and her type is known to be u .

- (1) *Definition.* We say that payoff y is *too high* for Alice, if there exists Bob's counteroffer

$$\exists y' \text{ s.t. } y' > \delta y \text{ and } \delta \Pi(y') > \Pi(y). \quad (4)$$

- (2) *Characterization.* One checks that any payoff that is higher than the solution to equation $\delta \Pi(\delta y) = \Pi(y)$ is too high and that the latter converges to the Nash solution as $\delta \rightarrow 1$.

¹⁴A similar mechanism is at play in Board and Pycia (2014) which considers a Coasian bargaining model where the informed player has an outside option. In equilibrium, the low types prefer to exit the market, and the rejection of an offer is not meaningful in itself unless reinforced by exit.

- (3) *“Induction” bound on Bob’s payoffs.* Suppose that y is too high, and higher than the highest equilibrium payoff in the game were Alice makes the first offer. Then, the lowest equilibrium payoff of Bob must be strictly higher than $\Pi(y)$. Otherwise, Bob can deviate by rejecting Alice’s offer and, in the next period, presenting her with counter-offer y' . Because the latter would give a higher payoff than the discounted value of Alice’s highest equilibrium payoff, Alice would accept offer y' . But because $\delta\Pi(y') > \Pi(y)$, Bob’s deviation would be profitable, which contradicts the equilibrium condition.

We refer to this part of the argument as an induction bound because it can also be used to show that the SPE payoffs can be found using a dynamic version of rationalizability (Fudenberg and Tirole (2005)).

- (4) *Lower bound on Bob’s payoffs.* The “induction” bound implies that the highest equilibrium payoff cannot be too high. Together with the characterization step, we obtain that in $\delta \rightarrow 1$ limit, Bob’s equilibrium payoffs cannot be lower than his Nash payoffs.

An analogous definition for payoffs that are too low and the remaining steps complete the argument from the other side.

We divide the proof of Theorem 1 into two main parts. First, we show that Bob’s limit payoff is not smaller than (3); otherwise, he has a profitable deviation in the form of a menu. Second, we show that Alice’s payoffs cannot be lower than her complete information benchmark; otherwise, she has a profitable counteroffer in the form of a menu of menus.

5.2. Lower bound on Bob’s payoffs. There are two difficulties in extending the argument from Section 5.1: Alice’s payoffs are not totally ordered, which means that

one cannot easily find the highest equilibrium payoff, and Bob's payoffs Π depend not only on Alice's payoffs but also on his beliefs. To deal with these difficulties, we need an additional preliminary step.

Dominance relation. We are working with (not necessarily equilibrium) menu-belief pairs (Y, μ) . Define a partial order for some $\eta > 0$. We say that (Y, μ) is η -dominated by (Y', μ') if

- $\text{supp}\mu' \subseteq \text{supp}\mu$, i.e., the support of the dominating beliefs is nested in the support of the dominated beliefs and
- for each $u \in \text{supp}\mu'$, $y(u; Y') \geq y(u; Y) + \eta$, i.e. the dominating menu has higher payoffs.

The dominance relation is transitive. Moreover, because Alice's payoffs are bounded, any ordered chain of menu-belief pairs must be finite.

Using the dominance relation, we define a notion of payoffs that corresponds to “higher than the highest equilibrium payoffs” from the complete information case. We say that the menu-belief pair (Y, μ) is η -undominated by equilibrium payoffs if they are not dominated by some (Y', μ') that is an equilibrium menu-belief pair in the game in which Alice makes the first offer (see the comments after Lemma 1).

Definition. We propose the following counterpart of definition (4). Fix $\delta < 1$ and $\eta > 0$. We say that Alice's menu-belief pair (Y, μ) is (δ, η) -too high, if there exists a “deviation” menu Y_0 , such that

$$\forall u \in \text{supp}\mu \ y(u; Y_0) > \delta y(u; Y) + \eta \text{ and } \delta \Pi(\mu; Y_0) > \Pi(\mu; Y). \quad (5)$$

The first inequality says that Alice prefers to accept menu Y_0 instead of waiting for Y ; the second one says that Bob prefers to reject menu Y and wait for Y_0 .

Characterization. For any menu Y and any $\varepsilon > 0$, let $Y_\varepsilon = \bigcup_{x \in Y} \{x' : \max_{i,n} |x_{i,n} - x'_{i,n}| \leq \varepsilon\}$ be its ε -neighborhood. The next result shows that, if Alice's payoffs are strictly higher than Nash payoffs, then they are too high:

Lemma 2. *For each $\varepsilon > 0$, there exists $\delta_\varepsilon < 1$ such that, for any $\delta > \delta_\varepsilon$, there exists $\eta > 0$ such that, for any menu-belief pair (Y, μ) , if (Y, μ) η -dominates $(\mathcal{N}_\varepsilon, \mu)$, then it is (δ, η) -too high.*

In the proof, for each menu Y , we construct a deviation menu Y_0 , by appropriately modifying (stretching, shifting, or contracting) Y . This step heavily relies on the geometric details of the environment from Section 2.1 and is likely not easy to replicate in other environments. We postpone the proof until the Appendix.

“Induction” bound on Bob’s payoffs. The next Lemma contains the heart of the Rubinstein-type argument:

Lemma 3. *Fix $\delta < 1$ and $\eta > 0$. If (Y, μ) is η -undominated by equilibrium payoffs, and it is (δ, η) -too high, then $\Pi(Y, \mu) < \Pi(Y', \mu)$ for any (Y', μ) that is an equilibrium menu-belief pair in the game in which Alice makes the first offer.*

Proof. On the contrary, suppose that (Y, μ) is η -undominated by equilibrium payoffs, is (δ, η) -too high, and $\Pi(\mu; Y) > \Pi(\mu; Y')$ for some equilibrium menu-belief pair (Y', μ) . Consider an equilibrium that implements (Y', μ) . Bob has a deviation, in which he rejects any offer of Alice’s and counter-offers in the next period with Y_0 , where Y_0 is the deviation menu from the definition of too high payoffs. We claim that Alice must accept such a menu with probability 1. Indeed, if not, let (Z, ψ) be the menu-belief pair that is associated with the continuation equilibrium in the third period; it must be that ψ is absolutely continuous wrt μ , hence $\text{supp } \psi \subseteq \text{supp } \mu$. Because Y is η -undominated,

some of the rejecting types in $\text{supp}\psi$ must expect discounted continuation payoff that is not higher than $\delta y(u; Z) \leq \delta y(u, Y) + \delta\eta$, which, by the definition of the “too high” property, is strictly smaller than their payoff from menu Y_0 . Because the payoffs are continuous in types, there must be a strictly positive ψ -mass of types who have strictly higher payoffs from Y_0 , which means that rejection of Y_0 cannot happen as a best response. Because $\delta\Pi(\mu; Y_0) > \Pi(\mu; Y) > \Pi(\mu; Y')$, the deviation leads to higher payoffs than the current equilibrium, which leads to a contradiction. \square

Lower bound on Bob’s payoffs. We first show that any (Y_ε, μ) that η -dominates $(\mathcal{N}_\varepsilon, \mu)$ is η -undominated by equilibrium payoffs. On the contrary, if (Y_ε, μ) is η -dominated by some equilibrium pair, it must also be dominated by an η -undominated equilibrium pair (Y, μ') (we can always form a chain of equilibrium pairs ordered by the dominance relation, and as we observed above, such a chain must be finite). Because of transitivity, (Y, μ') η -dominates $(\mathcal{N}_\varepsilon, \mu)$, which also implies that it η -dominates $(\mathcal{N}_\varepsilon, \mu')$ due to $\text{supp}\mu' \subseteq \text{supp}\mu$. Hence, (Y, μ') is (δ, η) -too high by Lemma 2. Lemma 3 implies that any equilibrium payoff of Bob (in the game where Alice makes the first offer and Bob’s beliefs are μ') is strictly larger than $\Pi(Y, \mu')$, which contradicts (Y, μ') being an equilibrium pair.

With the above observation, Lemma 3 implies that for any $\varepsilon > 0$, if $\delta > \delta_\varepsilon$, where δ_ε is defined in Lemma 2, in any equilibrium, Bob’s payoffs cannot be lower than $\max_{Y: Y \supseteq \mathcal{N}_\varepsilon} \Pi(\mu; Y)$. In the limit $\delta \rightarrow 1$, we can take $\varepsilon \rightarrow 0$, and the payoff bound converges to (3).

5.3. Lower bound on Alice’s payoffs. Next, we show that Alice’s equilibrium payoffs cannot be lower than her Nash payoffs; otherwise, Alice has a profitable deviation.

In comparison to the previous case, there are two additional difficulties: (a) a potential profitable deviation for Alice may depend on her type, and (b) as a response to the deviation, Bob may change beliefs in an arbitrary way, which may affect Alice's continuation payoff.

Definition. We focus on a particular type u . We say that payoff y is (δ, η) -too low for type u if, for any menu Y such that $y(u; Y) \geq y$ and for any belief $\psi \in \Delta\mathcal{U}$, there exists menu Y_0 such that

$$y(u; Y_0) \geq \frac{1}{\delta}y + \eta \text{ and } \Pi(\psi; Y_0) > \frac{1}{\delta}\Pi(\psi; Y).$$

The first inequality says that a type- u Alice would wait one period to obtain menu Y_0 instead of y ; the second inequality implies that Bob prefers to accept Y_0 rather than wait for one period to get Y . The definition ensures the existence of an appropriate counter-offer menu Y_0 , but the counter-offer may depend on Bob beliefs.

“Induction” bound. Fix type u , and let

$$y_u^* = \inf \left\{ e_A(u) : (e_A, e_B) \in E^B(\delta, \mu; \mathcal{M}), \mu \in \Delta\mathcal{U} \right\}$$

be the lowest payoff of Alice's type u across all of possible equilibria and all of initial Bob beliefs in the game where Bob makes the first offer.

Lemma 4. *For each u , y_u^* is not (δ, η) -too low for type u .*

Proof. On the contrary, suppose that y_u^* is (δ, η) -too low for type u . Because y_u^* is an infimum over equilibrium payoffs, we can find an equilibrium with payoffs (e_A, e_B) such that $e_A(u) < y_u^* + \delta\eta$. Consider a deviation in which Alice rejects any offer from Bob

in the first period, and in the subsequent period proposes a menu of menus:

$$W = \left\{ Y_0 \subseteq X : y(u; Y_0) \geq \frac{1}{\delta} y_u^* + \eta \right\}. \quad (6)$$

If mechanism W is accepted, Bob is free to design any menu Y_0 , as long as type u can always find an allocation in that menu that gives her at least $\frac{1}{\delta} y_u^* + \eta$. We claim that Bob strictly prefers to accept such an offer. Indeed, if his beliefs following Alice's rejection and counteroffer W are ψ and the continuation equilibrium payoffs after Bob rejects W are attained by some menu-belief pair (Y, ψ) , then, because $y(u, Y) \geq y_u^*$, the too-low definition implies that there exists menu $Y_0 \in W$ such that Bob's expected payoff from accepting it is higher than the discounted continuation payoff, $\Pi(\psi; Y_0) > \frac{1}{\delta} \Pi(\psi; Y)$. Anticipating that Bob accepts the menu of menus, Alice's deviation is profitable as her discounted continuation payoff $y_u^* + \delta\eta$ is greater than $e_A(u)$. \square

The menu of menus W protects Alice from adversarial choice of Bob's beliefs: regardless of his beliefs, he can find a menu $Y_0 \in W$ that is more attractive for him than the continuation payoff.

