BARGAINING OVER HETEROGENEOUS PIE WITH MECHANISMS AND INCOMPLETE INFORMATION

MARCIN PĘSKI

ABSTRACT. This paper 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 gets at least her payoff under complete information. We explain that this outcome is equivalent to a constrained commitment. With more than two dimensions, there are 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 a simple 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]

Date: March 17, 2020.

Department of Economics, University of Toronto. I am grateful for helpful comments from V Bhaskar, Laura Doval, 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.

[1]See 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

They 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." All such offers can be represented as mechanisms, i.e., general games, an 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. We address 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's 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. Bob has 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 finite (extensive-form) game with perfectly observable actions, where players' choices determine final allocations. When the offer is accepted, the mechanism is

and the representatives of faculty and staff. Among others, the parties negotiated the size of 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.

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

[3]The EU's accession negotiations typically take the form of independent bargaining over 30-40 areas.

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

Any strategic model of bargaining under incomplete information must deal with two type of problems. Due to a screening problem, a player's offer may be acceptable for 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 of the threat of being 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 without requiring equilibrium refinements. Our main result is that, when $N = 2$ and players are patient, there is a unique limit of PBE outcomes. The limit outcome is equivalent to Bob proposing a screening menu m^* that is optimal for him subject to the constraint that each of Alice's types receives at least her complete information payoff. The latter is equal to the Nash bargaining payoff. 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 has two parts: first, we show that any offer from Bob that gives Alice types their Nash payoffs will be accepted; second, we show that each Alice type can ensure the Nash payoff. 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, by offering menus (for Bob to choose) of menus (of allocations, for Alice to choose), Alice can protect herself from "punishment with beliefs".

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 without commitment. 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. The type of uncertainty considered here is very important. Since there are no best or worst types, types are simply different, and thus belief threat can be tested by a mechanism that is acceptable for both Bob with his punishment beliefs and Alice's true type.

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 about Alice 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 Rubinstein's model and the Nash solution. We explain in the paper that the "commitment ability" seems to result from Bob exploiting Alice's fear of revealing too much information about herself.

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 uniformed 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 (among other differences, such as the fact that players can only propose menus rather than arbitrary mechanisms). Interestingly, more commitment leads to a Coasian-type result: in the unique (rational and patient limit) equilibrium, Bob proposes a menu $m^{1/2}$ of all allocations that give him at least his worst possible Nash payoff (equal to $\frac{1}{2}$). Bob is typically strictly worse off than under m^*; 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).

Although most of the literature is restricted to the two-dimensional case, 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, the informed principal must benefit from incomplete information due to the collapse of the agent incentive and individual rationality constraints.[4] 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

[4] I am grateful to V Bhaskar for this observation.

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, appear also in some strategic papers on bargaining under one-sided incomplete information. With the exception of Jackson et al. (2018), all related papers that we are aware of work solely with two types. Sen (2000) 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 (the incentives of the low type go in the correct direction). Inderst (2003) studies a similar setting, but assumes that the two types have the incentives to separate to complete information payoffs. 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, 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''.

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). There are few 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. 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 with limited commitment.

2. Model

2.1. **Bargaining.** Two players, Alice and Bob, bargain over a heterogeneous pie with $N \geq 2$ parts. An allocation is defined as a tuple $x = (x_{i,n}) \in X = \left\{x \in [0,1]^{2N} : \sum_{i,n} x_{i,n} \leq 1\right\}$, where $x_{i,n}$ is player i's share of the nth part of the pie. We allow for allocations with waste, but this does not affect our results. The main result is about $N = 2$, in which case, we refer to the two parts of the pie as chocolate and strawberry, $n = c, s$. (Section 4.2 comments on $N > 2$.)

Each player i has a linear preference $u_i(x) = u_i \cdot x_i$, where $u_i \in \mathcal{U} = \left\{u \in [0,1]^N : \sum u_n = 1\right\}$ over allocations x. We normalize the preferences so that the coefficients add up to 1; this normalization is w.l.o.g. because multiplying payoffs by a constant does not alter the strategic behavior. Extreme preferences are denoted as $\omega^n \in \mathcal{U}$, where $\omega^n_m = 1_{n=m}$ for each n, m. Bob's preferences, denoted as v, are commonly known. Alice's preferences, denoted as u, are privately known by her; Bob's beliefs are denoted by $\mu \in \Delta\mathcal{U}$.

In alternating periods, one player offers to choose an allocation with a mechanism m; the other player either accepts or reject. The first offer is made by player j. If the offer is accepted, the mechanism m is implemented, the allocation is determined in a continuation equilibrium, and the game ends with players receiving payoffs from the allocation. If the offer is rejected, the game moves to the next period, with the other player making an offer. A *mechanism* is formally defined as an extensive, finite-horizon game with observable actions $m = \left((S_i^t)_{i=A,B}^{t\leq T}, \chi\right)$, where $T < \infty$, S_i^t is a finite set of actions for player i in period t, , and $\chi : \prod_i^T S_i^t \to X$ is an allocation function. Let

$$\mathcal{M}_F = \bigcup_{T<\infty, \left(S_i^t\right)_{i=A,B}^{t\leq T}:\text{finite } S_i^t \subseteq \mathbb{N}} X^{\prod_{i,t\leq T} S_i^t}$$

denote the space of all mechanisms. $\mathcal{M}_F$ is a countable union of Euclidean spaces. We do not allow for transfers, but otherwise allow for arbitrary (finite) mechanisms, including:

- *simple offers:* $T = |S_A^1| = |S_B^1| = 1$. Each simple offer can be identified with a single allocation $\chi \in X$,
- *(Alice's) menus*: $T = 1 = |S_B^1|$. Each menu m is characterized by a finite set of allocations $Y_m = \{\chi(s_A) : s_A \in S_A^1\} \in CX$, where CX is the space of closed subsets of X with Hausdorff distance.
- *(Bob's) menus of (Alice's) menus*: $T = 2, |S_A^1| = |S_B^2| = 1$. Menu m of menus are characterized by a finite set of finite sets of allocations: $W_m = \{Y(s_B) : s_B \in S_B^1\} \in C^2X$, where $Y(s_B) = \{\chi(s_A, s_B) : s_A \in S_A^2\}$. Here, Bob first chooses a menu $Y \in W$, and then Alice chooses an allocation in Y. An example is the "Bob divides, Alice chooses" mechanism.

All the results go through as long as the space of available mechanisms contains menus and menus of menus.

Because we do not have the existence of equilibrium result for infinite-action games, we choose to work with finite approximations.[5] Because $\mathcal{M}_F$ is separable (as a countable union of closed subsets of Euclidean spaces), it can be approximated with an increasing sequence of finite subsets $\mathcal{M}_k \subseteq \mathcal{M}_{k+1} \subseteq ...\mathcal{M}_F$ such that $\text{cl} \bigcup_k \mathcal{M}_k = \mathcal{M}_F$. We write $\mathcal{M}_k \to \mathcal{M}_F$.

The players discount with a common factor $\delta < 1$. We are interested in the case of frequent offers, i.e., $\delta \to 1$. For some results (that we explicitly state), we assume that the players observe the outcome of a public randomization device before taking any action. Let $\Gamma^j(\delta, \mathcal{M}_k, \mu)$ denote the bargaining game in which player j makes the first offer, Bob's initial beliefs are given by μ, and the players choose their offers from set $\mathcal{M}_k$.

2.2. **Strategies and equilibrium.** Let $T_j = \{t \in \mathbb{N} : t \text{ odd}\}$ be the periods in which initial player j makes the offer. Let $T_{-j} = \{t \in \mathbb{N} : t \text{ even}\}$. For $t \geq 1$, let $H_t = \mathcal{M}_k^{t-1}$ be the set of histories at the beginning of period t. A (complete information) pure strategy of player i is a tuple $\sigma = \left(\sigma^M, \sigma^D, \sigma^m\right)$, where $\sigma^M : \bigcup_{t \in T_i} H_t \to \Delta\mathcal{M}_k$ is the choice of mechanism when player i makes an offer, $\sigma^D : \bigcup_{t \in T_{-i}} H_t \times \mathcal{M}_k \to \Delta\{A, R\}$ is the decision about player $-i$'s offer , and $\sigma^m : \bigcup_t H_t \to \Delta\Sigma_i^m$ is the behavior in the proposed and accepted mechanism m. Here, Σ^m is the set of pure strategies in the mechanism m. Let Σ_i be the set of complete information strategies of player i.

An *assessment* is defined as a tuple of $(\sigma_A, \sigma_B, \mu)$, where measurable mapping $\sigma_A : \mathcal{U} \to \Delta\Sigma_A$ is Alice's strategy, $\sigma_B \in \Delta\Sigma_B$ is Bob's strategy, and $\mu_t : \bigcup_{t \in T_A} H_t \times \mathcal{M}_k \cup \bigcup_{t \in T_B} H_t \times \mathcal{M}_k \times \{A, R\} \to \Delta\mathcal{U}$ is a *belief function* that specifies Bob's beliefs about

[5]There are two main reasons for the lack of existence result in our model. First, the space of mechanisms is not compact, and hence the existence of a best response is not guaranteed. More importantly, there are well-known problems with the existence of sequential equilibrium in signaling games with infinitely many actions (see Myerson and Reny (2015)).

Alice's types, either after she offers a mechanism or after she makes a decision about Bob's offer. The implicit restriction is that the beliefs are updated only after Alice's actions.

A *Perfect Bayesian equilibrium* (or, simply, equilibrium) is an assessment such that (a) the players best respond to their strategies and beliefs, and (b) $\mu = \mu_0$ and the beliefs are updated through Bayes's theorem after each of Alice's decisions (mechanism choice or acceptance) that has a positive probability given history and strategies. Because the action choices are finite at each decision node, the PBE exists by the standard argument due to Selten (1975).

An equilibrium outcome (e_A, e_B) is a (measurable) function $e_A : \mathcal{U} \to [0,1]$ 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 $E^j(\delta, \mathcal{M}_k, \mu)$ be the set of expected equilibrium outcomes in game $\Gamma^j(\delta; \mathcal{M}_k, \mu)$. We are interested in the equilibrium outcomes as, first, the mechanism space becomes well approximated as $k \to \infty$, followed by offers become increasingly frequent as $\delta \to 1$:

$$E^j(\delta, \mu) = \sup_{(\mathcal{M}_k):\mathcal{M}_k \to \mathcal{M}_F} \limsup_{k\to\infty} E^j(\delta, \mathcal{M}_k, \mu) = \bigcup_{(\mathcal{M}_k):\mathcal{M}_k \to \mathcal{M}_F} \bigcap_n \mathrm{cl} \bigcup_{k \geq n} E^j(\delta, \mathcal{M}_k, \mu),$$

$$E^j(\mu) = \limsup_{\delta \to 1} E^j(\delta, \mu) = \bigcap_n \mathrm{cl} \bigcup_{\delta \geq 1 - \frac{1}{n}} E^j(\delta, \mu).$$

The closure is taken with respect to the topology of uniform convergence. We show below (see the comment after Lemma 2) that each of the closed sets in the above definitions is compact, and the intersections of compact sets are not empty.

3. Main result

In this section, we assume that $N = 2$.

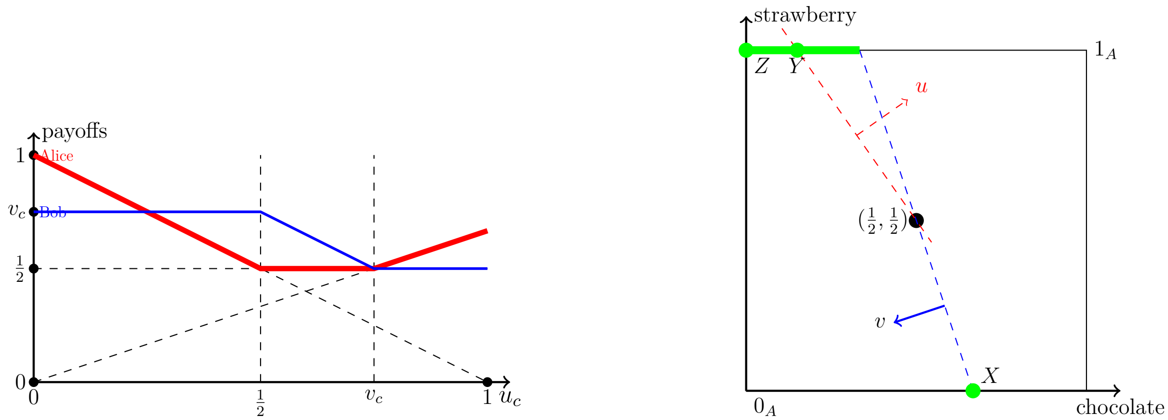

FIGURE 1. Nash payoffs and allocations.