Characterization. In the reminder of the argument, we show that if some Alice type has payoffs below her complete information payoffs, then there is a type (possibly, a different one) whose payoffs are too low. As in the previous part of the proof, this step heavily relies on the geometric details of the environment from Section 2.1 and, consequently, may not be easy to replicate elsewhere. We postpone the proof till the Appendix. The key step is contained in the following result.

Lemma 5. *For each $\varepsilon > 0$, there exists $\delta_\varepsilon < 1$ such that, for any $\delta > \delta_\varepsilon$, there exists $\eta > 0$ such that*

- if $u_c \in \{0, 1\}$ (i.e., Alice either only likes strawberry or only likes chocolate), then any $y \leq \mathcal{N}_A(u) - \varepsilon$ is (δ, η) -too low for u ;
- if $u_c \in [0, 1]$, then any $y \leq \frac{1}{2} - \varepsilon$ is (δ, η) -too low for u .

The lemma says that (a) payoffs below the complete information payoffs are too low for extreme types, and (b) payoffs below $\frac{1}{2}$ are too low for any type. The idea of the proof is to construct menus Y_0 by modifying (shifting or contracting) menu Y . We postpone the details to the Appendix.

Lower bound on Alice's payoffs. Together with the above argument, Lemma 5 shows that the limit equilibrium payoffs of the two extreme types must be higher than their Nash payoffs and that, in the limit, each type should get at least payoff $\frac{1}{2}$. The geometry of the problem implies that these bounds are sufficient:

Lemma 6. *Suppose that $Y \subseteq X$ is a menu such that (a) $y(u; Y) \geq \mathcal{N}_A(u)$ for extreme types u s.t. $u_c \in \{0, 1\}$, and (b) $y(u; Y) \geq \frac{1}{2}$ for all types. Then, $y(u; Y) \geq \mathcal{N}_A(u)$ for all types.*

Together with the discussion above, Lemma 6 implies that, in the limit, each Alice type must receive at least her complete information payoff.

5.4. Proof of Theorem 1. The first claim follows from the two parts. While Bob's limit payoff cannot be smaller than (3), it cannot be higher either, as Alice types must receive at least the same payoffs as from the Nash menu \mathcal{N} . Thus, Bob's limit payoff must be equal to (3).

The continuity of Bob's limit payoff is a standard consequence of the continuity of the objective function and compactness of the domain. The convexity is a consequence of the fact that $\Pi(\mu; Y)$ is linear in μ for each Y : Let $\mu_0, \mu_1 \in \Delta\mathcal{U}$ be two probability

distributions and let $\alpha \in (0, 1)$. Let $Y_\alpha \in \arg \max_{Y \supseteq \mathcal{N}} \Pi(\alpha\mu_1 + (1 - \alpha)\mu_0; Y)$. Then, because $Y_\alpha \supseteq \mathcal{N}$,

$$\begin{aligned} \max_{Y \supseteq \mathcal{N}} \Pi(\alpha\mu_1 + (1 - \alpha)\mu_0; Y) &= \Pi(\alpha\mu_1 + (1 - \alpha)\mu_0; Y_\alpha) = \alpha\Pi(\mu_1; Y_\alpha) + (1 - \alpha)\Pi(\mu_0; Y_\alpha) \\ &\leq \alpha \max_{Y \supseteq \mathcal{N}} \Pi(\mu_1; Y) + (1 - \alpha) \max_{Y \supseteq \mathcal{N}} \Pi(\mu_0; Y). \end{aligned}$$

6. COMMENTS¹⁵


6.1. Simple offers. We consider a special case of our model wherein players are only allowed to make simple offers. Let \mathcal{X} be the collection of all single-offer mechanisms, i.e., mechanisms in which players do not choose any action and some allocation $x \in X$ is implemented.

Proposition 1. *Suppose that $v_c > v_s$ and $\mathcal{M} = \mathcal{X}$, i.e., only single offers are available. Fix $u^* \in \mathcal{U}$ s.t. $u_c^* < v_c$. There exists δ_0 such that, for each $\delta \geq \delta_0$, and any belief μ s.t. $u_c^* = \inf \{u_c : u_c \in \text{supp } \mu\}$, there is $(e_A, e_B) \in E^B(\delta, \mu_0)$ such that*

$$e_B \geq \delta \mathcal{N}_B(u^*) \text{ and for each } u \text{ } e_A(u) \leq \max_{x: v(x) \geq \delta \mathcal{N}_B(u^*)} u(x).$$

Let type u^* be the type with the strongest preference for strawberry in the support of Bob's beliefs. If players are sufficiently patient, then there exists a limit equilibrium in which Bob receives his complete information payoff $\mathcal{N}_B(u^*)$ as if facing type u^* for sure. This is his best complete information payoff across Alice's types. On the other hand, each Alice type receives her best payoff subject to the constraint that Bob's payoff is at least $\mathcal{N}_B(u^*)$. The right panel of Figure 2 illustrates the result for type u^* who prefers strawberry to chocolate. The Nash allocation of type u^* gives her the strawberry part of the pie; the chocolate goes to Bob (allocation x). The blue line is

¹⁵All the proofs of the results from this section are in the Appendix.

FIGURE 3. Nash payoffs when $N = 3$.

Bob's indifference curve. Generically, Alice types choose one of the two allocations x or y .

The proof constructs an equilibrium, in which Alice offers either the allocations x or y (or anything in between). If she deviates, she is punished with a belief that she is type u^* . From now on, Bob expects nothing less than allocation x .

The equilibrium has an Anti-Coasian flavor, as Bob receives his best possible complete information payoff across all of Alice's types. With all Alice types weakly and some strictly worse off than under complete information, Alice would therefore benefit from being able to credibly reveal her type.

6.2. Case $N > 2$. We demonstrate with an example that the thesis of Theorem 1 does not hold when $N = 3$, or when the pie has a third part, vanilla. Let $v = (v_s, v_c, v_v) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Figure 3 presents the Nash payoffs for all types who only care about the first two dimensions. Let $\tau_1 = (\frac{2}{3}, \frac{1}{3}, 0)$, $\tau = (\frac{1}{2}, \frac{1}{2}, 0)$, $\tau_2 = (\frac{1}{3}, \frac{2}{3}, 0)$ be three distinctive Alice types. We have $\mathcal{N}_A(\tau_1) = \mathcal{N}_A(\tau_2) = \frac{2}{3}$ and $\mathcal{N}_A(\tau) = \frac{3}{4}$. Notice that the Nash payoffs are not convex.¹⁶

Proposition 2. *For any belief $\Delta\{\tau_1, \tau_2, \tau\}$ that assigns a strictly positive probability to types τ_1 and τ_2 , there exists $(e_A^{j,\delta}, e_B^{j,\delta}) \in E^j(\delta, \mu)$ such that $e_A^{A,\delta}(\tau) \rightarrow \frac{2}{3}$. In particular,*

¹⁶It follows that, with $N \geq 3$, one cannot implement Alice's Nash payoffs in an incentive compatible way (see Section 3.2).

for sufficiently high δ , type τ receives a payoff substantially lower than her Nash payoff of $\frac{3}{4}$.

We construct an equilibrium, in which player j always offers menu Y^j , and the offer is accepted. The menus are chosen in such a way that the optimal allocation $x_l^j \in Y^j$ of type τ_l is the same as the allocation chosen in complete information bargaining between Alice type τ_l and Bob where player j makes the first offer. Type τ chooses the best allocation for her subject to respecting incentive constraints for other types. Any offer from Alice of a mechanism that has an equilibrium with a higher payoff for type τ is rejected and punished with beliefs that Alice has a type in set $\{\tau_1, \tau_2\}$; the exact belief is chosen so that Bob prefers to reject Alice's offer and wait for the continuation game.

Although we are unable to fully characterize the set of payoffs when $N \geq 3$, we have an informative payoff bound.

Theorem 3. *Suppose that $N \geq 2$ and \mathcal{M} contains all menus and menus of menus. Then, for any $j = A, B$, any belief $\mu \in \Delta U$, any limit payoff $(e_A, e_B) \in E^j(\mu, \mathcal{M})$, any type $u \in U$, we have $e_A(u) \geq \frac{1}{2}$ and $e_B \geq \frac{1}{2}$.*

At any limit of equilibria, Bob and each Alice type receive larger payoff than their worst possible complete information payoff (the worst possible across all Bob's preferences in the case of Alice types).

6.3. Buyer-seller environment. We illustrate the applicability of methods developed in Section 5 in another environment. Suppose that a single buyer (Alice) wants a single unit of good supplied by the single seller (Bob). The seller decides on the quality of the good. The cost of quality to the seller is known, but the buyer's preference for the quality is unknown. An allocation is a pair $(q, t) \in X = \{(q, t) : q \in [0, 1], t \in \mathbb{R}\}$, where q is the quality, and t is a transfer. Alice's utility is $uq - t$, where $u \in [u_{\min}, u_{\max}]$

is a privately known value, and Bob's utility is $t - cq$, where $c > 0$ is (known) cost of providing quality. For simplicity, we assume the “no-gap” case $c < u_{\min}$. The rest of the description of the model from Section 2 and the developed notation and definitions remain the same.

Because of transferable utility, Nash bargaining splits the surplus equally:

$$\mathcal{N}_A^{BS}(u) = \frac{1}{2}(u - c).$$

Such payoffs can be implemented by single-element menu $\mathcal{N}^{BS} = \left\{\left(\frac{1}{2}, \frac{c}{2}\right)\right\}$.

To formally state the main result, we need some notation: for each of Bob's belief μ , define

$$H(\mu) = \max_{u_0} (u_0 - c)(1 - \mu(u : u \geq u_0))$$

and let $u_0(\mu)$ be the solution to the above maximization problem. It is well-known that $H(\mu)$ is the value of a price-discriminating monopolist, and $u_0(\mu)$ is the optimal price. For each $\alpha \in [0, 1]$, let

$$Y_{\alpha, \mu} = \{(\alpha, \alpha c), (1, (1 - \alpha)u_0(\mu) + \alpha c)\}$$

be a menu wherein Alice can choose to receive the good either (a) with quality α and pay its cost, or (b) with a maximum quality and the price chosen so that type $u_0(\mu)$ is indifferent between the two options. It is easy to see that, for each belief μ and each α ,

$$\Pi(\mu; Y_{\alpha, \mu}) = (1 - \alpha)H(\mu). \tag{7}$$

Bulow and Roberts (1989) shows that $Y_{0, \mu}$ is the optimal mechanism in this environment.

Theorem 4. *Suppose that \mathcal{M} contains all menus and menus of menus. Then,*

$$E^{BS,j}(\mu) \subseteq \left\{ (y(\cdot; Y^*), \Pi(\mu; Y^*)) : Y^* \in \arg \max_{Y \supseteq \mathcal{N}^{BE}} \Pi(\mu; Y) \right\}.$$

Bob's limit equilibrium payoff is equal to

$$\max_{Y \supseteq \mathcal{N}} \Pi(\mu; Y) = \Pi(\mu; Y_{1/2, \mu}) = \frac{1}{2} H(\mu). \quad (8)$$

The first part of Theorem 4 is identical to the statement of Theorem 1. The second part characterizes the optimal screening menu: in the unique limit of PBEs of the bargaining game, Bob receives a half of his monopoly payoff, and Alice pays $\frac{1}{2}c$ for a half-quality good, and, if she is willing, she pays $\frac{1}{2}u_0(\mu)$ for the maximum quality option..

The proof of Theorem 4 is contained in Appendix F. The argument relies on methods developed in Section 5, including the definitions of too-high and too-low payoffs as well as the “induction” step stated in Lemmas 3 and 4. Due to a different geometry of the allocation space, the characterization in Lemmas 2 and 5 does not carry over to the present environment and must be amended. In particular, we show that menus $Y_{\alpha, \mu}$ are too high if $\alpha > \frac{1}{2}$ and payoffs $\alpha(u - c)$ are too low if $\alpha < \frac{1}{2}$.

6.4. Renegotiation. So far, we have assumed that allocations determined by an accepted mechanism are final and cannot be renegotiated. At first glance, it may seem that the ability to commit jointly is responsible for Bob's high constrained-commitment payoff, and that allowing for renegotiation might introduce forces that would reduce Bob's payoff to, say, the Coasian outcome \mathcal{C} .

In order to examine the effect of renegotiation, we consider the following modification of the basic model. Suppose that after a mechanism is implemented and an allocation

is chosen, one of the players can request renegotiation and the other player either accepts or rejects. (Who requests the renegotiation is not relevant, but the decision to renegotiate must be made jointly.) If the renegotiation request is rejected, the game ends, and the original allocation prevails. If the request is accepted, the previous agreement is forgotten, and the players restart the bargaining game (with a possibility for future renegotiation(s)) in the next period.

We claim that Bob's equilibrium payoff under renegotiation cannot be lower than (3), i.e., it cannot be lower than the equilibrium payoff without renegotiation. The argument described in Section 5.2 remains valid under the following modification of the proof of Lemma 3 ("Induction" step).

Proof. Suppose that (Y, μ) is η -undominated by equilibrium payoffs, it is (δ, η) -too high, and $\Pi(\mu; Y) > \Pi(\mu; Y')$ for some equilibrium menu-belief pair (Y', μ) . Let Y_0 be Bob's deviation menu from the definition of the (δ, μ, η) -too high property. A potential complication due to renegotiation is that, if Bob offers a menu Y_0 and it is accepted, the payoffs of the agents depend not only on the payoffs in Y_0 , but possibly also on the continuation game in which renegotiation occurs. In particular, Alice may choose a sub-optimal allocation because she anticipates it to be renegotiated. However, we claim that the problem is not relevant here, and, with probability 1, Alice accepts Y_0 , chooses an allocation optimal for her type, and refuses renegotiation (if requested). On the contrary, suppose that Alice accepts the menu, one of the players requests renegotiation, which is, in turn, accepted. Let ψ be Bob's belief following Alice's request or acceptance of the renegotiation, and let Z be the menu that is associated with the continuation equilibrium; it must be that ψ is absolutely continuous wrt μ and hence $\text{supp}\psi \subseteq \text{supp}\mu$. Because Y is undominated, some of the rejecting types in $\text{supp}\psi$ must expect a discounted continuation payoff that is no higher than $\delta y(u; Z) \leq \delta y(u, Y) + \delta\eta$,

which, by definition of the “too high” property, is strictly smaller than their payoff from menu Y_0 . Because the payoffs are continuous in types, there must be a strictly positive ψ -mass of types who have strictly higher payoffs from Y_0 , which means that the acceptance of a renegotiation offer of Y_0 (following a possibly sub-optimal behavior in Y_0) cannot happen as a best response. Hence, Y_0 is accepted, Alice behaves as if it is final, she chooses optimally, and the outcome is not renegotiated. Because Bob prefers to wait for menu Y_0 than to accept Y immediately (the second inequality in (5)), the deviation is profitable and Y cannot be an equilibrium menu. The rest of the argument from Section 5.2 remains unchanged. \square

Although we do not know the upper bound on Bob’s payoff, the argument in Section 5.3 is not valid under renegotiation due to the problem described above. In particular, if Bob accepts Alice’s counter-offer, Alice’s behavior in the menu of menus may be sub-optimal, and lead to subsequent renegotiation. If the payoff from the continuation game is sufficiently low, Bob will reject Alice’s counter-offer in equilibrium, which may lead Alice to accept Bob’s offer in the previous period.

The fact that a reduction in commitment abilities does not reduce the uninformed party’s bargaining power is surprising. At the same time, we note that there are alternative ways of modeling renegotiation, under which a more Coasian-type result may be obtained. For example, as in Strulovici (2017), one can assume that an allocation chosen in an accepted mechanism becomes a status quo for future bargaining proposals.¹⁷ We leave these investigations for future research.

¹⁷Strulovici’s model applies to situations where players derive utility from an allocation while they renegotiate.

REFERENCES

- ABREU, D. AND F. GUL (2000): “Bargaining and Reputation,” *Econometrica*, 68, 85–117.
- AUSUBEL, L. M., P. CRAMTON, AND R. J. DENECKERE (2002): “Bargaining with incomplete information,” *Handbook of game theory*, 3, 1897–1945, publisher: Elsevier Amsterdam.
- AUSUBEL, L. M. AND R. J. DENECKERE (1989a): “A direct mechanism characterization of sequential bargaining with one-sided incomplete information,” *Journal of Economic Theory*, 48, 18–46, publisher: Elsevier.
- (1989b): “Reputation in bargaining and durable goods monopoly,” *Econometrica: Journal of the Econometric Society*, 511–531, publisher: JSTOR.
- BOARD, S. AND M. PYCIA (2014): “Outside options and the failure of the Coase conjecture,” *American Economic Review*, 104, 656–71.
- BULOW, J. AND J. ROBERTS (1989): “The simple economics of optimal auctions,” *Journal of political economy*, 97, 1060–1090.
- DOVAL, L. AND V. SKRETA (2018): “Mechanism Design with Limited Commitment,” SSRN Scholarly Paper ID 3281132, Social Science Research Network, Rochester, NY.
- FUDENBERG, D. AND J. TIROLE (2005): *Game Theory*, New Delhi: ANE Books, 1st edition ed.
- GROSSMAN, S. J. AND M. PERRY (1986): “Perfect sequential equilibrium,” *Journal of Economic Theory*, 39, 97–119.
- GUL, F., H. SONNENSCHN, AND R. WILSON (1986): “Foundations of dynamic monopoly and the coase conjecture,” *Journal of Economic Theory*, 39, 155–190.
- HARSANYI, J. C. AND R. SELTEN (1972): “A Generalized Nash Solution for Two-Person Bargaining Games with Incomplete Information,” *Management Science*, 18,

80–106.

INDERST, R. (2003): “Alternating-offer bargaining over menus under incomplete information,” *Economic Theory*, 22, 419–429.

JACKSON, M. O., H. SONNENSCHN, Y. XING, C. TOMBAZOS, AND O. AL-UBAYDLI (2018): “Negotiation as the Art of the Deal,” SSRN Scholarly Paper ID 3153853, Social Science Research Network, Rochester, NY.

LIU, Q., K. MIERENDORFF, X. SHI, AND W. ZHONG (2019): “Auctions with limited commitment,” *American Economic Review*, 109, 876–910.

MASKIN, E. AND J. TIROLE (1990): “The principal-agent relationship with an informed principal: The case of private values,” *Econometrica: Journal of the Econometric Society*, 379–409.

MOODIE, J. (2019): “No talks planned as teachers ratchet up the pressure in Sudbury and Ontario,” *Sudbury Star*.

MYERSON, R. B. (1979): “Incentive compatibility and the bargaining problem,” *Econometrica: journal of the Econometric Society*, 61–73, publisher: JSTOR.

——— (1983): “Mechanism design by an informed principal,” *Econometrica: Journal of the Econometric Society*, 1767–1797.

——— (1984): “Two-Person Bargaining Problems with Incomplete Information,” *Econometrica*, 52, 461–487.

MYERSON, R. B. AND P. J. RENY (2020): “Perfect conditional ϵ -equilibria of multi-stage games with infinite sets of signals and actions,” *Econometrica*, 88, 495–531, publisher: Wiley Online Library.

MYLOVANOV, T. AND T. TRÖGER (2012): “Informed-principal problems in environments with generalized private values,” *Theoretical Economics*, 7, 465–488.

- PESKI, M. (2019): “Bargaining over Heterogeneous Good with Structural Uncertainty,” .
- RUBINSTEIN, A. (1982): “Perfect equilibrium in a bargaining model,” *Econometrica*, 50.
- RUSHOWY, K. (2020): “Ontario government offers mediation amid faltering talks with teacher unions,” *Toronto Star*.
- SEN, A. (2000): “Multidimensional Bargaining Under Asymmetric Information,” *International Economic Review*, 41, 425–450.
- SKRETA, V. (2006): “Sequentially Optimal Mechanisms,” *Review of Economic Studies*, 73, 1085–1111.
- STRULOVICI, B. (2017): “Contract Negotiation and the Coase Conjecture: A Strategic Foundation for Renegotiation-Proof Contracts,” *Econometrica*, 85, 585–616.
- WANG, G. H. (1998): “Bargaining over a Menu of Wage Contracts,” *The Review of Economic Studies*, 65, 295–305.