3.1. **Complete information benchmark.** A special case of our model is when Alice's preferences are commonly known to be u. The argument from Rubinstein (1985) implies that our game has the unique subgame perfect equilibrium payoffs . When $\delta \to 1$, the payoffs converge to the Nash bargaining payoffs $(\mathcal{N}_A(u), \mathcal{N}_B(u))$.

When $v_c \geq v_v$, the Nash payoffs are given by function

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

The Nash payoffs and allocations are illustrated in Figure 1. Suppose Bob prefers chocolate. If Alice likes chocolate even more, she is going to get her favorite allocation subject to the constraint that Bob's payoff is at least $\frac{1}{2}$ (allocation X). Bob's payoff is $\frac{1}{2}$; Alice gets a larger payoff. If Alice has the same preference as Bob, allocation $\left(\frac{1}{2}, \frac{1}{2}\right)$ or any other allocation which gives both of them payoff of $\frac{1}{2}$ is a Nash allocation. If Alice prefers chocolate to strawberry, but likes chocolate to a lesser extent than Bob, Bob receives his favorite allocation subject to the constraint that Alice's payoff is at least $\frac{1}{2}$ (allocation Y if Alice's preferences are u). The allocation and Bob's payoff depends on Alice's preference; Alice's payoff is $\frac{1}{2}$. Finally, if Alice prefers strawberry to chocolate, each player receives his or her favorite part of the pie (allocation Z).

Figure 1 makes clear that Alice does not always have the incentive to honestly reveal her type. In particular, if she likes strawberry more than Bob does it, she is best off if Bob thinks that she her preferences are as close to his as possible.

3.2. **Menus.** (Alice's) finite and infinite menus play an important role in the analysis. Formally, a menu is any compact subset $Y \in \mathcal{C}X$. Each menu has its dual characterization through the payoff function $y(u;Y) = \max_{x \in Y} u(x)$. Any payoff function obtained from a menu is called a *menu function.* Let $\mathcal{Y}$ be the set of all menu functions.

For each u, and payoff function $y : \mathcal{U} \to [0,1]$, let $D_u y$ be the set of all affine functions $l : \mathcal{U} \to [0,1]$ such that $l(u) = y(u)$ and $\forall_{u'} y(u') \geq l(u')$.

Lemma 1. *Payoff function y is a menu function if and only if y is convex, continuous, and for each $u \in \mathcal{U}$, $D_u y$ is non-empty and closed. The set of menu functions $\mathcal{Y}$ is compact under the topology of the uniform convergence.*

In the dual approach, the "derivative" set $D_u y$ can be interpreted as the set of optimal choices l for Alice's type u, where Alice's share of the nth part of the pie is equal to $l(\omega^n)$. Given Alice's choice l, Bob's payoff is equal to $1 - l(v)$. This leads to a tight upper bound on Bob's expected payoff in a menu associated with a payoff function y: for each belief μ, let

$$\Pi(y,\mu) = \int \left(\max_{l \in D_u y} (1 - l(v)) \right) d\mu(u) = 1 - \int \left(\min_{l \in D_u y} l(v) \right) d\mu(u).$$

Lemma 2. *For each δ, μ, and k, if $(e_A, e_B) \in E^j(\delta, \mathcal{M}_k, \mu)$ is an equilibrium outcome, then e_A is a menu function, and $e_B \leq \Pi(e_A, \mu)$.*

Lemma 2 implies that $E^j(\delta, \mathcal{M}_k, \mu) \subseteq \mathcal{Y} \times [0,1]$. Because of standard arguments, the set of equilibrium outcomes is closed under uniform topology. Therefore, the equilibrium payoffs sets are closed subsets of a compact space. In particular, for each sequence $e_k \in E^j(\delta, \mathcal{M}_k, \mu)$, there exists a convergent subsequence with a limit $e_k \to e \in E^j(\delta, \mu)$. Similarly, for each sequence $e_{\delta_k} \in E^j(\delta_k, \mu)$ st. $\delta_k \to 1$, there exists a convergent subsequence with limit $e_k \to e \in E^j(\mu)$.

3.3. **Equilibrium payoffs.** For each belief μ and each function $c : \mathcal{U} \to [0,1]$, define

$$\Pi_{\text{opt}}(c, \mu) = \max_{y \in \mathcal{Y}, y \geq c} \Pi(y, \mu) \text{ and } \mathcal{M}_{\text{opt}}(c, \mu) = \arg\max_{y \in \mathcal{Y}, y \geq c} \Pi(y, \mu).$$

Here, $\Pi_{\text{opt}}(c, \mu)$ is the largest payoff that Bob can attain with a menu that ensures that each Alice's type u gets at least $c(u)$. Because of the compactness of $\mathcal{Y}$, the set of optimal menus $\mathcal{M}_{\text{opt}}(c, \mu)$ is non-empty. The optimal menu is unique for generic beliefs. We can state the main result of this paper.

Theorem 1. *Suppose that $N = 2$. Then, $E^j(\mu) \subseteq \mathcal{M}_{opt}(\mathcal{N}_A, \mu) \times \{\Pi_{opt}(\mathcal{N}_A, \mu)\}$.*

In the limit, as the space of available mechanisms becomes dense, and offers become increasingly frequent, regardless who makes the first offer, Bob's equilibrium payoff is equal to the expected payoff from the optimal screening menu subject to the constraint that each Alice type receives her Nash (i.e., complete information) payoff. If the optimal screening menu is unique, the payoff of each Alice's type is also unique. The outcome is ex ante efficient, but not ex post efficient. Bob's payoff is equal to the payoff that he would obtain if he were able to commit to an optimal mechanism subject to the complete information constraint. This is opposite to what one should expect from the literature on the Coase conjecture.

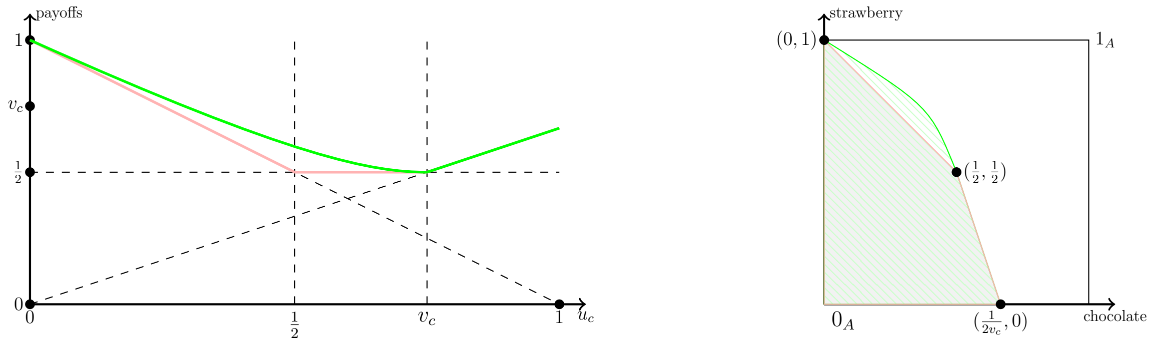

FIGURE 2. Payoffs and allocations in the constrained optimal menu.

The optimal menu is illustrated in Figure 2. If Alice likes chocolate more than Bob does, or if Alice only likes strawberry, the equilibrium allocation is equal to the complete information Nash Allocation. Otherwise, for generic beliefs, Alice receives a larger payoff.

We are unable to construct any equilibrium in this game. However, the equilibrium behavior cannot be too much different than as 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.

Bob's optimal payoff, $\mathcal{M}_{\text{opt}}(\mathcal{N}_A, \mu)$, is convex in μ, which implies that it has a natural comparative statics with respect to information: Bob is better off if his information improves. When $\mu \to \delta_u$ for some Alice type, Bob's payoff converges to $\mathcal{N}_B(u)$, the Nash outcome of bargaining against type u.

We briefly explain the structure of the proof, with emphasis on the roles of the mechanisms. The proof has two parts. In the first part, we start with the following key observation. We show that if Alice's payoffs are too large, in a sense that *each* of her types on the support receives a payoff that is significantly larger than her complete information payoff, then Bob has a profitable deviation in the form of a menu mechanism. The deviation is chosen so that Bob prefers to wait until it is accepted, and Alice prefers to accept it rather than wait for an opportunity to return to her original proposal. In order to find such a deviation, it is important that Bob has access to a sufficiently large set of menus. In general, it is not possible to find an optimal deviation in the form of a simpler mechanisms, like single offers. The reason is that, typically, a single offer is not acceptable to some of the types.

The above observation implies that the equilibrium payoff of at least some of types from Alice's support cannot be larger than her complete information payoff. This further implies that any menu with payoffs above complete information is going to be accepted by Alice. Indeed, if not all types from the support accept such a menu, than at least some of them must anticipate payoffs not higher than complete information in the continuation game. But then, today's rejection of the menu with payoffs strictly above complete information cannot be a best response.

If any menu with Alice's payoffs above complete information is accepted, Bob's payoff cannot be lower than the optimal payoff from such menus. In particular, this argument shows that Bob's ability to commit to optimal menu (subject to the constraint) is due to Alice's fear of revealing too much information about herself.

In the second part, we show that Alice's equilibrium payoffs cannot be significantly lower than the Nash payoffs. Otherwise, in the game in which Bob makes the first offer, each Alice's type u would have a profitable deviation to reject the Bob's offer and propose a new mechanism in the next period. We choose the deviation such that Alice's counteroffer is accepted by Bob and improves type u's payoffs. In the proof, Alice's counteroffer is a particular menu of menus: for each $y_u \in [0,1]$, let

$$W_{u,y_u} = \{\text{all menus with menu functions } y \text{ st. } y(u) \geq y_u\}. \tag{2}$$

An offer of W_{u,y_u} can be interpreted as Alice's request to Bob: "I am type u. You can design any menu as long as type u gets at least y_u." Formally, W_{u,y_u} consists of infinitely many menus, and hence it is not a mechanism as per our definition. However, we show that as $k \to \infty$, W_{u,y_u} can be well-approximated by a sequence of menus of menus $W^k_{u,y_u} \in \mathcal{M}_k$.

The menu of menus helps Alice to address the signaling problem. When Alice deviates from the equilibrium path (first by rejecting, and then, possibly, by making a particular counteroffer), Bob's beliefs are not constrained by the solution concept. In fact, his beliefs may maximize his incentives to reject the deviation. By offering an approximate version of W_{u,y_u}, Alice allows Bob to choose an acceptable mechanism (subject to the constraint $y(u) \geq y_u$), *whatever his beliefs.* There is a simple intuition. If Bob's believes that Alice's type is u, he will accept W_{u,y_u} as long as y_u is not larger than Alice's complete information payoff. If he believes that her type is $u' \neq u$, he will find a menu in W_{u,y_u} that extracts as much payoff form type u' as possible,

The above argument makes clear that the Theorem holds as long as the limit set of available mechanisms $\mathcal{M}$ contains all menus and menus of menus. Section 4.1 argues the thesis of the Theorem 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, one can show that $\Pi_{\text{opt}}(\mathcal{N}_A, \mu)$ is a lower bound on Bob's equilibrium payoffs.

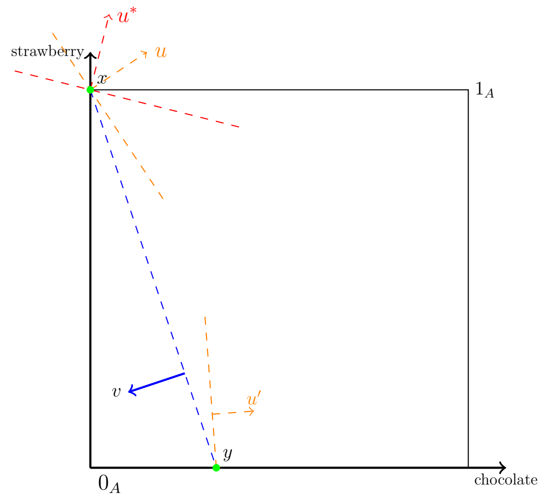

FIGURE 3. Allocations in the Anti-Coasian equilibrium.