APPENDIX A. PROOF OF THEOREM 1

A.1. Proof of Lemma 1. Any PBE with payoffs (e_A, e_B) induces, for each type u , a probability distribution p_u over allocation-agreement time pairs $(x, t) \in X \times \{0, 1, 2, \dots, \infty\}$. Because Alice’s preferences are linear, the expected payoff of type u is

$$\int \delta^t u(x) dp_u(t, x) = u \left(\int x_A \delta^t dp(t, x) + \mathbf{0}_A \left[p_u(\infty) + \int (1 - \delta^t) dp_u(t) \right] \right) = u(x^*(u)),$$

where $x^*(u)$ is equal to the expectation of the discounted stream of allocations. Let $\bar{x}(u)$ be the non-wasteful allocation in which Alice gets $\bar{x}_A(u) = x_A^*(u)$. Consider a menu $Y = \{\bar{x}(u) : u \in \mathcal{U}\}$. The incentive requirements of the equilibrium imply that

$y(u; Y) = e_A(u)$. For each u , we have

$$\int \delta^t v(x) dp_u(t, x) = v(x^*(u)) \leq v(\bar{x}(u)) \leq \pi(u; Y),$$

which implies that

$$e_B = \int \left(\int \delta^t v(x) dp_u(t, x) \right) d\mu(u) \leq \int \pi(u; Y) d\mu(y) \leq \Pi(\mu; Y).$$

A.2. Completed menus. One more piece of notation is going to be helpful. For each menu-belief pair (Y, μ) , define menu

$$Y_\mu = \bigcap_{u \in \text{supp}\mu} \{x : u(x) \leq y(u; Y)\}.$$

The next result summarizes the properties of

Lemma 7. *For each menu-belief pair (Y, μ) , $Y \subseteq Y_\mu$, menu Y_μ is convex, attains the same payoffs for types in the belief support, i.e., $y(u; Y) = y(u; Y_\mu)$ for each $u \in \text{supp}\mu$, and Bob's payoff under Y_μ is not lower than under Y , $\Pi(\mu; Y) \leq \Pi(\mu; Y_\mu)$.*

The inclusion comes from the fact that, because $y(u, Y) = \max_{x \in Y} u(x)$, it must be that, for each u , and each $x \in Y$, $u(x) \leq y(u, Y)$. The convexity is immediate. In order to show the equality of Alice's payoffs, given $Y \subseteq Y_\mu$, it is enough to notice that no type $u \in \text{supp}\mu$ can attain a higher payoff than $y(u; Y)$ on menu Y_μ . Finally, to verify the claim about Bob's payoffs, notice that $\arg \max_{x \in Y} u(x) \subseteq \arg \max_{x \in Y_\mu} u(x)$. Hence, for each type u in the belief support, Bob's optimal payoff among all u -optimal allocations is higher.

Menu Y_μ is the largest menu that attains the same payoffs for in-support types of Alice and preserves the incentive conditions for all types. We refer to Y_μ as the completion of menu Y .

A.3. Proof of Lemma 2. The idea of the proof is to take menu Y and to use it to construct the “deviation” menu Y_0 . The idea is illustrated in Figure 4. The elements of the constructions are quite straightforward and, as we describe below, correspond to contracting, shifting, or stretching of the original menu. The construction is complicated by the need to separately deal with the main case and two additional special cases, and because, in the main case, we need a different technique below and above the 45° diagonal.

We start with some preparation. Choose δ_ε so that, for each $\delta \geq \delta_\varepsilon$, we have

$$\delta \geq 1 - \frac{1}{4}\varepsilon^2, \quad \frac{\delta}{1+\delta} \geq \frac{1}{2} - \varepsilon, \quad \varepsilon \geq 2(1-\delta), \quad \text{and} \quad \frac{1+2\varepsilon}{1-2\varepsilon} - \frac{1}{\delta} \geq \frac{1}{2} \left(\frac{1+2\varepsilon}{1-2\varepsilon} - 1 \right).$$

Henceforth, we assume $\delta > \delta_\varepsilon$. Let

$$\eta_{\delta,\varepsilon} = (1-\delta) \min \left(\frac{2\varepsilon^2}{1-2\varepsilon}, \frac{1}{4}\varepsilon, 1 - \frac{1+\delta}{\delta} \left(\frac{1}{2} - \varepsilon \right), \frac{1}{4} \left(\frac{1+2\varepsilon}{1-2\varepsilon} - \frac{1}{\delta} \right) \varepsilon, \varepsilon \frac{1}{2\delta} - \frac{1-\delta}{\delta} \right).$$

The assumptions imply $\eta_{\delta,\varepsilon} > 0$.

Next, take any menu $Y \supseteq \mathcal{N}_\varepsilon$. Because the definition of too-high payoffs only depends on the payoffs of types in the support of μ , and $Y \subseteq Y_\mu$ by Lemma 7, we assume w.l.o.g. that $Y = Y_\mu$ (otherwise, we replace the menu by its completion). Then, Y is convex. Let

$$\alpha = \max \left\{ a : a\mathbf{1}_A + (1-a)\mathbf{0}_A \in \overline{Y} \right\}$$

be Alice’s best allocation along the 45° diagonal that belongs to completed menu Y .

We consider separately three cases:

- $\alpha < 1 - \frac{1}{2}\varepsilon$ and $v_c > v_s$,
- $\alpha < 1 - \frac{1}{2}\varepsilon$ and $v_c = v_s$, and




FIGURE 4. Left panel: Menus Y and Y_0^* . Right panel: Menu $Y \cup A$ from the proof of Theorem 2.

- $\alpha \geq 1 - \frac{1}{2}\varepsilon$.

In each of the cases, by perturbing (contracting, shifting or stretching) Y , we construct menu Y_0^* such that $y(u; Y_0^*) \geq \delta y(u; Y)$ for each u and $\Pi(\mu; Y_0) > \frac{1}{\delta}\Pi(\mu; Y) + \eta_{\delta, \varepsilon}$, which implies that Y_0^* satisfies one strict inequality of the deviation menu. To make the other inequality strict, we take $Y_0 = \left\{ \left(1 - \frac{1}{2}\eta_{\delta, \varepsilon}\right)x + \frac{1}{2}\eta_{\delta, \varepsilon}\mathbf{1}_A : x \in Y_0^* \right\}$. Then, Y_0 satisfies all the required properties of the deviation menu for $\eta = \frac{1}{2}\eta_{\delta, \varepsilon}$. It follows that Y is (δ, μ, η) -too high.

The remainder of the argument describes the construction of menu Y_0^* in the three cases above. Let $x(u) = \arg \max_{x \in x(u; Y)} v(x)$ be Bob's most preferred allocation among optimal allocations of Alice type u .

Case $\alpha < 1 - \varepsilon$ and $v_c > v_s$. In order to construct menu Y_0^* , we divide the Edgeworth box of non-wasteful allocations into two areas: above and below the 45° diagonal:

- For each u such that $x(u)$ is below the diagonal (i.e., $x_{A,c}(u) \geq x_{A,s}(u)$), define an non-wasteful allocation $x_0(u)$ so that $x_0(u) = \delta x(u) + (1 - \delta) \mathbf{0}_A$. Clearly, $x_0(u)$ is a well-defined allocation.
- For each u such that $x(u; Y)$ is above diagonal, define an non-wasteful allocation $x_0(u)$ so that $x_0(u) = x(u) + \rho(x(u) - \mathbf{1}_A)$, where constant $\rho = (1 - \delta) \frac{\alpha}{1 - \alpha}$ is chosen so that the definitions above and below agree on the diagonal. (See the left panel of Figure 4.)

We claim that $x_0(u)$ is a well-defined non-wasteful allocation. Notice that $x_{A,c}(u) \geq \varepsilon$ due to the fact menu \mathcal{N} contains all allocations that have 0 chocolate and \mathcal{N}'_ε contains all allocations that have at most ε chocolate. Also, by assumption, $\rho < (1 - \delta) \frac{1 - \frac{1}{2}\varepsilon}{\frac{1}{2}\varepsilon} < 2(1 - \delta) \frac{1}{\varepsilon}$. It follows that $x_{0,A,c}(u) \geq \varepsilon + 2(1 - \delta) \frac{1}{\varepsilon}(\varepsilon - 1) \geq \varepsilon - 2(1 - \delta) \frac{1}{\varepsilon} \geq \frac{1}{2}\varepsilon$ due to $\delta > \delta_\varepsilon$.

Let $Y_0^* = \{x_0(u) : u \in \mathcal{U}\}$. Thus, menu Y_0^* is constructed by a contraction of Y below the diagonal and expansion above the diagonal.

We show that $x_0(u)$ is the optimal allocation of Alice type u in menu Y_0^* . Indeed, take any allocation $x(u')$ for any other u' :

- If $x_0(u)$ and $x_0(u')$ are both above the diagonal, then, because $x(u)$ is the optimal allocation of type u in Y , we have

$$u(x_0(u)) = (1 + \rho)u(x(u)) - \rho \geq (1 + \rho)u(x(u')) - \rho = u(x_0(u')).$$

- If $x_0(u)$ and $x_0(u')$ are both below the diagonal, then, because $x(u)$ is the optimal allocation of type u in Y , we have

$$u(x_0(u)) = \delta u(x(u)) \geq \delta u(x(u')) = u(x_0(u')).$$

- Suppose that $x_0(u)$ is above the diagonal and $x_0(u')$ is below it. Then,

$$\begin{aligned}
u(x_0(u)) - u(x_0(u')) &= (1 + \rho)u(x(u)) - \rho - \delta u(x(u')) \\
&= \delta(u(x(u)) - u(x(u'))) + (1 + \rho - \delta)u(x(u)) - \rho \\
&\geq (1 + \rho - \delta)u(x(u)) - \rho \\
&= (1 - \delta) \left[u(x(u)) - \frac{\alpha}{1 - \alpha} (1 - u(x(u))) \right] \\
&\geq (1 - \delta) \left[\alpha - \frac{\alpha}{1 - \alpha} (1 - \alpha) \right] = 0.
\end{aligned}$$

We use $\rho = (1 - \delta) \frac{\alpha}{1 - \alpha}$; the second inequality follows from the fact that $\alpha \mathbf{1}_A + (1 - \alpha) \mathbf{0}_A \in Y$, which implies that $u(x(u)) \geq u(\alpha \mathbf{1}_A + (1 - \alpha) \mathbf{0}_A) = \alpha$.

- Finally, suppose that $x_0(u)$ is below the diagonal and $x_0(u')$ is above it. Then,

$$\begin{aligned}
u(x_0(u)) - u(x_0(u')) &= \delta u(x(u)) - (1 + \rho)u(x(u')) + \rho \\
&= \delta(u(x(u)) - u(x(u'))) + \rho - (1 + \rho - \delta)u(x(u')) \\
&\geq \rho - (1 + \rho - \delta)u(x(u')) \\
&= (1 - \delta) \left[\frac{\alpha}{1 - \alpha} (1 - u(x(u'))) - u(x(u')) \right].
\end{aligned}$$

Because $\alpha \mathbf{1} + (1 - \alpha) \mathbf{0}_A$ is a boundary point of the convex set Y that lies in-between $x(u)$ and $x(u')$, we have $u(x(u')) \leq u(\alpha \mathbf{1} + (1 - \alpha) \mathbf{0}_A) = \alpha$. Hence, the above is not smaller than

$$\geq (1 - \delta) \left[\alpha - \frac{\alpha}{1 - \alpha} (1 - \alpha) \right] = 0.$$

Next, we check that $y(u; Y_0^*) \geq \delta y(u; Y)$ for each u .

- If $x(u)$ is below the diagonal, the claim is obvious.