4. COMMENTS[6]

4.1. **Simple offers.** We consider a special case of our model wherein players are only allowed to make simple offers. Let $\mathcal{S} \subseteq \mathcal{M}$ be the collection of all single-offer mechanisms (all mechanisms in which no player chooses any action). Any such a mechanism can be identified with a single allocation x. Let $X(\mathcal{S})$ be the collection of all such offers. Let $\mathcal{S}_1 \subseteq \mathcal{S}_2 \subseteq ... \subseteq \mathcal{S}$ be an approximation sequence with finite sets, where $\mathcal{S}_k \to S$ in the sense of the Hausdorff distance.

Proposition 1. *Suppose that $v_c > v_s$. Fix $u^* \in \mathcal{U}$ st. $u_c^* < v_c$. There exists δ_0 and k_0 such that, for each $\delta \geq \delta_0$, $k \geq k_0$, and any belief μ st. $u_c^* = \inf \{u_c : u_c \in \operatorname{supp} \mu\}$, there is $(e_A, e_B) \in E^B(\delta, \mathcal{S}_k, \mu_0)$ such that $\forall_u e_A(u) \leq \max_{x:v(x)\geq \delta \mathcal{N}_B(u^*)} u(x)$ and $e_B \geq \delta \mathcal{N}_B(u^*)$.*

Suppose type u^* has the strongest preference for strawberry in the support of Bob's beliefs. If players are sufficiently patient, then there exists an equilibrium in which

[6] All the proofs of results from this section are in the Online Appendix.

Bob receives his complete information payoff $\mathcal{N}_B(u^*)$ as if facing type u^*, regardless of his beliefs. This is also his best complete information payoff across Alice's types. Alice types receive the best payoffs subject to the constraint that Bob's payoff is at least $\mathcal{N}_B(u^*)$. Figure 3 illustrates the result. 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 Bob's indifference curve. Alice types, generically choose between two allocation: x and y.

The proof constructs such an equilibrium with required properties. Roughly, the idea is that Alice must offer 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 Anti-Coasian flavor, as Bob receives his best possible complete information payoff across all Alice's types. All Alice types are weakly and some are strictly worse off than under complete information. Alice would therefore benefit from being able to credibly reveal her type. This observation may be surprising to the reader familiar with the informed principal literature. Maskin and Tirole (1990) claim that, in the private value case, the informed principal must benefit from incomplete information due to the collapse of agent incentive and individual rationality constraints. This observation does not necessarily hold with interdependent values. We work with private values, but, in the dynamic setting, the continuation value depends on the belief of the uninformed agent. Hence, there is endogenous interdependence.

There are potentially other equilibria. For some beliefs, it is easy to extend the construction from Proposition 1, in which Alice types that prefer strawberry receive allocation x and the types that prefer chocolate receive an allocation with more chocolate than y.

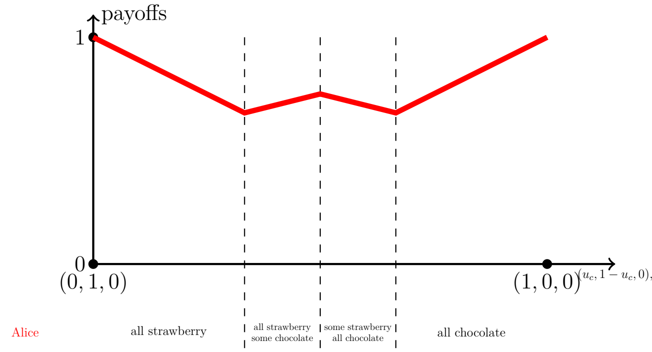

FIGURE 4. Nash payoffs when $N = 3$.

4.2. **Case $N > 2$.** The thesis of Theorem 1 does not hold when $N = 3$. We demonstrate this with an example. Let $v = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$. Figure 4 presents the Nash payoffs for all types who only care about the first two dimensions. Let $\tau_1 = \left(\frac{2}{3}, \frac{1}{3}, 0\right), \tau = \left(\frac{1}{2}, \frac{1}{2}, 0\right), \tau_2 = \left(\frac{1}{3}, \frac{2}{3}, 0\right)$ be three distinguished 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.

Proposition 2. *Fix $\delta < 1$, and suppose that players have an access to a public randomization. Suppose that $\mathcal{M}' \subseteq \mathcal{M}$ is a finite set of mechanisms that contains two menus:*

$$Y^A = \left\{\left(1, 2\frac{1-\delta}{\delta}, 0\right), \left(\frac{1}{\delta}\frac{2}{3}, \frac{1}{\delta}\frac{2}{3}, 0\right), \left(2\frac{1-\delta}{\delta}, 1, 0\right)\right\},$$
$$Y^B = \left\{(1,0,0), \left(\frac{2}{3}, \frac{2}{3}, 0\right), (0,1,0)\right\}.$$

(The tuples in the menus correspond to Alice's shares; Bob receives the complementary shares.) Then, for any belief $\Delta\{\tau_1, \tau_2, \tau\}$ that assigns a strictly positive probability to types τ_1, τ_2, there exists $\left(e^j_A, e^j_B\right) \in E^j(\delta, \mathcal{M}', \mu)$ such that $e^A_A(\tau) = \frac{1}{\delta}\frac{2}{3}, e^B_A(\tau) = \frac{2}{3}$. In particular, if $\delta < 1$ is sufficiently high, type τ receives a payoff substantially lower than her Nash payoff of $\frac{3}{4}$.

We construct a (sequential) equilibrium with the required features. In the equilibrium, player j always offers menu Y^j, and the offer is accepted. If Alice proposes some other mechanism (as a possible deviation), we show that there are Bob's beliefs such that either the mechanism has an equilibrium where all Alice's types receive payoffs smaller than $\frac{1}{\delta}\frac{2}{3}$, or Bob's payoffs are smaller than the discounted continuation equilibrium payoff of $\delta\frac{2}{3}$, in which case Bob rejects it. Similarly, we show that if Bob deviates, his proposal either has an equilibrium where he receives less than his equilibrium payoff of $e_B^B = \frac{2}{3}$, or all Alice types receive less than $\frac{2}{3}$, in which case Alice would reject it.

Although we are unable to fully characterize the set of payoffs when $N \geq 3$, we derive the following bound.

Theorem 2. *Suppose that $N \geq 2$. Then, for any $j = A, B$, any belief $\mu \in \Delta U$, any limit payoff $(e_A, e_B) \in E^j(\mu)$, 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).

5. Proof of Theorem 1

We divide the proof into two parts. In the first part, we show that any menu giving each Alice type payoffs strictly above her complete information payoffs will be accepted. In the second part, we show that a menu giving each Alice type less than her complete information payoffa will be rejected.

5.1. **Upper bound.** We show that if Alice's payoffs are too large when she makes the first offer, then Bob has a profitable deviation.

We proceed in two steps. The first step can be understood as a generalization of the method from Rubinstein (1982) to situations where player's payoffs are described by a function rather than a single number. We define a property of a payoff function, and we show that the payoff of any Alice's type in the belief support cannot be larger than a value of a payoff function with the property. Take any $\gamma < 1$ and $\delta < 1$. We say that menu function $h : \mathcal{U} \to [0,1]$ *has $UB(\gamma, \delta)$-property* if $\inf h > 0$ and, for each menu function $y \geq h$ and each belief μ, there exists a menu function y' such that $y' \geq_{\text{supp}\mu} (1 - \gamma(1-\delta)) y$ and $\gamma\delta\Pi(y', \mu) \geq \Pi(y, \mu)$.

Lemma 3. *For all $\gamma, \delta_0 < 1$, there exists k_0 such that, if function h has $UB(\gamma, \delta)$-property, then if $(e_A, e_B) \in E^A(\delta, \mathcal{M}_k, \mu)$ for some $k \geq k_0$ and μ, then $\inf_{u \in supp\mu} e_A(u) < h(u)$.*

The Lemma says that for any equilibrium, there must be an Alice's type in the belief support such that her payoff is strictly smaller than the value of a function with UB-property. The proof is by contradiction. Suppose that there exists an equilibrium where Alice attains payoffs above h. We show that we can find maximal equilibrium payoffs which such a property. Let y be Alice's payoff function in such an equilibrium. Because y is larger than h, the UB-property implies the existence of a menu function y' such that each Alice's type receives more than δy and such that Bob receives more that $\frac{1}{\delta}\Pi(y, \delta)$, or more than $\frac{1}{\delta}$ times his equilibrium payoff (by Lemma 2). If Bob rejects Alice's current offer, and proposes menu with payoffs y' instead, his counteroffer is accepted by all types in the support. (If not, because payoffs y are maximal, some rejecting types have to receive less than y in the continuation equilibrium, and rejecting more than $y' > \delta y$ is not profitable.) Because the deviation leads to payoffs that are larger than $\frac{1}{\delta}$ times the current equilibrium payoff, this contradicts the existence of equilibrium with payoffs y.

The need to use $\gamma < 1$ in the above definition and the result is due to the fact that we work with approximate spaces of mechanisms rather than all mechanisms.

Lemma 4. *Suppose that function h has property $UB(\gamma, \delta)$ for some $\gamma, \delta < 1$. It follows that, if $(e_A, e_B) \in E^B(\delta, \mu)$, then $e_B \geq \Pi_{opt}(\delta h, \mu)$.*

Proof. We present the proof in-text, because it is short and intuitive. Let k_0 be as in Lemma 3 . We are going to show that, for any $k \geq k_0$, in game $\Gamma^B(\delta, \mathcal{M}_k, \mu)$, Alice accepts with probability 1 any menu $m \in \mathcal{M}_k$ yielding payoffs strictly larger than δh for each of her types in the support. On the contrary, if a positive probability set of types rejected such a menu, they would face a continuation equilibrium described in Lemma 3. By Lemma 3, at least some of those types would receive a continuation payoff that is strictly lower than h; so, their rejection of m could not have been a best response.

For each $k \geq k_0$, let $\mathcal{M}_k^* \subseteq \mathcal{M}_k$ be the set of such menus. The argument implies that, if $(e_A, e_B) \in E^B(\delta, \mu, \mathcal{M}_k)$, it must that $e_B \geq \max_{m \in \mathcal{M}_k:\mathrm{m}} \Pi(y_m, \mu)$. By the approximation results (Lemma 12 from the Appendix), the RHS off the above inequality converges to $\Pi_{\mathrm{opt}}(\delta h, \mu)$ as $k \to \infty$. □

In the second step, we show that an approximation to the Nash payoffs has the UB-property.

Lemma 5. *Suppose that $N = 2$. Then, for each $\varepsilon > 0$, there is a function h such that $\sup_u |h(u) - \mathcal{N}_A(u)| \leq \varepsilon$ and $\gamma, \delta_0 < 1$ such that h has $UB(\gamma, \delta)$ property for each $\delta \geq \delta_0$.*

We sketch the intuition. We want to show that menu function $h > \mathcal{N}_A$ has the UB property. Suppose that $y > h$ is a menu function associated with menu Y. Let $x(u) \in \arg\max_{x\in Y} u(x)$ be an optimal choice of type u in menu Y. We use the 45° diagonal to divide the space of allocations into two areas. For each u such that $x(u)$ is below the diagonal, define a wasteless allocation $x'(u)$ so that $x'_A(u) = \delta x_A(u) + (1-\delta) x_A(u)$. Similarly, for each u such that $x_Y(u)$ is above diagonal, define a wasteless allocation $x'(u)$ so that $x'_A(u) = x_A(u) + \rho(x_A(u) - \mathbf{1}_A)$, where constant $\rho = (1-\delta)\frac{\alpha}{1-\alpha}$ is chosen so that the two definitions agree for u^* st. $x(u^*) = \alpha\mathbf{1}_A + (1-\alpha)\mathbf{0}_A$ lies exactly on the diagonal (see Figure 5). Because $y > \mathcal{N}_A$, we check that $\alpha > \frac{1}{2}$, and $\rho > 1-\delta$. (In the latter case, so defined $x'(u)$ will only exist if $x(u)$ is sufficiently inside X, or, in other words, if $1-\delta$ is small and $u(x(u)) \geq h(u)$ is bounded away from $\mathcal{N}_A(u)$. The details can be found in the Appendix.) Let $Y' = \{x'(u) : u \in \mathcal{U}\}$. Because the new allocations are obtained by (partially) linear operations, it is straightforward to show that $x'(u)$ is the type u optimal choice in menu Y'. Further, we check by direct calculations that $u(x'(u)) \geq \delta u(x(u))$ and $\delta v(x'(u)) \geq v(x(u))$ for each u. For instance, suppose that $x(u)$ is below the diagonal. Then, the second inequality follows easily from the fact that $\rho > 1$. For the first inequality, notice that

$$u(x'(u)) - \delta u(x(u)) = (1+\rho-\delta)\, u(x(u)) - \rho$$
$$= (1-\delta)\frac{1}{1-\alpha}(u(x(u)) - \alpha) \geq (1-\delta)\frac{1}{1-\alpha}(u(x(u^*)) - \alpha) = 0,$$

where the inequality comes from the fact that $x(u^*)$ is one of the choices available to type u. The two inequalities imply that the menu function y' satisfies $y' \geq \delta y$ and $\delta\Pi(y', \mu) \geq \Pi(y, \mu)$.

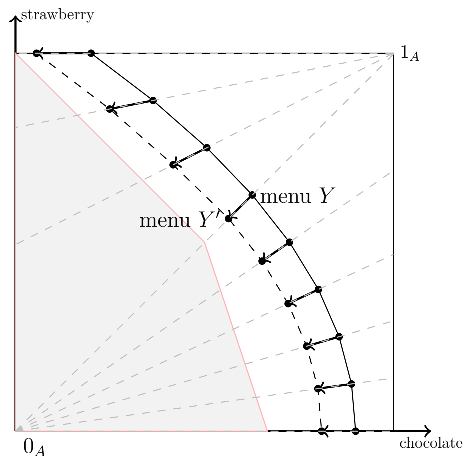

FIGURE 5. Menus Y and Y'.

5.2. **Lower bound.** As in the upper bound case, we begin with a certain generalization of the Rubinstein method. For $\gamma, \delta < 1$, say that payoff function $h : \mathcal{U} \to [0,1]$ has $LB(\gamma,\delta)$-property if for each type w, each constant $y_w \le h(w)$, any belief $\psi \in \Delta\mathcal{U}$,

$$\Pi_{\text{opt}}\left(\left(\frac{1}{\gamma\delta}y_w\right)\mathbb{1}_{.=w},\psi\right) \ge (1-\gamma(1-\delta))\,\Pi_{\text{opt}}\left(y_w\,\mathbb{1}_{.=w},\psi\right). \tag{3}$$

Here, $\mathbb{1}_{.=w}$ is a function of types equal to 1 if $u = w$ and to 0 otherwise.

Lemma 6. *For all $\gamma, \delta < 1$, there exists $k_0 < \infty$ such that if h has $LB(\gamma,\delta)$ property, it follows that for each $k \ge k_0, \mu \in \Delta\mathcal{U}$, if $(e_A, e_B) \in E^B(\delta, \mathcal{M}_k, \mu)$, then, $e_A \ge h$.*

To see the intuition, take an arbitrary type $w \in U$. Let y_w be the minimum of the equilibrium payoffs of type w across all equilibria and all beliefs in a game where Bob makes the first offer assuming for simplicity that such an minimum is attained in some equilibrium. Suppose that $y_w < h$. Given such an equilibrium with payoffs $y_w = e_A(w)$, we consider Alice's deviation to reject any offer in the first period and to propose menu of menus $W_{w,\frac{1}{\gamma\delta}y_w}$ in the next period. Such an offer induces (possibly, off-path) equilibrium beliefs ψ. If Bob accepts the menu, his payoff will be equal to the left-hand side of (3). In contrast, if Bob rejects Alice's offer, the continuation equilibrium in the next period yields at least y_w to Alice type y_w. Thus, Bob's expected and discounted continuation payoff is equal to $\delta\Pi_{\text{opt}}\left(\mathbb{1}_{.=w}\, y_w, \psi\right)$. Inequality (3) implies that Bob prefers to accept the menu of menus. However, such a menu of menus leads to a payoff of $\frac{1}{\gamma\delta}y_w$ for Alice type w. Hence, the deviation is profitable, which contradicts the existence of equilibrium with payoff y_w.

Lemma 7. *For any $\varepsilon > 0$, there exist $\gamma, \delta < 1$ such that the following h functions have $LB\left(\gamma, \delta\right)$ property:*

(1) $h^0\left(u\right) = \left(1-\varepsilon\right)\frac{1}{2}$,
(2) $h^k\left(u\right) = \left(1-\varepsilon\right)\min\left(1, \frac{1}{2v_k}\right)u_k$ *for any $k = c, s$.*

In both cases, we take arbitrary menu function y such that $y\left(w\right) \geq y_w$, where $y_w \leq h\left(w\right)$, and use it to construct a menu function y' such that $y'\left(w\right) \geq \frac{1}{\gamma\delta}y_w$ and such that Bob with beliefs ψ attains larger expected payoffs from accepting menu y' than waiting for the next period and menu y, $\Pi\left(y', \psi\right) \geq \delta\Pi\left(y, \psi\right)$. For the first claim, it is enough to replace menu associated with y by its convex combination with allocation $\mathbf{1}_A$ with weight $1-\delta$, or in other words, by a menu associated with $y' = \delta y + 1 - \delta$. Similarly, in the case of the second claim, we also replace the original menu, but with an appropriate convex combination with an allocation that gives the entire part k of the pie to Alice.

5.3. **Proof of Theorem 1.** Assume w.l.o.g. that Bob likes chocolate more, i.e., $v_c \geq v_s$ It follows from Corollary 4 and Lemma 5 that, if $\left(e_A, e_B\right) \in E^B\left(\mu\right)$, then $e_B \geq \Pi_{\text{opt}}\left(\mathcal{N}_A, \mu\right)$. Similarly, it follows from Lemmas 6 and 7 that, for each $\left(e_A, e_B\right) \in E^B\left(\mu\right)$, it must be that

$$e_A(u) \geq \max\left(\frac{1}{2}, \min\left(1, \frac{1}{2v_c}\right)u_c, \min\left(1, \frac{1}{2v_s}\right)u_s\right) = \max\left(\frac{1}{2}, \frac{1}{2v_c}u_c, u_s\right) = \mathcal{N}_A(u),$$

where the last equality is a consequence of (1). The result follows from the definition of the value $\Pi_{\text{opt}}(.)$ and the solution $\mathcal{M}_{opt}(.)$ to Bob's optimization problem.

References

Abreu, D. and F. Gul (2000): "Bargaining and Reputation," *Econometrica*, 68, 85–117.

Doval, L. and V. Skreta (2018): "Mechanism Design with Limited Commitment," SSRN Scholarly Paper ID 3281132, Social Science Research Network, Rochester, NY.

Grossman, S. J. and M. Perry (1986): "Perfect sequential equilibrium," *Journal of Economic Theory*, 39, 97–119.

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. Sonnenschein, 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.

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

Myerson, R. B. and P. J. Reny (2015): "Sequential equilibria of multi-stage games with infinite sets of types and actions," *manuscript, University of Chicago*.

Peski, M. (2019): "Bargaining over Heterogeneous Good with Structural Uncertainty," .

Rubinstein, A. (1982): "Perfect equilibrium in a bargaining model," *Econometrica*, 50.

——— (1985): "A Bargaining Model with Incomplete Information About Time Preferences," *Econometrica*, 53, 1151–1172.

Rushowy, K. (2020): "Ontario government offers mediation amid faltering talks with teacher unions," *Toronto Star*.

Selten, R. (1975): "Reexamination of the perfectness concept for equilibrium points in extensive games," *International Journal of Game Theory*, 4, 25–55.

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. Menu functions

A.1. Menu functions.

Lemma 8. *Suppose that $N = 2$. Suppose that y is a convex function and types $u, u' \in \mathcal{U}$ such that $u_c < u'_c$. If $D_u y$ and $D_{u'} y$ are non-empty, then $D_w y$ is non-empty for any type w such that $u_c < w_c < u'_c$.*

Proof. Because y is convex, for any $u, u' \in \mathcal{U}$ st. $u_c < u'_c$, and any $l \in D_u y, l' \in D_{u'} y$, we have $l(0) > l'(0)$, and $l(1) < l'(1)$. □

Lemma 9. *For each menu function y_0, each belief μ, and each $\alpha \in [0, 1]$, we have $\Pi_{opt}(\alpha y_0 + (1 - \alpha)\mathbf{1}, \mu) \geq \alpha \Pi_{opt}(y_0, \mu)$.*

Proof. Observe that if $y \in \mathcal{M}_{\text{opt}}(y_0, \mu)$, then $y' = \alpha y + (1 - \alpha)\mathbf{1} \geq \alpha y_0 + (1 - \alpha)\mathbf{1}$. Because $D_u y' = \{\alpha l(.) + 1 - \alpha : l \in D_u y\}$, we have

$$\Pi(y', \mu) = 1 - \int \left(\max_{l \in D_u(\alpha y + (1-\alpha)\mathbf{1})} l(v) \right) d\mu(u) = 1 - \int \left(\max_{l \in D_u(\alpha y + (1-\alpha)\mathbf{1})} (\alpha l(v) + 1 - \alpha) \right) d\mu(u)$$

$$= \alpha - \alpha \int \left(\max_{l \in D_u(\alpha y + (1-\alpha)\mathbf{1})} l(v) \right) d\mu(u) = \alpha \Pi(y_0, \mu).$$

□

Lemma 10. *For any two menu functions $y, y' \in \mathcal{Y}$, $\max(y, y')$ is also a menu function.*

Proof. The convexity and the payoff restriction are immediate. For the bound on the derivative, notice that $D_u \max(y, y') \subseteq D_u y \cup D_u y'$. □

A.2. Proof of Lemma 1.

For the "only if" direction, notice that

$$D_u y = \left\{ l : \forall_{u'} l(u') = u'(x) \text{ for some } x \in \arg\max_{z \in Y} u(z) \right\}.$$

For the "if" direction, for each affine function $l : \mathcal{U} \to [0,1]$, define an allocation $x(l)$ such that $x_{A,n}(l) = l(\omega^n)$ and $x_{B,n}(l) = 1 - x_{A,n}(l)$. Let $Y = \mathrm{cl}\{x(l) : l \in \bigcup_u D_u y\}$. Then, y is the payoff function for menu Y.

Because of the characterization, each menu function is Lipschitz with bounded payoffs. Hence, set $\mathcal{Y}$ is equicontinuous, and the Arzelà–Ascoli theorem implies that it is precompact under the topology of the uniform convergence. Let $y_n \in \mathcal{Y}$ be a convergent sequence and, for each u, let $l_n^u \in D_u y_n$ be a supporting affine function.

A.3. **Proof of Lemma 2.** Any profile of (complete information) strategies $\sigma \in \prod_i \Sigma_i$ induces a probability distribution $\mu(\sigma) \in \Delta(\mathbb{N} \times X \cup \{\infty\})$, with the interpretation that $\mu(t, A|\sigma)$ for measurable $A \subseteq X$ is the probability that the game ends in period t with outcome in A, and $\mu(\infty|\sigma)$ is the probability that the bargaining never ends. Let $x(\sigma) \in X$ be defined so that $x_i(\sigma) = (1-\delta)\sum \delta^t \int x_i \mu(t, dx)$ is the expected and discounted award to player i.

Suppose that $(\sigma_A, \sigma_B, \mu)$ is an equilibrium with payoffs $(e_A, e_B) \in E^j(\delta, \mathcal{M}_k, \mu)$. For each u, let $x_u = x(\sigma_A(u), \sigma_B)$ be the expected and discounted equilibrium allocation. Then, $u(x_{u'})$ is the expected payoff of Alice's type u from mimicking the strategy of type u'. For each u, define an affine $l_u : \mathcal{U} \to [0,1]$ so that $l_u(u') = u'(x_u)$. The incentive compatibility implies that for each u and u', $e_A(u') = l_{u'}(u') \geq l_u(u')$. It follows that e_A is convex (as it is a supremum over affine functions l_u), and $l_u \in D_u e_A$. By Lemma 1, e_A is a menu function. Moreover,

$$e_B = \int v(x(u))\, d\mu(u) \leq \int (1 - l_u(v))\, d\mu(u) \leq \Pi(e_A, \mu).$$

Appendix B. Mechanisms and approximations

This part of the Appendix is devoted to approximations of mechanisms.