- If $x(u)$ is above the diagonal, we have

$$y(u; Y_0^*) - \delta y(u; Y) = (1 + \rho) y(u; Y) - \rho \delta y(u; Y) = (1 + \rho - \delta) u(x(u)) - \rho \geq 0,$$

where the inequality is established above.

Finally, we check that $\pi(u; Y_0) > \frac{1}{\delta} \pi(u; Y) + \eta_{\delta, \varepsilon}$ for each type u .

- If $x(u)$ is below the diagonal, then, because $(u) \notin \mathcal{N}'_\varepsilon$, $v(x(u)) \leq \frac{1}{2} - \varepsilon$ (see Figure 4). Thus,

$$\begin{aligned} \pi(u; Y_0) - \frac{1}{\delta} \pi(u; Y) &= v(x_0(u)) - \frac{1}{\delta} v(x(u)) = \left(\delta - \frac{1}{\delta} \right) v(x(u)) + 1 - \delta \\ &= (1 - \delta) \left(1 - \frac{1 + \delta}{\delta} v(x(u)) \right) \\ &\geq (1 - \delta) \left(1 - \frac{1 + \delta}{\delta} \left(\frac{1}{2} - \varepsilon \right) \right) \geq \eta_{\delta, \varepsilon}. \end{aligned}$$

- If $x(u)$ is above the diagonal, we have

$$\begin{aligned} \pi(u; Y_0) - \frac{1}{\delta} \pi(u; Y) &= v(x_0(u)) - \frac{1}{\delta} v(x(u)) = \left(1 + \rho - \frac{1}{\delta} \right) v(x(u)) \\ &= (1 - \delta) \left(\frac{\alpha}{1 - \alpha} - \frac{1}{\delta} \right) v(x(u)) \\ &\geq (1 - \delta) \frac{1}{4} \left(\frac{1 + 2\varepsilon}{1 - 2\varepsilon} - \frac{1}{\delta} \right) \varepsilon \geq \eta_{\delta, \varepsilon}, \end{aligned}$$

where, in the first inequality, we use $\alpha \geq \frac{1}{2} + \varepsilon$ (because $(\frac{1}{2}, \frac{1}{2}) \in \mathcal{N}$ and $\mathcal{N}_\varepsilon \subseteq Y$), and $v(x(u)) \geq v_c \frac{1}{2} \varepsilon \geq \frac{1}{4} \varepsilon$ (because $\alpha < 1 - \frac{1}{2} \varepsilon$, hence, for each allocation $x(u)$ above the diagonal, it must be that $x_{A,c}(u) \leq 1 - \frac{1}{2} \varepsilon$), in the second inequality, we use the choice of $\delta \geq \delta_\varepsilon$ and $\eta_{\delta, \varepsilon}$.

Case $\alpha < 1 - \frac{1}{2}\varepsilon$ and $v_c = v_s$. In such a case, we define $x_0(u) = \delta x(u) + (1 - \delta)\mathbf{0}_A$ and $Y_0^* = \{x_0(u) : u \in \mathcal{U}\}$. The below-diagonal part of the argument from above applies and demonstrates that menu Y_0^* satisfies all the required properties.

Case $\alpha \geq 1 - \frac{1}{2}\varepsilon$. We start with an observation that Bob's highest ex post payoff is not larger than

$$\max_u v(x(u)) \leq \max_{x \in \mathcal{N}_\varepsilon} v(x) = \max_{x \in \mathcal{N}} v(x) - \varepsilon = v_c - \varepsilon.$$

For each type u , define a non-wasteful allocation $x_0(u)$ so that

$$\begin{aligned} x_{0,A,c}(u) &= x_{A,c}(u) - \left(1 - \frac{1}{2}\varepsilon\right)(1 - \delta), \\ x_{0,A,s}(u) &= x_{A,s}(u). \end{aligned}$$

Because $x_{A,c}(u) \geq \varepsilon$, $x_0(u)$ is well-defined. Let $Y_0^* = \{x_0(u) : u \in \mathcal{U}\}$.

It is immediate that $x_0(u)$ is the optimal choice of type u in set Y_0 .

For each type u , $u(x(u)) \geq u((1 - \varepsilon)\mathbf{1}_A + \varepsilon\mathbf{0}_A) = 1 - \frac{1}{2}\varepsilon$. It follows that

$$y(u; Y_0) = u(x_0(u)) \geq u(x(u)) - \left(1 - \frac{1}{2}\varepsilon\right)(1 - \delta) \geq \delta u(x(u)) = \delta y(u; Y).$$

Finally, notice that for each type u ,

$$\begin{aligned} &\pi(u; Y_0) - \frac{1}{\delta}\pi(u; Y) \\ &= v(x_0(u)) - \frac{1}{\delta}v(x(u)) = -\frac{1 - \delta}{\delta}v(x(u)) + \left(1 - \frac{1}{2}\varepsilon\right)(1 - \delta)v_c \\ &\geq (1 - \delta) \left[\left(1 - \frac{1}{2}\varepsilon - \frac{1}{\delta}\right)v_c + \frac{1}{\delta}\varepsilon \right] \geq (1 - \delta) \left[\varepsilon \frac{1}{2\delta} - \frac{1 - \delta}{\delta} \right] \geq \eta_{\delta, \varepsilon}. \end{aligned}$$

We use the above observation in the first inequality, that $v_c \leq 1$ in the second, and the choice of δ_ε and $\eta_{\delta, \varepsilon}$ in the third.

A.4. Proof of Lemma 5. Let $\delta > \max\left(\frac{1}{1+\varepsilon}, 1 - \frac{1}{2}\varepsilon\right)$. For each $\delta > \delta_\varepsilon$, choose $\eta > 0$ so that

$$\eta \leq \min\left(\frac{1}{2}\left(\frac{1}{2}\varepsilon - (1 - \delta)\right), \frac{1 - \delta}{\delta}\varepsilon, \frac{1}{\delta}(\varepsilon - 2(1 - \delta)), \frac{1 - \delta}{\delta}\left[\varepsilon - \frac{1}{2}(1 - \delta)\right]\right)$$

We consider each case of the lemma separately. In each case, for any menu Y such that $y(u; Y) \geq y$, we construct menu Y_0 such that (b) $y(u; Y_0) \geq \frac{1}{\delta}y + \eta$, and (c) $\Pi(\psi, Y_0) > \delta\Pi(\psi, Y)$ for any belief ψ .

There are a few preliminary observations. First, we can assume that $y(u, Y) \leq \frac{1}{\delta}y + \eta$; otherwise, we can take $Y_0 = Y$. Second, notice that if Bob's expected payoff from menu Y is strictly smaller than $\frac{1}{2}(1 + \varepsilon)$, we can take Y_0 to be the approximate Coasian menu

$$Y_0 = \left\{x : v(x) \geq \frac{1}{2}(1 + \varepsilon)\right\}.$$

In such a case, Bob's payoff $\Pi(\psi, Y_0) \geq \frac{1}{2}(1 + \varepsilon) > \Pi(\psi, Y)$, and, because the completion of the Coasian menu contains the Nash menu, for each Alice type,

$$\begin{aligned} y(u; Y_0) &\geq \mathcal{N}_A(u) - \frac{1}{2}\varepsilon \geq \frac{1}{\delta}(\mathcal{N}_A(u) - \varepsilon) + \varepsilon\frac{1}{\delta} - \left(\frac{1}{\delta} - 1\right)\mathcal{N}_A(u) - \frac{1}{2}\varepsilon \\ &\geq \frac{1}{\delta}(\mathcal{N}_A(u) - \varepsilon) + \frac{1}{\delta}\left(\frac{1}{2}\varepsilon - (1 - \delta)\right) \geq \frac{1}{\delta}(\mathcal{N}_A(u) - \varepsilon) + \eta. \end{aligned}$$

Thus, we assume from now on that $\Pi(\psi, Y) \geq \frac{1}{2}(1 + \varepsilon)$.

Third, for each part of the pie $k = s, c$, each $\gamma \geq 0$, and each allocation x such that $x_{A,k} \leq 1 - \gamma$, we construct a new non-wasteful allocation $D_\gamma^k x$, in which Alice gets a γ -larger share of part k of the pie:

$$\left(D_\gamma^k x\right)_{A,k} = x_{A,k} + \gamma \text{ and } \left(D_\gamma^k x\right)_{A,-k} = x_{A,-k}.$$

Notice that $v(D_\gamma^k x) = v(x) - \gamma v_k$ and $u(D_\gamma^k x) = u(x) + \gamma u_k$ for each Alice type u .

- Suppose that $u_c = 0$. Let $y \leq \mathcal{N}_A(u) - \varepsilon = 1 - \varepsilon$. For each menu Y such that $y(u; Y) \geq y$, let

$$Y_0 = \left\{ D_{\frac{1-\delta}{1-\delta}}^s x : x \in Y \right\}$$

be a menu obtained from Y by giving Alice an extra $\frac{1-\delta}{\delta}$ -slice of strawberry. Because $x_{A,s} \leq 1 - \varepsilon$ (as the utility of type u is less than $1 - \varepsilon$) and $\frac{1-\delta}{\delta} \leq 1 - \varepsilon$, menu Y_0 is well-defined. Then, for any beliefs ψ ,

$$\begin{aligned} \Pi(\psi, Y_0) &= \Pi(\psi, Y) - \frac{1-\delta}{\delta} v_s \geq \delta \Pi(\psi, Y) + (1-\delta) \left[\Pi(\psi, Y) - \frac{1}{\delta} v_s \right] \\ &\geq \delta \Pi(\psi, Y) + (1-\delta) \left[\frac{1}{2} (1+\varepsilon) - \frac{1}{\delta} v_s \right], \end{aligned}$$

where we use $\Pi(\psi, Y) \geq \frac{1}{2} (1+\varepsilon)$. Because $v_s \leq \frac{1}{2}$ and $\frac{1}{\delta} < 1+\varepsilon$, the term in the square bracket is strictly positive and $\Pi(\psi, Y_0) \geq \delta \Pi(\psi, Y)$. Moreover,

$$\begin{aligned} y(u; Y_0) &= y(u; Y) + \frac{1-\delta}{\delta} \geq y + \frac{1-\delta}{\delta} \\ &\geq \frac{1}{\delta} y + \frac{1-\delta}{\delta} (1-y) \geq \frac{1}{\delta} y + \frac{1-\delta}{\delta} \varepsilon \geq \frac{1}{\delta} y + \eta. \end{aligned}$$

- Suppose that $u_c = 1$. Let $y \leq \mathcal{N}_A(u) - \varepsilon = \frac{1}{2v_c} - \varepsilon$. Take a menu Y such that $y(u, Y) \geq y$. Let

$$Y_0 = \left\{ D_{\frac{1-\delta}{\delta} \frac{1}{2v_c}}^c x : x \in Y \right\}$$

be a menu obtained from Y by giving Alice an extra $\frac{1-\delta}{\delta} \frac{1}{2v_c}$ -slice of chocolate. Because $x_{A,c} \leq \frac{1}{2v_c} - \varepsilon$ (as the utility of type u is less than $\frac{1}{2v_c} - \varepsilon$), menu Y_0 is

well-defined. Then, for any beliefs ψ ,

$$\begin{aligned}
\Pi(\psi, Y_0) &= \Pi(\psi, Y) - \frac{1-\delta}{\delta} \frac{1}{2v_c} v_c \\
&\geq \delta \Pi(\psi, Y) + (1-\delta) \left[\Pi(\psi, Y) - \frac{1}{\delta} \frac{1}{2v_c} v_c \right] \\
&\geq \delta \Pi(\psi, Y) + (1-\delta) \left[\Pi(\psi, Y) - \frac{1}{\delta} \frac{1}{2} \right] \\
&\geq \delta \Pi(\psi, Y),
\end{aligned}$$

where the last inequality follows from the fact that $\Pi(\psi, Y) \geq \frac{1}{2}(1+\varepsilon) \geq \frac{1}{2}\frac{1}{\delta}$.

Moreover,

$$\begin{aligned}
y(u; Y_0) &= y(u; Y) + \frac{1-\delta}{\delta} \frac{1}{2v_c} \geq \frac{1}{\delta} y + \frac{1-\delta}{\delta} \left[\frac{1}{2v_c} - y \right] \\
&\geq \frac{1}{\delta} y + \frac{1-\delta}{\delta} \varepsilon \geq \frac{1}{\delta} y + \eta.
\end{aligned}$$

- Finally, suppose that u is arbitrary and $y < \frac{1}{2} - \varepsilon$. Take a menu Y such that $y(u, Y) \geq y$, and let $Y_0 = \{\delta x + (1-\delta) \mathbf{1}_A : x \in Y\}$ be a menu obtained from Y by replacing a fraction $1-\delta$ of her current allocation by a fraction $1-\delta$ of the whole pie. Clearly, Y_0 is a well-defined menu. For any beliefs ψ , $\Pi(Y_0; \psi) = \delta \Pi(Y; \psi)$. Moreover,

$$\begin{aligned}
y(u; Y_0) &= \delta y(u; Y) + (1-\delta) \geq \frac{1}{\delta} y - \left(\frac{1}{\delta} - \delta \right) y + (1-\delta) \\
&\geq \frac{1}{\delta} y + \frac{1-\delta}{\delta} [\delta - (1+\delta)y] \geq \frac{1}{\delta} y + \frac{1-\delta}{\delta} \left[\delta - \frac{1}{2}(1+\delta) + \varepsilon \right] \\
&= \frac{1}{\delta} y + \frac{1-\delta}{\delta} \left[\varepsilon - \frac{1}{2}(1-\delta) \right] \geq \frac{1}{\delta} y + \eta.
\end{aligned}$$

A.5. Proof of Lemma 6. Take any menu Y that satisfies the two conditions (a) and (b).

- Condition (a) applied to u such that $u_c = 1$ implies that there exists an allocation $x \in Y$ such that $x_c \geq \frac{1}{2v_c}$. But then, for each u , $y(u, Y) \geq \frac{1}{2v_c}u_c$.
- Condition (a) applied to u such that $u_s = 1$ implies that there exists an allocation $x \in Y$ such that $x_s = 1$. But then, for each u , $y(u, Y) \geq u_s$.

Together with condition (b), we obtain

$$y(u; Y) \geq \max\left(\frac{1}{2v_c}u_c, \frac{1}{2}, 1 - u_c\right) = \mathcal{N}_A(u).$$

APPENDIX B. PROOF OF THEOREM 2

For each menu Y , define its “convexified” and “completed” version as $\bar{Y} = \bigcap_{u \in \mathcal{U}} \{x : u(x) \leq y(u; Y)\}$. By the proof of Lemma 7, $Y \subseteq \bar{Y}$, which implies that menu \bar{Y} leads to the same or higher payoffs for all players.

First, we show that the optimal menu must be contained in menu $\bar{\mathcal{C}} = \{x : v(x) \geq 1/2\}$ of all allocations in which Bob gets at least $\frac{1}{2}$. Hence, we can restrict ourselves to menus in the subspace $\mathcal{Y}_0 = \{Y \subseteq X, Y = \bar{Y} : \mathcal{N} \subseteq Y \subseteq \bar{\mathcal{C}}\}$. Next, we show that the space of all menus can be equipped with a natural convex structure, under which Bob’s payoffs are affine. The main step of the proof is to show that menus $\mathcal{N} \cup \{x\}$ for $x \in \bar{\mathcal{C}}$ are the only extreme points of collection \mathcal{Y}_0 .

We show that any menu $Y \supseteq \mathcal{N}$ can be replaced by a menu Y' with the same or higher payoffs for all players and such that $Y' \in \mathcal{Y}_0$. Notice first that we can replace Y by a menu \bar{Y} without reducing anybody payoffs. Let $Y' = \bar{Y} \cap \bar{\mathcal{C}}$. Because $\mathcal{N} \subseteq \bar{\mathcal{C}}$, we have $Y' \supseteq \mathcal{N}$. For each type u , either the optimal choice of Alice in menu Y belongs to menu $\bar{\mathcal{C}}$, in which case the same choice is optimal in menu Y' , or, if not, the optimal choice from menu Y leads to Bob payoff being strictly less than $\frac{1}{2}$, which is worse than

any choice from menu Y' . It follows that $\Pi(Y'; \mu) \geq \Pi(\bar{Y}; \mu) \geq \Pi(Y; \mu)$. Hence, Bob's optimal menu belongs to the class \mathcal{Y}_0 .

Define a convex structure on the space of all menus: for any Y and Y' and each $\alpha \in [0, 1]$, let

$$\alpha Y + (1 - \alpha) Y' = \bigcap_{u \in \mathcal{U}} \{x : u(x) \leq \alpha y(u; Y) + (1 - \alpha) y(u; Y')\}.$$

The convex combination is continuous with respect to the Hausdorff topology on menus. Also, for any type u , the set of optimal choices in menu $\alpha Y + (1 - \alpha) Y'$ is a convex combination of u -optimal choices in Y and Y' . It follows that for any beliefs μ ,

$$y(u; \alpha Y + (1 - \alpha) Y') = \alpha y(u; Y) + (1 - \alpha) y(u; Y'),$$

$$\Pi(\alpha Y + (1 - \alpha) Y'; \mu) = \alpha \Pi(Y, \mu) + (1 - \alpha) \Pi(Y', \mu).$$

Let $\mathcal{Y}_0^m = \{\overline{\mathcal{N} \cup A} \subseteq X : A \subseteq \bar{\mathcal{C}}, |A| \leq m\}$ be a set of the completions of menus that are unions of the Nash menu and at most m other allocations. Let $\mathcal{Y}_0^f = \bigcup_m \mathcal{Y}_0^m$. For each $m \geq 1$, \mathcal{Y}_0^m is a closed convex subset of compact space of all menus.

We have the following lemma:

Lemma 8. *For each $m \geq 2$, the set of extreme points of \mathcal{Y}_0^m is contained in \mathcal{Y}_0^1 . It follows that the set of extreme points of \mathcal{Y}_0^m is contained in \mathcal{Y}_0^1 .*

Proof. Take $m \geq 2$, and menu $Y = \overline{\mathcal{N} \cup A} \in \mathcal{Y}_0^m \setminus \mathcal{Y}_0^1$ for some $A \subseteq \bar{\mathcal{C}}, |A| = m$. We will show that Y is not an extreme point. Because $Y \in \mathcal{Y}_0^m \setminus \mathcal{Y}_0^1$, there are two preference types u^0, u^1 and allocations $a^0, a^1 \in A, a^0 \neq a^1$ such that

- $u^0(a^0) = u^0(s)$ and, for each u such that $u_c < u_c^0$, we have $\arg \max_{x \in Y} u(x) = \{s\}$;

- $u^1(a^1) = u^1(s)$ and, for each u such that $u_c > u_c^1$, we have $\arg \max_{x \in Y} u(x) \subseteq \{h, p\}$.

Define allocation a^* such that

$$u^0(a^0) = u^0(a^*) \text{ and } u^1(a^1) = u^1(a^*).$$

See the right panel of Figure 4.

For each $\alpha \in \mathbb{R}$, each $a \in A$, let $a^\alpha = \alpha a^* + (1 - \alpha)a$. Then, if $|\alpha|$ is sufficiently small, a^α is a well-defined allocation. Also, for each $x \in \mathcal{N}$, let $x^\alpha = x$. Define menu

$$Y^\alpha = \{x^\alpha : a \in Y \cup A\}.$$

Notice that

- for each $x \in \mathcal{N}$, x is the optimal choice of type u in menu Y if and only if $u_c \leq u_c^0$ or $u_c \geq u_c^1$. But then, $x = x^\alpha$ remains the optimal choice of type u in menu Y^α ;
- for each $x \in A$, x is the optimal choice of type u in menu Y if and only if $u_c^0 \leq u_c \leq u_c^1$ and x is the optimal choice in menu A . But then, x^α is the optimal choice in menu $\{x^\alpha : a \in A\}$ and (weakly) preferable to all allocations in the Nash menu.

Because payoffs $u(x^\alpha)$ are affine in α , it follows that menu Y is a convex combination of menus $Y^{-\alpha}$ and Y^α for some small but strictly positive $\alpha > 0$. \square

The Choquet Theorem implies that for each $Y \in \mathcal{Y}_0^f = \bigcup_m \mathcal{Y}_0^m$, there exists a probability measure $\mu_Y \in \Delta \mathcal{Y}_0^1$ such that, for any affine function $f : \mathcal{Y}_0 \rightarrow (\mathbb{R})$, we have

$$f(Y) = \int f(\mathcal{N} \cup \{x\}) d\mu_Y(\mathcal{N} \cup \{x\}).$$

Because \mathcal{Y}_0^f is a dense subset of a compact space \mathcal{Y}_0 , the existence of measure $\mu \in \Delta \mathcal{Y}_0^1$ with such properties extends to all menus $Y \in \mathcal{Y}_0$. The result follows from the above observation that Bob's expected payoffs are affine in Y .

APPENDIX C. PROOF OF PROPOSITION 1

Let $(x^{j,\delta})$ be the Rubinstein allocation for type u^* , i.e., the outcome of the complete information game of Bob and Alice type u where player j makes the first offer and players can make any simple offer, $\Gamma^j(\delta, \mathcal{S}, \delta_{u^*})$. Let $A = \{x : v(x) \geq \delta v(x^{B,\delta})\}$ be the set of allocations x that are acceptable to Bob, who expects allocation $x^{B,\delta}$ in the continuation game if he rejects the current offer x .