Let $\mathcal{M}_{\text{menu}}$ denote the class of menu mechanisms, and let $\mathcal{M}_{menu}(n)$ denote the subset of menus with n actions for Alice. Let $\mathcal{M}_{\text{mmenu}}$ denote the class of (Bob's) menus of (Alice's menus) and let $\mathcal{M}_{mmenu}(n_B, n_A)$ denote the subset of such mechanisms with n_i actions for player i. Let d_{CX} and d_{C^2X} be the Hausdorff distances on spaces, respectively, CX and C^2X, induced by the Euclidean distance on X. For ease of notation, we drop the subscripts in the definition of the distance.

For any mechanism m and any beliefs $\mu \in \Delta\mathcal{U}$, let $E(m, \mu) \subseteq \mathcal{Y} \times [0,1]$ denote the set of outcomes (e_A, e_B) that can be obtained in equilibrium.

B.1. **Menus.** For each menu Y and each $\eta > 0$, define menu

$$Y^\eta = \{(1-\eta)x + \eta v(x)\mathbb{1}_A + \eta(1 - v(x))\mathbb{0}_A : x \in Y\}.$$

Lemma 11. *For each $\eta > 0$, there exists $\varepsilon > 0$ such that for each menu Y, each menu Y' such that $d(Y^\eta, Y') \leq \varepsilon$, and each $u \in \mathcal{U}$, if $x \in \arg\max_{x \in Y} u(x)$ and $x' \in \arg\max_{x \in Y'} u(x)$, then*

$$u(x') \geq u(x) - 2\eta \text{ and } v(x') \geq v(x) - 2\eta. \tag{4}$$

Proof. Choose $\varepsilon > 0$ such that $\eta > \varepsilon$ and $\eta(\eta - \varepsilon) > 2\varepsilon$. Take any two menus Y and Y' such that $d(Y^\eta, Y') \leq \varepsilon$. Fix u.

For the first inequality in (4), notice that for each $x \in Y$, there exists $x' \in Y'$ such that $\|(1-\eta)x + \eta(1 - v(x))\mathbb{0}_A + \eta v(x)\mathbb{1}_A - x'\| \leq \varepsilon$. But

$$\begin{aligned} u(x') &\geq u((1-\eta)x + \eta(1 - v(x))\mathbb{0}_A + \eta v(x)\mathbb{1}_A) - \varepsilon \\ &= u(x) + \eta(v(x) - u(x)) - \varepsilon \geq u(x) - 2\eta, \end{aligned}$$

where the last inequality comes from the choice of ε. Hence, if $x' \in Y'(u)$, it must be that $u(x') \geq u(x) - 2\eta$.

For the second inequality in (4), suppose on the contrary that there is $x \in \arg\max_{x\in Y} u(x)$ and $x' \in \arg\max_{x\in Y'} u(x)$ such that $v(x') < v(x) - 2\eta$. Because $x' \in Y'$ and $d(Y^\eta, Y') \leq \varepsilon$, there exists $x_0 \in Y^\eta$ such that $\|x' - x_0\| \leq \varepsilon$, which implies that $u(x') \leq u(x_0) + \varepsilon$ and $v(x_0) \leq v(x') + \varepsilon \leq v(x) - 2\eta + \varepsilon$. Because $x_0 \in Y^\eta$, there exists $x_1 \in Y$ such that

$$x_0 = (1-\eta)\,x_1 + \eta\,(1 - v(x_1))\,\mathbb{0}_A + \eta v(x_1)\,\mathbb{1}_A\,.$$

Recall that $v(\mathbb{0}_A) = 1 = 1 - v(\mathbb{1}_A)$. Hence,

$$v(x_1) \leq v(x_1) + \eta\,(1 - 2v(x_1)) + \eta = v(x_0) + \eta \leq v(x) - \eta + \varepsilon.$$

(The first inequality comes from $v(x_1) \leq 1$; the second from $v(x_0) \leq v(x) - 2\eta$.)Moreover, because $x \in Y(u)$, we have $u(x_1) \leq u(x)$, which implies

$$\begin{aligned} u(x') &\leq u(x_0) + \varepsilon = u((1-\eta)\,x_1 + \eta\,(1 - v(x_1))\,\mathbb{0}_A + \eta v(x_1)\,\mathbb{1}_A) + \varepsilon \\ &= (1-\eta)\,u(x_1) + \eta v(x_1) + \varepsilon \leq (1-\eta)\,u(x_1) + \eta\,(v(x) - \eta + \varepsilon) + \varepsilon \\ &\leq (1-\eta)\,u(x) + \eta v(x) - \eta\,(\eta - \varepsilon) + \varepsilon. \end{aligned} \tag{5}$$

On the other hand, because $d(Y^\eta, Y') \leq \varepsilon$, there is $x_2 \in Y'$ such that

$$\begin{aligned} u(x_2) &\geq u((1-\eta)\,x + \eta\,(1 - v(x))\,\mathbb{0}_A + \eta v(x)\,\mathbb{1}_A) - \varepsilon \\ &= (1-\eta)\,u(x) + \eta v(x) - \varepsilon. \end{aligned} \tag{6}$$

Because of the choice of ε, inequalities (5) and (6) contradict $x' \in \arg\max_{x\in Y'} u(x)$. The contradiction demonstrates that $v(x') \geq v(x) - 2\eta$. □

Lemma 12. *For each $\eta > 0$, there exists k_0 such that for each $k \geq k_0$, each $y \in \mathcal{Y}$, there is a mechanism $m \in \mathcal{M}_k$ such that for each $\mu \in \Delta\mathcal{U}$, each $(e_A, e_B) \in E(m, \mu)$, we have $e_A(u) \geq y(u) - \eta$ for each u, and $e_B \geq \Pi(y, \mu) - \eta$.*

Proof. Because CX is compact in the Hausdorff metric, and finite menus form a dense subset, there is $n(\varepsilon) < \infty$ such that for each $Y \in CX$, there exists $m \in \mathcal{M}_{\text{menu}}(n(\varepsilon))$ such that $d(Y, Y_m) \leq \frac{1}{2}\varepsilon$. Further, because $\mathcal{M}_{menu}(n(\varepsilon))$ is a compact subset of Euclidean space, for each $\varepsilon > 0$, there exists $k_0(\varepsilon)$ such that for any $k \geq k_0(\varepsilon)$, any $m \in \mathcal{M}_{\text{menu}}(n(\varepsilon))$, there exists $m' \in \mathcal{M}_k \cap \mathcal{M}_{\text{menu}}(n(\varepsilon))$ such that $d(Y_m, Y_{m'}) \leq \frac{1}{2}\varepsilon$. It follows that for each $k \geq k_0(\varepsilon)$ and any $Y \subseteq X$, there is $m \in \mathcal{M}_k \cap \mathcal{M}_{\text{menu}}(n(\varepsilon))$ such that $d(Y, Y_m) \leq \varepsilon$.

Let $\varepsilon > 0$ be small enough so that the thesis of Lemma 11 holds for $\frac{1}{2}\eta$. Take arbitrary menu function y and the associated menu Y. Find mechanism m such that $d\left(Y^{\frac{1}{2}\eta}, Y_m\right) \leq \varepsilon$. The claim follows from Lemma 11. □

B.2. **Menus of menus.**

Lemma 13. *For each $\eta > 0$, there exists k_0 such that for each $k \geq k_0$ and each $(u, y_u) \in \mathcal{U} \times [0, 1]$, there is a mechanism $m \in \mathcal{M}_k$ such that for each $\mu \in \Delta\mathcal{U}$ and each $(e_A, e_B) \in E(m, \mu)$, we have $e_A(u) \geq y_u - \eta$ and $e_B \geq \max_{y \in \mathcal{Y}: y(u) \geq y_u} \Pi(y, \mu) - \eta$.*

Proof. Because the space of all menus of menus C^2X is compact in the Hausdorff metric, and finite menus of menus form a dense subset, there is $n_B(\varepsilon) < \infty$ such that for each $W \in C^2X$, there exists $m \in \mathcal{M}_{\text{mmenu}}(n_B(\varepsilon), n(\varepsilon))$ (where $n(\varepsilon)$ is as in the proof of Lemma 12) such that $d(W, W_m) \leq \frac{1}{2}\varepsilon$. Further, because the $\mathcal{M}_{menu}(n_B(\varepsilon), n(\varepsilon))$ is a compact subset of Euclidean space, for each $\varepsilon > 0$, there exists $k_0(\varepsilon)$ such that for any $k \geq k_0(\varepsilon)$ and any $m \in \mathcal{M}_{\text{mmenu}}(n_B(\varepsilon), n(\varepsilon))$, there exists $m' \in \mathcal{M}_k \cap$

$\mathcal{M}_{\text{mmenu}}\left(n_B\left(\varepsilon\right), n\left(\varepsilon\right)\right)$ such that $d\left(W_m, W_{m'}\right) \leq \frac{1}{2}\varepsilon$. It follows that for each $k \geq k_0\left(\varepsilon\right)$ and any $W \subseteq C^2X$, there is $m \in \mathcal{M}_k$ such that $d\left(W, W_m\right) \leq \varepsilon$.

Let $\varepsilon > 0$ be small enough so that the thesis of Lemma 11 holds for $\frac{1}{2}\eta$. Fix $k \geq k_0\left(\varepsilon\right)$. Take arbitrary $\left(u, y_u\right) \in \mathcal{U} \times [0,1]$. Construct menu W_{u,y_u} of all menus Y associated with menu functions y such that $y\left(u\right) \geq y_u$. Further, construct menu of menus $W^{\eta}_{u,y_u} = \left\{Y^{\frac{1}{2}\eta} : Y \in W_{u,y_u}\right\}$. Find mechanism $m \in \mathcal{M}_k$ such that $d\left(W^{\eta}, W_m\right) \leq \varepsilon$. The latter means that:

- for each $s_B \in S^1_B\left(m\right)$, there exists $Y \in W_{u,y_u}$ such that $d\left(Y^{\frac{1}{2}\eta}, Y_m\left(s_B\right)\right) \leq \varepsilon$. If y is the menu function associated with Y and $y_m\left(s_B\right)$ is associated with $Y_m\left(s_B\right)$, then by Lemma 11, we have that for any s_B, $y_m\left(u; s_B\right) \geq y\left(u\right) - \eta \geq y_u - \eta$,
- for each $Y \in W_{u,y_u}$, there exists $s_B \in S^1_B\left(m\right)$ such that $d\left(Y^{\frac{1}{2}\eta}, Y_m\left(s_B\right)\right) \leq \varepsilon$, and, if y is the menu function associated with Y and $y_m\left(s_B\right)$ is associated with $Y_m\left(s_B\right)$, then by Lemma 11, we have that $\Pi\left(y_m\left(s_B\right), \mu\right) \geq \Pi\left(y, \mu\right) - \eta$ for any belief $\mu \in \Delta U$. Therefore, $\max_{s_B} \Pi\left(y_m\left(s_B\right), \mu\right) \geq \max_{y \in \mathcal{Y}: y(u) \geq y_u} \Pi\left(y, \mu\right) - \eta$.

□

Appendix C. Proofs of upper bound of Theorem 1

C.1. **Basic equilibrium bound.**

Lemma 14. *For each* δ, k,

- $e_A\left(u\right) \geq 1 - \delta, e_B \geq \delta\left(1-\delta\right)$ *for each* u *and each* $\left(e_A, e_B\right) \in E^A\left(\delta, \mathcal{M}_k, \mu\right)$,
- $e_A\left(u\right) \geq \delta\left(1-\delta\right), e_B \geq 1 - \delta$ *for each* u *and each* $\left(e_A, e_B\right) \in E^B\left(\delta, \mathcal{M}_k, \mu\right)$.

Proof. Notice that Alice will accept menu $\delta\,\mathbb{1}_A + \left(1-\delta\right)\mathbb{0}_A$, which sets a lower bound on Bob's payoff. Anticipating that, Bob is going to reject any continuation equilibrium that gives him less that $\delta\left(1-\delta\right)$. The other case is analogous. □

C.2. **Proof of Lemma 3.** Let $\eta = \frac{1}{5}(\inf h)(1-\gamma)(1-\delta) > 0$. Let k_0 be such that Lemma 12 holds. Suppose that h has $UB(\gamma,\delta)$ property. Fix a belief μ.