We construct an equilibrium such that

- in the Alice offer subgame,
 - Alice type u always offers her best allocation from set A , $x(u) \in \arg \max_{x \in A} u(x)$.
(Notice that Alice is indifferent between two allocations only when $u = v$, in which case Bob is also indifferent among her optimal choices.)
 - Bob accepts any offer $x \in A$.
 - If Alice offers $x \notin A$, Bob rejects, changes his beliefs to $\mu^* = \delta_{u^*}$, and expects from now on to follow the complete information equilibrium with payoff $v(x^{B,\delta})$ in the next period.
- Bob's offer subgame,
 - Bob (with beliefs μ) chooses an offer so to maximize his expected payoff

$$\int_{u: u(x) \geq \delta u(x(u))} v(u(x)) d\mu(u) + \delta v(x^{B,\delta}) \int_{u: u(x) \geq \delta u(x(u))}^2 d\mu(u),$$

subject to Alice's above strategy.

- Alice type u accepts any offer x such that $u(x) \geq \delta u(x(u))$ and rejects otherwise.

Clearly, the strategies are best responses and the equilibrium payoffs are as required.

APPENDIX D. PROOF OF PROPOSITION 2

Fix $\delta < 1$. For each player j and each $l = 1, 2$, define non-wasteful allocations x_l^j so that x_l^j is the equilibrium allocation in the complete information bargaining between Alice type τ_l and Bob in the subgame where player l makes the first offer. Let

$$A^j = \{x : \tau_l(x) \leq \tau_l(x_l^j) \text{ for each } l\}$$

be the set of allocations such that each type τ_l (weakly) prefers x_l^j to x^j . Let $x^j \in \arg \max_{x \in A} \tau(x)$ be the best non-wasteful allocation for Alice type τ among all allocations in A . It is easy to check that, for each j ,

$$\tau_1(x_1^j) = \tau(x^j) = \tau_2(x_2^j) =: u^j \text{ and } \delta u^A = u^B.$$

(Note that $\tau_l(x_l^j) = \tau_l(x^j)$ by the choice of x^j and symmetry. Moreover, $\tau(x^j) = \frac{1}{2}\tau_1(x^j) + \frac{1}{2}\tau_2(x^j)$.) The symmetry of two extreme types τ^l implies that

$$v(x_1^j) = v(x_2^j) = v^j \text{ and } v^A = \delta v^B.$$

Also, let $v_0^j = v(x^j)$. For each j , define a menu

$$Y^j = \{x_1^j, x^j, x_2^j\} \text{ for each } j.$$

Recall the definition of equilibrium correspondence E from footnote 8. Let \mathcal{M}^A be a class of mechanisms m such that, for each belief $\mu \in \Delta\{\tau_1, \tau_2\}$, and each continuation

equilibrium $(e_A, e_B) \in E(m, \mu)$, we have $e_A(t) > u^A$ for some $t \in \{\tau_1, t, \tau_2\}$. We have the following lemma:

Lemma 9. *For each $m \in \mathcal{M}^A$, there exists a belief $\mu^m \in \Delta\{\tau_1, \tau_2\}$ and continuation equilibrium payoffs $e^m \in E(m, \mu)$ such that $e_B^m < v^A$.*

Proof. Fix $m \in \mathcal{M}^A$. The choice of allocation x^A as the τ -best allocation that satisfies incentive-compatibility conditions for types τ_l implies that, for any allocation x , if $\tau(x) > \tau(x^A) = u^A$, then $\tau_l(x) > u^A$ for some l . Hence, for any belief $\mu \in \Delta\{\tau_1, \tau_2\}$ and any mechanism equilibrium payoffs $(e_A, e_B) \in E(m, \mu)$, there is $l = 1, 2$ such that $e_A(\tau_l) > u^A$.

The compactness assumption (Footnote 8) implies that

$$E(m) = \{(e, \mu) : \mu \in \Delta\{\tau_1, \tau_2\}, e \in E(m, \mu)\}$$

is a compact and connected subset of compact Polish space $E^* \times \Delta\{\tau_1, \tau_2\}$. Let

$$W_l = \{(e, \mu) \in E(m) : e(\tau_l) > u^A\}.$$

Sets W_l are open, cover $E(m)$, and hence either

- there is l_0 such that $E(m) \subseteq W_{l_0}$. In such a case, take $\mu^m = \delta(\tau_{l_0})$ as the measure assigning full probability to type τ_{l_0} , and take any $e^m \in E(m, \mu^m)$. Then, the choice of $x_{l_0}^A$ implies that $e_B^m < v^A$. Or,
- $W_1 \cap W_2 \neq \emptyset$, in which case take any $(e^m, \mu^m) \in W_1 \cap W_2$. The choice of x_l^A for each l implies that $e_B < v^A$.

□

Construct an equilibrium:

- Alice's offer subgame: Alice offers menu Y^A and each of her type expects payoff u^A .
 - If Alice offers menu Y^A , Bob accepts it. If not, his continuation payoff is $\delta v^B = v^A$. Hence, accepting is a best response, and Alice's payoff is u^A .
 - If Alice offers mechanism $m \in \mathcal{M}^A$, Bob updates his beliefs to μ^m , and he expects that, if m is accepted, the equilibrium with payoffs e^m is played, which leads to payoff $e_B^m < v^A$ for Bob. Given that, Bob rejects the mechanism and expects discounted continuation payoff $\delta v^B = v^A$. In such a case, the expected payoff of each Alice type t is no higher than $\delta u^B < u^A$.
 - If Alice offers mechanism $m \notin \mathcal{M}^A$, an arbitrary continuation equilibrium is played. By definition of class \mathcal{M}^A , each Alice type receives a payoff no higher than u^A .

Because no deviation leads to a higher payoff for any type, offer of menu Y^A is a best response.

- Bob's offer subgame: Bob offers menu Y^B and expects payoff $\mu \{\tau_1, \tau_2\} v^B + (1 - \mu \{\tau_1, \tau_2\}) v_0^B$.
 - If Bob offers menu Y^B , Alice accepts it and each of her types receives payoff u^B . If she were to reject, each of her types expects continuation payoff $\delta u^A = u^B$. Hence, accepting is a best response.
 - If Bob offers any other mechanism $m \neq Y^B$, and Alice rejects it, each of her types expects continuation payoff $\delta u^A = u^B$. If any of her types accepts it, it must be that she receives a payoff of at least u^B . In any case, whatever is the continuation equilibrium payoff in the subgame after Bob offers m , each Alice type receives at least u^B as her continuation payoff. The choice

of allocations x_1^B, x^B, x_2^B implies that Bob's payoff cannot be larger than v^B .

Hence, offering Y^B is a best response.

APPENDIX E. PROOF OF THEOREM 3

The lower bound on Alice's payoff is a consequence of the proof of the second part of Lemma 5, which does not depend on N .

We show the bound on Bob's payoff. For each Alice type u , and each $v \geq 0$, define $u^*(v) = \max_{x: v(u) \geq v} u(x)$ as the largest payoff of type u that is consistent with Bob receiving at least v . For each $v \in [0, 1]$, let $Y(v)$ be a menu $Y(v) = \{x : v(x) \geq v\}$. Let

$$e_B^{A,\delta} = \inf \left\{ e_B : (e_A, e_B) \in E^A(\delta, \mu; \mathcal{M}) \text{ for any } \mu \in \Delta\mathcal{U} \right\}$$

be Bob's lowest equilibrium payoff across all possible beliefs in the game in which Alice makes the first offer.

We show that, in the subgame where Bob makes the first offer, Alice will accept any menu that contains menu $Y(\delta e_B^{A,\delta} + 1 - \delta)$ in its interior. In any equilibrium of the game where Alice makes the first offer, Bob's expected payoff is not lower than $e_B^{A,\delta}$. Hence, it is impossible for every type u of Alice to expect more than $u^*(e_B^{A,\delta})$. It follows that, in the game where Bob makes the first offer, a positive-measure fraction of Alice's types must accept any offer that is strictly larger than $\delta u^*(e_B^{A,\delta})$. An argument similar to the one used in the "Equilibrium cannot be too high" part of Section 5.2 shows that all types u must accept any menu with payoffs $y(u) > \delta u^*(e_B^{A,\delta})$. (If some types reject, then a positive fraction of them would receive tomorrow's payoffs that are lower than $u^*(e_B^{A,\delta})$. But then, a rejection would not be a best response.) Due to

linearity of u and v ,

$$\begin{aligned}
\delta u^*(e_B^{A,\delta}) &= \delta \max_{x:v(x) \geq e_B^{A,\delta}} u(x) = \max_{x:v(x) \geq e_B^{A,\delta}} u(\delta x + (1-\delta)\mathbf{0}_A) \\
&= \max_{x:v(\delta x + (1-\delta)\mathbf{0}_A) \geq \delta e_B^{A,\delta} + 1 - \delta} u(\delta x + (1-\delta)\mathbf{0}_A) \\
&\leq \max_{x:v(x) \geq \delta e_B^{A,\delta} + 1 - \delta} u(x) = u^*(\delta e_B^{A,\delta} + 1 - \delta).
\end{aligned}$$

(The inequality comes from the fact that the set of allocations X is convex, and, specifically, for each $x \in X$, $\delta x + (1-\delta)\mathbf{0}_A \in X$.) We conclude that Alice accepts any menu that contains menu $Y(\delta e_B^{A,\delta} + 1 - \delta)$ in its interior.

We are going to show that $e_B^{A,\delta} \geq \frac{\delta}{1+\delta}$. On the contrary, suppose that $e_B^{A,\delta} < \frac{\delta}{1+\delta}$. Then, there exists an equilibrium of the game where Alice makes the first offer with Bob's expected payoffs $e_B \leq e_B^{A,\delta}$. Consider a deviation where Bob rejects any offer from Alice and, instead, proposes menu $Y(\delta e_B^{A,\delta} + 1 - \delta)$. The above paragraph implies that such a menu is accepted for sure. Bob's deviation is strictly profitable: notice that, in the previous period,

$$\begin{aligned}
\delta(\delta e_B^{A,\delta} + 1 - \delta) &= e_B^{A,\delta} - (1 - \delta^2)e_B^{A,\delta} + \delta(1 - \delta) = e_B^{A,\delta} + (1 - \delta)(\delta - (1 + \delta)e_B^{A,\delta}) \\
&> e_B^{A,\delta} + (1 - \delta)\left(\delta - (1 + \delta)\frac{\delta}{1 + \delta}\right) = e_B^{A,\delta}.
\end{aligned}$$

It follows that $\lim_{\delta \rightarrow 1} \inf e_B^{A,\delta} \geq \lim_{\delta \rightarrow 1} \frac{\delta}{1 + \delta} = \frac{1}{2}$.

APPENDIX F. PROOF OF THEOREM 4

F.1. Bob's lower payoff bound. The first part of the proof establishes a lower bound on Bob's payoffs. We are going to show that for each $\varepsilon > 0$, there is $\delta_\varepsilon < 1$ such that, for each $\delta \geq \delta_\varepsilon$ and each belief μ , Bob's equilibrium payoff is not smaller than

$(\frac{1}{2} - \varepsilon) H(\mu)$. Due to formula (7), it is enough to show that $\alpha_{\max} \leq \frac{1}{2} + \varepsilon$, where we define

$$\alpha_{\max} = \sup \{ \alpha : \Pi(\mu; Y) \geq \Pi(\mu; Y_{\alpha, \mu}) \text{ for each equilibrium pair } (Y, \mu) \}.$$

The next result is a straightforward consequence of the payoff formula (7).

Lemma 10. *For each $\varepsilon > 0$, there is $\delta_\varepsilon < 1$ such that, for each $\delta > \delta_\varepsilon$, there exists $\eta > 0$, such that for each μ , and each $\alpha > \frac{1}{2}(1 + \varepsilon)$, pair $(Y_{\alpha, \mu}, \mu)$ is (δ, η) -too high.*

Proof. Choose δ_ε so that $\frac{1}{1+\delta_\varepsilon} \leq \frac{1}{2}(1 + \frac{1}{2}\varepsilon)$ and, for each $\delta \geq \delta_\varepsilon$, choose η so that

$$(1 - \delta^2) \left(\alpha - \frac{1}{1 + \delta} \right) > (1 + (u_{\min} - c)) \eta.$$

That can be done as $\alpha - \frac{1}{1+\delta} > \frac{1}{4}\varepsilon$. Simple calculations show that if we take $\alpha_0 = \delta\alpha + \eta$, then, for each u ,

$$y(u; Y_{\alpha_0, \mu}) \geq \delta y(u; Y_\alpha) + \eta \text{ and } \delta(1 - \alpha_0)((u_{\min} - c)) \geq (1 - \alpha)((u_{\min} - c)) + \eta.$$

Due to $H(\mu) \geq u_{\min} - c$ and the payoff formula (7), the last inequality implies that $\delta\Pi(Y_{\alpha_0, \mu}, \mu) > \Pi(Y_{\alpha, \mu}, \mu) + \eta$. \square

Lemma 10 shows that each menu Y_α for which Alice's payoffs are strictly larger than her Nash payoffs for all her types (i.e., if $\alpha > \frac{1}{2}$) is too high.

Compared to Lemma 2, the range of menus identified as too-high is narrower and limited to menus $Y_{\alpha, \mu}$. The next result that shows that such menus are “optimal” in some sense:

Lemma 11. *For each menu-belief pair (Y, μ) , if, for some $\alpha \in [0, 1]$, $y(u; Y) \geq \alpha(u - c)$ for each $u \in \text{supp}\mu$, then $\Pi(\mu; Y) \leq (1 - \alpha)H(\alpha) = \Pi(\mu; Y_{\alpha, \mu})$.*

Proof. Construct a menu $Y' = \{(\alpha, \alpha c)\} \cup Y \setminus \{(q, t) : t < qc\}$. Then, for each type u , $y(u; Y') \geq \alpha(u - c)$. Moreover, notice that, for each type u , if $\pi(u; Y) > 0$, then, the optimal choice of this type in menu Y is the same as the optimal choice in menu Y' , which implies that $\pi(u; Y') = \pi(u; Y)$. Because $\pi(u; Y') \geq 0$ for each type u , we have

$$\Pi(\mu; Y') = \int \pi(u; Y') d\mu(u) \geq \int \pi(u; Y) d\mu(u) = \int \pi(u; Y) d\mu(u).$$

Hence, it is sufficient to show that $\Pi(\mu; Y') \leq (1 - \alpha) H(\mu)$.

For each $u \in \text{supp}\mu$, let (q_u, t_u) be Bob's optimal choice among all Alice type u optimal allocations in menu Y' . Then, $t_u \geq q_u c$ and $q_u u - t_u \geq \alpha(u - c)$. Because $u > c$, the two inequalities imply that $q_u \geq \alpha$. Hence,

$$\begin{aligned} \Pi(\mu; Y') &= \int (t_u - q_u c) d\mu(u) = \int (t_u - \alpha c - (q_u - \alpha) c) d\mu(u) \\ &= (1 - \alpha) \int (\tau_u - p_u c) d\mu(u), \end{aligned} \tag{9}$$

where we denote $p_u = \frac{1}{1-\alpha}(p_u - \alpha) \in [0, 1]$ and $\tau_u = \frac{1}{1-\alpha}(t_u - \alpha c)$. Note that the scheme $(p_u, \tau_u)_{u \in \text{supp}\mu}$ must satisfy standard incentive compatibility conditions: for any u, u' ,

$$p_u u - \tau_u - (p_{u'} u - \tau_{u'}) = \frac{1}{1-\alpha} (q_u u - t_u - (q_{u'} u - t_{u'})) \geq 0,$$

where the inequality follows from the incentive conditions implied by the optimal choice from menu Y' . Hence, the analysis of Bulow and Roberts (1989) implies that the value of the integral in (9) is no larger than $H(\mu)$, or that

$$\Pi(\mu; Y') \leq (1 - \alpha) H(\mu).$$

□

Armed with the two results, we can proceed with the proof of the first part of Theorem 4. Suppose that $\alpha_{\max} > \frac{1}{2} + \varepsilon$. Let $\delta \geq \delta_\varepsilon$ and $\eta > 0$ be such as in Lemma 10. Take α s.t. $\max\left(\frac{1}{2}(1 + \varepsilon), \alpha_{\max} - \frac{1}{2} \frac{\eta}{u_{\max} - c}\right) < \alpha < \alpha_{\max}$. We are going to show that each menu-belief pair $(Y_{\alpha, \mu}, \mu)$ is η -undominated by equilibrium payoffs. If not, then there is equilibrium menu-belief pair (Z, ψ) such that $\text{supp}\psi \subseteq \text{supp}\mu$ and for each $u \in \text{supp}\psi$,

$$y(u; Z) \geq y(u; Y_{\alpha, \mu}) + \eta \geq \alpha(u - c) + \eta \geq \left(\alpha + \frac{\eta}{u_{\max} - c}\right)(u - c).$$

But then, Lemma 11 implies that

$$\Pi(\psi; Z) \leq \Pi\left(\mu; Y_{\alpha + \frac{\eta}{u_{\max} - c}, \mu}\right) < \Pi(\mu; Y_{\alpha_{\max}, \mu}),$$

which contradicts the definition of α_{\max} .

Because pairs $(Y_{\alpha, \mu}, \mu)$ are η -undominated, Lemma 3 implies that a payoff in any equilibrium is not lower than $\Pi(\mu; Y_{\alpha, \mu})$. But this means that $\alpha_{\max} \leq \alpha$, which contradicts the choice of α . The contradiction shows that $\alpha_{\max} \geq \frac{1}{2} + \varepsilon$.

F.2. Alice's payoff bound. Another consequence of the payoff formula (7) is that, if $\alpha < \frac{1}{2}$, then payoffs $\alpha(u - c)$ are too low for Alice's type u :

Lemma 12. *For each $\varepsilon > 0$, there is $\delta_\varepsilon < 1$ such that, for each $\delta > \delta_\varepsilon$, there exists $\eta > 0$, such that, for each $\alpha < \frac{1}{2}(1 - \varepsilon)$, payoff $\alpha(u_{\min} - c)$ is (δ, η) -too low for type u_{\min} .*

Proof. Take any menu Y such that $y(u_{\min}; Y) \geq \alpha(u_{\min} - c)$. As in the proof of Lemma 11, we can replace Y with menu $Y' = \{(\alpha, \alpha c)\} \cup Y \setminus \{(q, t) : t < qc\}$ without reducing the payoffs of either Bob's or type u_{\min} . Let (q_u, t_u) be an optimal choice of type $u \geq u_{\min}$. Then, $q_u \geq \alpha$ (see the proof of Lemma 11). Incentive compatibility

implies that

$$q_u u - t_u \geq q_{u_{\min}} u - t_{u_{\min}} \geq q_{u_{\min}} (u - u_{\min}) + \alpha (u_{\min} - c) \geq \alpha (u - c).$$

Lemma 11 implies that $\Pi(\mu; Y) \leq \Pi(\mu; Y') \leq \Pi(\mu; Y_{\alpha, \mu})$.

Choose δ_ε so that $\frac{1+\delta_\varepsilon}{\delta_\varepsilon} \leq \frac{2}{1-\frac{1}{2}\varepsilon}$ and, for each $\delta \geq \delta_\varepsilon$, choose η so that

$$\left(1 + \frac{1}{u_{\min} - c}\right) \eta < \frac{1}{2} (1 - \delta) \varepsilon.$$

That can be done as $\alpha - \frac{1}{1+\delta} > \frac{1}{4}\varepsilon$. Take α_0 so that $1 - \alpha_0 = \delta(1 - \alpha) + \eta$. Then, due to the payoff formula (7), $\Pi(\mu; Y_{\alpha_0, \mu}) > \delta \Pi(\mu; Y_{\alpha, \mu})$. Moreover,

$$\begin{aligned} y(u_{\min}; Y_{\alpha_0, \mu}) &\geq \alpha_0 (u_{\min} - c) = (1 - \delta(1 - \alpha) - \eta) (u_{\min} - c) \\ &= \frac{1}{\delta} \alpha (u_{\min} - c) + \eta + \left(1 - \delta(1 - \alpha) - \eta - \frac{1}{\delta} \alpha - \frac{\eta}{u_{\min} - c}\right) (u_{\min} - c) \\ &\geq \frac{1}{\delta} \alpha (u_{\min} - c) + \eta, \end{aligned}$$

where the inequality follows from the fact that

$$\begin{aligned} 1 - \delta(1 - \alpha) - \eta - \frac{1}{\delta} \alpha - \frac{\eta}{u_{\min} - c} &= 1 - \delta - \left(\frac{1 - \delta^2}{\delta}\right) \alpha - \left(1 + \frac{1}{u_{\min} - c}\right) \eta \\ &= (1 - \delta) \left(1 - \frac{1 + \delta}{\delta} \alpha\right) - \left(1 + \frac{1}{u_{\min} - c}\right) \eta \\ &\geq (1 - \delta) \left(1 - \frac{1 - \varepsilon}{1 - \frac{1}{2}\varepsilon}\right) - \left(1 + \frac{1}{u_{\min} - c}\right) \eta \\ &\geq \frac{1}{2} (1 - \delta) \varepsilon - \left(1 + \frac{1}{u_{\min} - c}\right) \eta \geq 0. \end{aligned}$$

Hence, menu $Y_0 = Y_{\alpha_0, \mu}$ is a deviation menu from the definition of the too-low payoffs.

□

Together with Lemma 4, the above result implies that each Alice type u receives at least payoff $\frac{1}{2}(1 - \varepsilon)(u - c) = (1 - \varepsilon)\mathcal{N}_A^{BS}(u)$ in the equilibrium of the game in which Bob makes the first offer and with discount factor $\delta \geq \delta_\varepsilon$. In the limit $\delta \rightarrow 1$, each type u must receive at least $\mathcal{N}_A^{BS}(u)$. Together with the lower bound on Bob's payoffs, this ends the proof of the theorem.