Let $E = \left\{(e_A,\mu) : (e_A,e_B) \in E^A(\delta,\mathcal{M}_k,\mu)\right\}$. Suppose that there exists $(e_A,e_B,\mu) \in E$ such that $\forall_{u\in\text{supp}\mu} e_A(u) \geq h(u)$. By the remark after Lemma 2, E is compact, and we can find the equilibrium payoffs and the beliefs (e_A,μ) that are undominated in E the following sense: there is no (e'_A,μ') such that $e_A(u)+\eta < e'_A(u)$ for some $u \in \text{supp}\mu' \subseteq \text{supp}\mu$. Let $y = \max(h,e_A)$ and notice that y is a menu function by Lemma 10. Also, $y(u) = e_A(u)$ for each $u \in \text{supp}\mu$, and, by Lemma 14, $\Pi(y,\mu) = \Pi(e_A,\mu) \geq e_B \geq \delta(1-\delta)$.

By the definition of the UB property, there exists a menu function y' such that

$$y' \geq (1-\gamma(1-\delta))\,y \geq (1-\gamma(1-\delta))\,e_A, \tag{7}$$

and

$$\gamma\delta\Pi(y',\mu) \geq \Pi(y,\mu) = \Pi(e_A,\mu). \tag{8}$$

By Lemma 12, there exists a mechanism m such that for any $\mu \in \Delta\mathcal{U}$, each μ', each $(e'_A,e'_B) \in E(m,\mu')$, each u, we have

$$e'_A(u) \geq y'(u) - \eta \geq (1-\gamma(1-\delta))\,e_A(u) - \eta \geq \left(1-\left(1-\frac{1}{2}(1-\gamma)\right)(1-\delta)\right)e_A(u) - \eta.$$

The last inequality follows from the choice of

$$\eta \leq \frac{1}{4}(1-\gamma)(1-\delta)\,h(u) \leq \frac{1}{4}(1-\gamma)(1-\delta)\,e_A(u).$$

Additionally, we have

$$e'_B \geq \Pi(y',\mu) - \eta \geq \frac{1}{\gamma\delta}\Pi(e_A,\mu) - \eta = \frac{1}{\delta}\Pi(e_A,\mu) + \left(\frac{1-\gamma}{\gamma}\right)\frac{1}{\delta}\Pi(e_A,\mu) - \eta > \frac{1}{\delta}\Pi(e_A,\mu),$$

where the choice of η implies that $\frac{1-\gamma}{\gamma}\frac{\Pi(e_A,\mu)}{\delta} \geq \frac{1-\gamma}{\gamma}(1-\delta) > \eta$.

Consider an equilibrium that supports an outcome (e_A, μ). Let (f, ψ) be an continuation equilibrium outcome with beliefs starting from period 3 after a history such that in period 1, Bob rejects; in period 2, Bob proposes mechanism m with continuation payoffs $(e'_A, e'_B) \in E(m, \mu')$ that Alice rejects. (Here, μ' is a belief after the mechanism m is accepted.) Let $A_R \subseteq \text{supp}\mu$ be the set of Alice types for whom the rejection in period 2 is a (possibly weak) best response. For each $u \in A_R$, it must be that

$$f(u) \geq \frac{1}{\delta} e'_A(u). \tag{9}$$

Suppose that the rejection in period 2 occurs with positive probability. If so, then ψ is absolutely continuous wrt μ, and, in particular, $\text{supp}\psi \subseteq A_R \subseteq \text{supp}\mu$. Because (e_A, μ) is undominated in E, there must be $u_0 \in \text{supp}\psi$ such that $y(u_0) \geq f(u_0) - 2\eta$. By (7) and (9) ,

$$\begin{aligned} e_A(u_0) &\geq f(u_0) - 2\eta \geq \frac{1}{\delta} e'_A(u_0) - 2\eta \geq \frac{1}{\delta}\left(1 - \left(1 - \frac{1}{2}(1-\gamma)\right)(1-\delta)\right) e_A(u_0) - 2\eta \\ &\geq e_A(u_0) + \frac{1}{\delta}\frac{1}{2}(1-\gamma)(1-\delta)\, e_A(u_0) - 2\eta \\ &\geq e_A(u_0) + (1-\gamma)(1-\delta)\left[\frac{1}{2} e_A(u_0) - \frac{2}{5}(\inf h)\right] > e_A(u_0), \end{aligned}$$

where the last inequality follows from the fact that $e_A(u_0) \geq h(u_0)$. The contradiction shows that period 2 offer m is accepted with probability 1. By Lemma 2, $e_B \leq \Pi(e_A, \mu)$. On the other hand, Bob's strategy to reject any offer in period 1 and propose m leads to the expected discounted payoff of $\delta e'_B > \Pi(e_A, \mu)$. But this leads to a contradiction with choice (y, μ) as an equilibrium.

C.3. **Proof of Lemma 5.** We assume w.l.o.g. that $v_c \geq v_s$. We consider the following two cases separately:

- $v_c > v_s$, or Bob prefers chocolate to strawberry; this case is further divided into two sub-cases that depend on the offered menu,
- $v_c = v_s$, or Bob is indifferent between chocolate and strawberry.

Because we work with $N = 2$, it is possible and convenient to redefine any function of type $u = (u_c, 1 - u_c)$ (i.e., $l(u)$, $\mathcal{N}_A(u)$, etc.) as a function of the first coordinate (i.e., $l(u_c)$, $\mathcal{N}_A(u_c)$, etc.).

For each menu function y, each $u \in \mathcal{U}$, define

$$d(y, u) = \max_{l \in D_u y} l(v). \tag{10}$$

C.3.1. *Case A:* $v_c > v_s$*.* Assume that $\varepsilon > 0$ is sufficiently small that

$$\mathcal{N}_A(1)(1+\varepsilon) = \frac{1}{2v_c}(1+\varepsilon) \leq 1 \text{ and } v_c > \frac{1}{2-\varepsilon}. \tag{11}$$

Define function h as

$$h(u_c) = \begin{cases} (1-\varepsilon)\mathcal{N}_A(u_c) + \varepsilon & \text{if } u_c \leq \frac{1}{2} \\ \mathcal{N}_A(1)(1+\varepsilon) & \text{if } u_c \geq \frac{1}{2}. \end{cases}$$

Notice that function h is continuous due to the fact that $\mathcal{N}_A\left(\frac{1}{2}\right) = \frac{1}{2}$.

Find δ_0 such that for all $\delta \geq \delta_0$ and all $\gamma \geq \delta$, we have

$$\frac{\gamma^2\delta(1-\delta)}{1-\gamma\delta} \geq \frac{1-\varepsilon}{1+\varepsilon},\ \frac{\varepsilon^2}{1+\varepsilon} > \gamma(1-\delta),\ v_c \geq \frac{1}{2-\varepsilon+\frac{1}{\varepsilon}\gamma(1-\delta)},\ \text{and } \varepsilon > \frac{1-\gamma\delta}{\gamma\delta}. \tag{12}$$

Take any menu function $y \geq h$. Find $u_c^* \in \arg\min_{u \in \mathcal{U}} y(u_c)$ and take $y^* := y(u_c^*) \geq \inf h = \frac{1}{2}(1+\varepsilon)$. We consider two sub-cases separately:

Subcase A1: $y^* < 1 - \frac{1}{\varepsilon}\gamma(1-\delta)$. Construct function

$$y'(u_c) = \begin{cases} 1 - \left(1 + \frac{y^*}{1-y^*}\gamma(1-\delta)\right)(1 - y(u_c)), & \text{if } u_c \le u_c^*, \\ (1-\gamma(1-\delta))\, y(u_c), & \text{if } u_c \ge u_c^*. \end{cases}$$

We are going to show that (a) y' is a menu function, (b) $y' \ge (1-\gamma(1-\delta))\, y$ and (c) for each type u,

$$\gamma\delta(1 - d(u_c, y')) \ge 1 - d(u_c, y). \tag{13}$$

(See equation (10).) That shall verify that h has $UB(\gamma, \delta_0)$ property.

Ad (a). We start by checking that function y' is convex. First, notice that function y' is continuous at u_c^*, and, as y, minimized at u_c^*. It follows that for each $a \le u_c^* \le b$, we have $y'(u_c^*) \le \frac{b-u_c^*}{b-a} y'(a) + \frac{u_c^*-a}{b-a} y'(b)$. Further, for any $\alpha < \frac{b-u_c^*}{b-a}$, we have

$$\begin{aligned}
& y'(\alpha a + (1-\alpha) b) \\
\le & \frac{\alpha a + (1-\alpha) b - u_c^*}{b - u_c^*} y'(b) + \frac{b - \alpha a + (1-\alpha) b}{b - u_c^*} y'(u_c^*) \\
\le & \frac{\alpha a + (1-\alpha) b - u_c^*}{b - u_c^*} y'(b) + \frac{b - \alpha a + (1-\alpha) b}{b - u_c^*} \left[\frac{b - u_c^*}{b-a} y'(a) + \frac{u_c^* - a}{b-a} y'(b)\right] \\
= & \left(\frac{\alpha a + (1-\alpha) b - u_c^*}{b - u_c^*} + \frac{b - \alpha a + (1-\alpha) b}{b - u_c^*} \frac{u_c^* - a}{b-a}\right) y'(b) + \frac{b - \alpha a + (1-\alpha) b}{b-a} y'(a) \\
= & \frac{1}{b - u_c^*} \left(\frac{(\alpha a + (1-\alpha) b)(b - u_c^*) - u_c^* a - ba}{b-a}\right) y'(b) + \frac{b - \alpha a + (1-\alpha) b}{b-a} y'(a) \\
= & \frac{\alpha a + (1-\alpha) b - a}{b-a} y'(b) + \frac{b - \alpha a + (1-\alpha) b}{b-a} y'(a).
\end{aligned}$$

An analogous calculation shows the same inequality when $\alpha > \frac{b-u_c^*}{b-a}$. Finally, if $a, b \le u_c^*$ and $a, b \ge u_c^*$, then $y'(\alpha a + (1-\alpha) b) \le \alpha y'(a) + (1-\alpha) y'(b)$ for any $\alpha \in [0,1]$, due to the construction of y' as a piecewise-linear transformation of convex y.

Because y' is convex, $D_u y'$ is closed for each u. We show that it is non-empty. By Lemma 8, it is sufficient to check the non-emptiness for $u_c = 0, 1$.

- For any $l \in D_1 y$, we have $(1-\gamma(1-\delta)) l \in D_1 y$. To see it, observe that the derivative of the convex function is multiplied by a constant $(1-\gamma(1-\delta)) \in (0,1)$ and so-obtained affine function satisfies the payoff restriction.
- Let $l \in D_0 y$. Because $1 \geq y \geq h$ and $h(0) = 1$, it must be that $y(0) = l(0) = 1$ and $l(1) \geq \Delta_0 h(1)$, where $\Delta_0 h$ is the affine function tangent to h at 0. By definition of h, we have $\Delta_0 h(1) = \varepsilon$. Consider affine l' defined by $l'(u_c) = 1 - \left(1 + \frac{y^*}{1-y^*}\gamma(1-\delta)\right)(1 - l(u_c))$. Clearly, l' supports y' at 0. By construction, $l'(0) = 1$, and

  $$1 \geq l'(1) \geq 1 - \left(1 + \frac{y^*}{1-y^*}\gamma(1-\delta)\right)(1-\varepsilon).$$

  Because we work with the case $y^* < 1 - \frac{1}{\varepsilon}\gamma(1-\delta)$, the above is not larger than

  $$\geq 1 - (1 + \varepsilon - \gamma(1-\delta))(1-\varepsilon) \geq 1 - (1+\varepsilon)(1-\varepsilon) \geq 0.$$

Lemma 1 implies that y' is a proper menu function.

Ad (b). For $u_c \leq u_c^*$, we have

$$\begin{aligned}
&1 - \left(1 + \frac{y^*}{1-y^*}\gamma(1-\delta)\right)(1 - y(u_c)) - (1-\gamma(1-\delta)) y(u_c) \\
=&1 - \left(1 + \frac{y^*}{1-y^*}\gamma(1-\delta)\right) + \left(1 + \frac{y^*}{1-y^*}\gamma(1-\delta) - 1 + \gamma(1-\delta)\right) y(u_c) \\
=&\gamma(1-\delta)\left[\frac{y(u) - y^*}{1-y^*}\right] \geq 0.
\end{aligned}$$

The claim is immediate for $u_c \geq u_c^*$.

Ad (c). We have

$$d(u_c, y') = \begin{cases} 1 - \left(1 + \frac{y^*}{1-y^*}\gamma(1-\delta)\right)(1 - d(u_c, y)), & \text{if } u_c \le u_c^*, \\ (1 - \gamma(1-\delta))\, d(u_c, y), & \text{if } u_c \ge u_c^*. \end{cases}$$

For $u_c \le u_c$, we check that

$$\begin{aligned} &\gamma\delta\left(1 + \frac{y^*}{1-y^*}\gamma(1-\delta)\right)(1 - d(u_c, y)) - (1 - d(u_c, y)) \\ &= \left(\frac{y^*}{1-y^*}\gamma^2\delta(1-\delta) - (1-\gamma\delta)\right)(1 - d(u_c, y)) \\ &= (1-\gamma\delta)\left(\frac{y^*}{1-y^*}\frac{\gamma^2\delta(1-\delta)}{1-\gamma\delta} - 1\right)(1 - d(u_c, y)) \ge 0, \end{aligned}$$

where the last inequality follows from the fact that $d(u_c, y) \le 1$, $\frac{y^*}{1-y^*} \ge \frac{1+\varepsilon}{1-\varepsilon}$, and inequality (12).

For $u_c \ge u_c^*$, notice first that $d(u_c, y) \le \min(d(u_c^*), d(1))$ due to the fact that function $d(., y)$ is increasing below v_c and decreasing above v_c. Moreover, $d(u_c^*, y) = y^* \ge \frac{1}{2}(1+\varepsilon)$ and,

$$d(1, y) \ge v_c y(1) \ge v_c (1+\varepsilon)\mathcal{N}_A(1) = \frac{1}{2}(1+\varepsilon).$$

(See (1).) Therefore,

$$\begin{aligned} &\gamma\delta(1 - (1 - \gamma(1-\delta))\, d(u_c, y)) - (1 - d(u_c, y)) = d(u_c, y)(1 - \gamma\delta(1 - \gamma(1-\delta))) - (1-\gamma\delta) \\ &= (1-\gamma\delta)\left(\left(1 + \frac{\gamma^2\delta(1-\delta)}{1-\gamma\delta}\right) d(u_c, y) - 1\right) \ge (1-\gamma\delta)\left(\left(1 + \frac{1-\varepsilon}{1+\varepsilon}\right)\frac{1}{2}(1+\varepsilon) - 1\right) \ge 0. \end{aligned}$$

Subcase A2: $y^* \ge 1 - \frac{1}{\varepsilon}\gamma(1-\delta)$. Construct function $y'(u_c) = y(u_c) - \frac{1-\gamma\delta}{\gamma\delta}(1-\varepsilon)u_c$. As in the subcase A1 above, we are going to establish (a), (b), and (c).

Ad (a) Notice that function y' is convex because it is the sum of convex y and an affine function. We check that sets $D_{u_c}y'$ are non-empty for each u_c. Let $l_{u_c} \in D_{u_c}y$ and define $l'_{u_c}(u) = l_{u_c}(u) - \frac{1-\gamma\delta}{\gamma\delta}(1-\varepsilon)u_c$. Notice that $l'_{u_c}(0) = l_{u_c}(0) \in [0,1]$. Moreover, $l'_{u_c}(1) \leq l_{u_c}(1) \leq 1$ and $l'_{u_c}(1) \geq l'_0(1)$. As in the sub-case A1, we determine that $l_0(1) \geq \Delta_0 h(1) = \varepsilon$. Hence, due to (12), we have

$$l'_{u_c}(1) \geq \varepsilon - \frac{1-\gamma\delta}{\gamma\delta}(1-\varepsilon) \geq \varepsilon - \frac{1-\gamma\delta}{\gamma\delta} > 0.$$

Ad (b). We have

$$\begin{aligned}
&y'(u_c) - (1-\gamma(1-\delta))\,y(u_c) = y(u_c)\,\gamma(1-\delta) - \frac{1-\gamma\delta}{\gamma\delta}(1-\varepsilon)\,u_c \\
\geq &\left(1 - \frac{1}{\varepsilon}\gamma(1-\delta)\right)\gamma(1-\delta) - \frac{1-\gamma\delta}{\gamma\delta}(1-\varepsilon) = \gamma(1-\delta)\left[1 - \frac{1}{\varepsilon}\gamma(1-\delta) - \frac{1-\gamma\delta}{\gamma^2\delta(1-\delta)}(1-\varepsilon)\right] \\
\geq &\gamma(1-\delta)\left[1 - \frac{1}{\varepsilon}\gamma(1-\delta) - \frac{1}{1+\varepsilon}\right] \geq \frac{\gamma(1-\delta)}{\varepsilon}\left[\frac{\varepsilon^2}{1+\varepsilon} - \gamma(1-\delta)\right] \geq 0.
\end{aligned}$$

where we used the fact that $y(u_c) \geq y^* \geq 1 - \frac{1}{\varepsilon}\gamma(1-\delta)$ and inequalities (12).

Ad (c). Because $d(u_c, y)$ initially increases and then decreases, we have $d(u_c, y) \geq \min(d(0,y), d(1,y))$. Notice that $d(1,y) \geq v_c y^*$ and $d(0,y) \geq 1 - v_c(1-\varepsilon)$. (For the first inequality, notice that if $l \in D_1 y$, then it must be that $l(1) \geq y^*$ and $l(0) \geq 0$. For the second inequality, notice that if $l \in D_0 y$, then $l(0) = 1$ and, as in part (a) of this case, $l(1) \geq \Delta_u h = \varepsilon$.) Because of (12), we have

$$\begin{aligned}
&v_c y^* - (1 - v_c(1-\varepsilon)) \geq v_c\left(1 - \frac{1}{\varepsilon}\gamma(1-\delta) + (1-\varepsilon)\right) - 1 \\
\geq &v_c\left(2 - \frac{1}{\varepsilon}\gamma(1-\delta) - \varepsilon\right) - 1 \geq 0.
\end{aligned}$$

Hence, $d(u_c, y) \geq 1 - v_c(1-\varepsilon)$. Because

$$d(u_c, y') = d(u_c, y) - \frac{1-\gamma\delta}{\gamma\delta}(1-\varepsilon)v_c,$$

we have

$$\gamma\delta(1 - d(u_c, y')) - (1 - d(u_c, y)) = \gamma\delta\left(1 - d(u_c, y) + \frac{1-\gamma\delta}{\gamma\delta}(1-\varepsilon)v_c\right) - (1 - d(u_c, y))$$

$$= (1-\gamma\delta)(1-\varepsilon)v_c - (1 - d(u_c, y))(1-\gamma\delta) \geq (1-\gamma\delta)(1-\varepsilon)v_c - v_c(1-\varepsilon)(1-\gamma\delta) \geq 0.$$

C.3.2. *Case B:* $v_c = v_s = \frac{1}{2}$. In such a case, we define $h(u_c) = (1-\varepsilon)\mathcal{N}_A(u_c) + \varepsilon$. Find $\gamma, \delta < 1$ such that

$$\varepsilon > 1 - \gamma\delta, \frac{\gamma^2\delta(1-\delta)}{1-\gamma\delta} \geq \frac{1-\varepsilon}{1+\varepsilon}. \tag{14}$$

Take any menu function $y \geq h$. Then, for each u_c

$$y(u_c) \geq (1-\varepsilon)\frac{1}{2} + \varepsilon = \frac{1}{2}(1+\varepsilon). \tag{15}$$

Let

$$y'(u_c) = 1 - \frac{1}{\gamma\delta}(1 - y(u_c)) = \frac{1}{\gamma\delta}y(u_c) - \frac{1-\gamma\delta}{\gamma\delta}.$$

As in the above cases, we are going to show (a), (b), and (c).

Ad (a). Function y' is trivially convex as the sum of a (scaled-up) convex function y and a constant. We check that sets $D_{u_c}y'$ are non-empty for each u_c. By Lemma 8, it is sufficient to check for $u_c = 0, 1$. Let $l_{u_c} \in D_{u_c}y$ and define $l'_{u_c}(u) = \frac{1}{\gamma\delta}l_{u_c}(u) - \frac{1-\gamma\delta}{\gamma\delta}$. We check that affine l'_{u_c} satisfies the required payoff restriction. Notice that $l'_{u_c}(u) \leq l_{u_c}(u) \leq 1$ for each u, including $u = 0, 1$. Because y is convex, $l_{u_c}(0) \geq l_1(0)$ and $l_{u_c}(1) \leq l_0(1)$. Becuase $1 \geq y \geq h$ and $h(0) = h(1) = 1$, it must be that $y(0) = y(1) = 1$ as well. It follows that $l_0(1), l_1(0) \geq \Delta_0 h(1) = \Delta_1 h(0)$, where $\Delta_x h$

is the affine function tangent to h at x. By definition of h, we have $\Delta_0 h(1) = \varepsilon$. To summarize, $l_{u_c}(0), l_{u_c}(1) \geq \varepsilon$, and

$$1 \geq l'_{u_c}(1) \geq \frac{1}{\gamma\delta}\varepsilon - \frac{1-\gamma\delta}{\gamma\delta} = \frac{1}{\gamma\delta}(\varepsilon - (1-\gamma\delta)) \geq 0,$$

where the last inequality comes from (12). Lemma 1 implies that y' is a proper menu function.

Ad (b). Notice that

$$y'(u_c) - (1-\gamma(1-\delta))\, y(u_c) = \left(\frac{1}{\gamma\delta} - (1-\gamma(1-\delta))\right) y(u_c) - \frac{1-\gamma\delta}{\gamma\delta}$$
$$= \frac{1-\gamma\delta}{\gamma\delta}\left[\left(1 + \frac{\gamma^2\delta(1-\delta)}{1-\gamma\delta}\right) y(u_c) - 1\right] \geq \frac{1-\gamma\delta}{\gamma\delta}\left[\left(1+\frac{1-\varepsilon}{1+\varepsilon}\right)\frac{1}{2}(1+\varepsilon) - 1\right] \geq 0,$$

which employs (15) and (14).

Ad (c). Fix u_c and notice that $d(u_c, y') = \frac{1}{\gamma\delta} d(u_c, y) - \frac{1-\gamma\delta}{\gamma\delta}$. Hence,

$$\gamma\delta(1 - d(u_c, y')) - (1 - d(u_c)) = \gamma\delta\left(1 - \frac{1}{\gamma\delta}d(u_c) + \frac{1-\gamma\delta}{\gamma\delta}\right) - (1 - d(u_c)) = 0.$$

Appendix D. Proofs of lower bound case of Theorem 1

D.1. **Proof of Lemma 6.** Let

$$\eta < \min\left(\frac{5}{11}\frac{1-\gamma}{\gamma}\frac{1-\delta}{\delta}, \frac{1}{2}(1-\gamma)\,\delta\,(1-\delta)^2\right) > 0. \tag{16}$$

Let k_0 be such that Lemma 13 holds for η. Fix $k \geq k_0$, and $\delta \geq \delta_0$. Define functions: for each $u \in \mathcal{U}$ and $\psi \in \Delta\mathcal{U}$,let

- $e_A^{\min}(u) = \lim_{k\to\infty} \inf_{(e_A, e_B) \in E^B(\delta, \mathcal{M}_k, \mu) \text{ for some } \mu} e_A(u)$ be the lowest equilibrium payoff of type u across equilibria for any beliefs,

- $g(u,\psi,y_u) = \Pi_{\text{opt}}\left(e_A^{\min}(u)\,\mathbf{1}_{.=u},\psi\right)$ be Bob's largest possible given beliefs ψ and subject to the constraint that Alice's type u receives at least $e_A^{\min}(u)$.

By Lemma 14, for any ψ and $(f_A, f_B) \in E^B(\delta, \mathcal{M}_k, \psi)$, we have

$$1-\delta \le f_B \le g(u,\psi,y_u), \text{ and } \delta(1-\delta) \le e_A^{\min}(u) \le f_A(u). \tag{17}$$

Suppose that $h(u) > e_A^{\min}(u)$ for some type u. Find a belief μ and equilibrium outcome $(e_A, e_B) \in E^B(\delta, \mathcal{M}_k, \mu)$ such that $e_A(u) < \min\left(e_A^{\min}(u) + \frac{1}{10}\eta, h(u)\right)$. By Lemma 13, there exists a mechanism $m \in \mathcal{M}_k$ such that for each ψ and each $(p_A, p_B) \in E(m,\psi)$,

$$\begin{aligned} p_A(u) &\ge \frac{1}{\gamma\delta} e_A^{\min}(u) - \eta \ge \frac{1}{\gamma\delta}\left(e_A(u) - \frac{1}{10}\eta\right) - \eta \\ &= \frac{1}{\delta} e_A(u) + \frac{1-\gamma}{\gamma}\frac{1}{\delta} e_A(u) - \frac{11}{10}\eta > \frac{1}{\delta} e_A(u), \end{aligned} \tag{18}$$

where the last inequality follows from (16) and (17), and

$$\begin{aligned} p_B &\ge \max_{y' \in \mathcal{Y}: y'(u) \ge \frac{1}{\gamma\delta} e_A^{\min}(u)} \Pi(y',\psi) - \eta = \Pi_{\text{opt}}\left(\frac{1}{\gamma\delta} e_A^{\min}(u)\,\mathbf{1}_{.=u}, \psi\right) - \eta \\ &\ge (1-\gamma(1-\delta))\,\Pi_{\text{opt}}\left(e_A^{\min}(u)\,\mathbf{1}_{.=u}, \mu\right) - \eta \\ &= \delta g(u,\psi,y_u) + (1-\gamma)(1-\delta)\,g(u,\psi,y_u) - \eta > \delta g(u,\psi,y_u), \end{aligned} \tag{19}$$

where the last inequality follows from (16) and (17).

Given an equilibrium with payoffs $(e_A, e_B) \in E^B(\delta, \mathcal{M}_k, \mu)$, consider a deviation by Alice, where she rejects Bob's first period offer (whatever it might be) and proposes mechanism m in the subsequent period. Let ψ be the continuation beliefs and $(f_A, f_B) \in E^B(\delta, \mathcal{M}_k, \psi)$ be the period 3 continuation equilibrium outcome after Bob rejects Alice's offer. By (17), the period 2 present value of rejection is not larger than $\delta f_B \leq \delta g(u, \psi, y_u)$, which, by (19), is strictly smaller than the payoff p_B from accepting m. Hence, in equilibrium, Bob will accept m. However, Alice's discounted period 1 payoff from her deviation to m is, by (18), strictly larger than her equilibrium payoff $e_A(u)$. This contradicts the definition of the equilibrium. The contradiction concludes the proof of the lemma.

D.2. **Proof of Lemma 7.** Choose $\gamma, \delta < 1$ so that

$$\frac{1-\gamma\delta}{\gamma\delta}, 2v_k\frac{1-\gamma\delta}{\gamma\delta} \leq \varepsilon, \frac{\gamma^2\delta(1-\delta)}{1-\gamma\delta} \geq \frac{1-\varepsilon}{1+\varepsilon}. \tag{20}$$

We consider the two cases separately. In each case, we take menu function y, type w st. $y(w) \leq h(w)$ and a belief $\mu \in \Delta\mathcal{U}$,and we define a new function y' such that (a) y' is a menu function, (b) $\gamma\delta y'(w) \geq y(w)$, and (c) $\Pi(y', \mu) \geq (1-\gamma(1-\delta))\Pi(y, \mu)$.

Part 1: $h(u) = (1-\varepsilon)\frac{1}{2}$. Take any w and any menu function such that $y(w) \leq \frac{1}{2}(1-\varepsilon)$. We can assume that $\Pi(y, \mu) \geq 1 - y(w)$; otherwise, we can replace y by $\hat{y}(u) = y(w)$ for each $u \in \mathcal{U}$. Define y': for each u,

$$y'(u) = 1 - (1-\gamma(1-\delta))(1-y(u)) = (1-\gamma(1-\delta))y(u) + \gamma(1-\delta)$$

Ad (a). y' is convex as an affine transformation of a convex function Let $l_u \in D_u y$ and define $l'_u(u') = (1-\gamma(1-\delta))l_u(u') + \gamma(1-\delta)$ for each u'. Because $l_u(0), l_u(1) \in [0,1]$, and because $\gamma(1-\delta) \in (0,1)$, we have $l'_u(0), l'_u(1) \in [0,1]$. It follows that $l'_u \in D_u y'$ and the sets $D_u y'$ are non-empty.

Ad (b) Observe that $\gamma\delta y'(u) - y(u) = \gamma(1-\delta)[1-y(u)] \geq 0$.

Ad (c). We have $d(y',u) = (1-\gamma(1-\delta))\, d(y,u) + \gamma(1-\delta)$. Hence,

$$\begin{aligned}&\Pi(y',\mu) - (1-\gamma(1-\delta))\,\Pi(y,\mu)\\ =&\int \left[1 - ((1-\gamma(1-\delta))\, d(y,u) + \gamma(1-\delta)) - (1-\gamma(1-\delta))(1-d(y,u))\right] d\mu(u) = 0.\end{aligned}$$

Part 2: $h(u) = (1-\varepsilon)\min\left(\frac{1}{2v_k},1\right) u_k$ for some k. Take menu function y and type w st. $y(w) \leq h(w)$. We can assume that

$$\Pi(y,w) \geq 1-(1-\varepsilon)\min\left(\frac{1}{2v_k},1\right) v_k = 1-(1-\varepsilon)\min\left(\frac{1}{2},1\right) \geq \frac{1}{2}(1+\varepsilon);$$

otherwise, we can replace y by $\hat{y}(u) = h(u)$. Define function $y'(u) = y(u) + \frac{1-\gamma\delta}{\gamma\delta} y(w)\, u_k$ for each u.

Ad (a). y' is convex as the sum of convex y and an affine function. We check that sets $D_u y'$ are non-empty. Let $l \in D_u y$ and define $l'(x) = l(x) + \frac{1-\gamma\delta}{\gamma\delta} y(w)\, x_k$. Clearly, $l'(u) = y'(u)$ and $l'(x) \leq y'(x)$ for each $x \in \mathcal{U}$. We check that affine l' satisfies the restriction $l'(\omega^n) \in [0,1]$ for each n. Clearly, $l'(\omega^n) \geq 0$. Notice that $l'(\omega^n) = l(\omega^n) + \frac{1-\gamma\delta}{\gamma\delta} y(w)\, \mathbb{1}_{n=k}$. Thus it is enough to show the claim for $n=k$. But then,

$$l'\left(\omega^k\right) = l\left(\omega^k\right) + \frac{1-\gamma\delta}{\gamma\delta} y(w) \leq y\left(\omega^k\right) + \frac{1-\gamma\delta}{\gamma\delta} \leq 1-\varepsilon + \frac{1-\gamma\delta}{\gamma\delta} \leq 1.$$

Ad (b). Immediate.

Ad (c). Notice that for each belief μ, $\Pi(y',\mu) = \Pi(y,\mu) - \frac{1-\gamma\delta}{\gamma\delta} y(w)\, v_k$. Hence, because $y(w) \leq (1-\varepsilon)\min\left(\frac{1}{2v_k},1\right) w_k \leq \frac{1}{2}(b1-\varepsilon)$ and because of (20),

$$\begin{aligned}&\Pi(y',\mu) - (1-\gamma(1-\delta))\,\Pi(y,\mu) = \gamma(1-\delta)\left(\Pi(y,\mu) - \frac{1-\gamma\delta}{\gamma^2\delta(1-\delta)} y(w)\, v_k\right)\\ \geq&\gamma(1-\delta)\left(\frac{1}{2}(1+\varepsilon) - \frac{1+\varepsilon}{1-\varepsilon}(1-\varepsilon)\min\left(\frac{1}{2},v_k\right) w_k\right) \geq \gamma(1-\delta)\left(\frac{1}{2}(1+\varepsilon) - \frac{1}{2}(1+\varepsilon)\right) = 0.\end{aligned}$$

Appendix E. Proofs of Section 4

This is an Online Appendix to "Bargaining over Heterogeneous Pie with Mechanisms and Incomplete Information." It contains the proofs of results from Section 4.

E.1. **Proof of Proposition 1.** The proof has three parts. The first part develops notation. The second part contains two intermediary steps. The last part constructs an equilibrium and verifies the equilibrium condition.

E.1.1. *Preliminaries.* Let $\eta^k = d(X, X(\mathcal{S}_k))$ be the quality of the approximation of the space of simple offers with $\mathcal{S}_k$. Let $\left(R_A^{j,\delta}(u), R_B^{j,\delta}(u)\right) \in [0,1]^2$ denote the unique complete information subgame perfect equilibrium payoffs of Alice's type u and Bob (i.e., the payoffs in the Rubinstein's game).

Let $\left(x^{j,\delta}\right)$ be the Rubinstein's allocation for type u^*, i.e., the outcome of the complete information game of Bob and Alice's type u with unrestricted (simple) offers $\Gamma^j(\delta, \mathcal{S}, \delta_{u^*})$. Let $\left(x^{j,k,\delta}\right)$ be the Rubinstein's allocation in the restricted game $\Gamma^j(\delta, \mathcal{S}_k, \delta_{u^*})$. Then, $x^{A,k,\delta} \in \arg\max_{x\in\mathcal{S}_k : v(x)\geq \delta v\left(x^{B,k,\delta}\right)} u^*(x)$ and $x^{B,k,\delta} \in \arg\max_{x\in\mathcal{S}_k : u^*(x)\geq \delta u^*\left(x^{A,k,\delta}\right)} v(x)$and $\lim_{k\to\infty} x^{j,k,\delta} = x^{j,\delta}$. It follows that $\lim_k v\left(x^{B,k,\delta}\right) = \mathcal{R}_B^{B,\delta}(u^*) \geq \delta\mathcal{N}_B(u^*)$.

We are going to construct an equilibrium in which Bob accepts any offer from Alice that gives him at least $\delta v\left(x^{B,k,\delta}\right)$. Let $A^{k,\delta} = \left\{x \in \mathcal{S}_k : v(x) \geq \delta v\left(x^{B,k,\delta}\right)\right\}$ be the set of allocations that are acceptable for Bob. For each u, each $\mu \in \mathcal{U}$, let

- $e_A^{A,k,\delta}(u) = \max_{x\in A^{k,\delta}} u(x)$ be type u best payoff among all acceptable allocations,
- $x^{A,k,\delta}(u) = \arg\max_{x\in\mathcal{S}_k : u(x)\geq e_A^{A,k,\delta}(u)} v(x)$ be Bob's optimal allocation among Alice's optimal choices,
- $e_B^{A,k,\delta}(u) = v\left(x^{A,k,\delta}(u)\right)$ be Bob's associated payoff given Alice's type u, and
- $e_B^{A,k,\delta}(\mu) = \int e_B^{A,k,\delta}(u)\, d\mu(u)$ be Bob's expected payoff given beliefs μ.

Each type u of Alice such that $u(x) < \delta e_A^{A,k,\delta}(u)$ will reject Bob's offer x. For each belief μ and each x, let $p^{x,k}(\mu) = \mu\left(\left\{u : u(x) < \delta e_A^{A,k,\delta}(u)\right\}\right)$ be the probability of rejection and $\mu^{x,k} = \mu\left(.|u : u(x) < \delta e_A^{A,k,\delta}(u)\right)$ be the updated belief after offer x is rejected. Let

- $e_B(x) = p^{x,k}(\mu)\,\delta e_B^{A,k,\delta}(\mu^x) + \left(1 - p^{x,k}(\mu)\right) v(x)$ be Bob's expected payoff from making offer x that can be rejected, in which case, he gets continuation payoff $\delta e_B^{A,k,\delta}(\mu^x)$,
- $x^{B,k,\delta}(\mu) = \arg\max_{x\in\mathcal{S}_k} e_B(x)$ be Bob's payoff-maximizing offer, and
- $e_B^{B,k,\delta}(\mu) = \max_{x\in\mathcal{S}_k} e_B(x)$ be the optimal payoff.

E.1.2. *Intermediary steps.*

Lemma 15. *Fix k, δ, $r \in [0,1]$, and $\rho > 0$. Let x^0 be the solution to equations $v(x^0) = r$ and $\frac{x^0_{A,c}}{x^0_{A,s}} = \rho$. Then, for each type u such that $u_c \leq v_c$, $x_0 \in \arg\max_{x:v(x)\geq r, \frac{x_{A,c}}{x_{A,s}}\geq\rho} u(x)$.*

Proof. The result has a simple intuition. Because type u likes chocolate less than Bob, her optimal payoff is achieved when Bob's constrain binds, and the allocation has as little chocolate as possible. □

The next result shows that Bob's optimal payoff is smaller than $v\left(x^{B,k,\delta}\right)$. In particular, when Alice makes an offer with an expected payoff that is not smaller than $\delta v\left(x^{B,k,\delta}\right)$, Bob will prefer to accept such an offer rather than wait for his period to receive $v\left(x^{B,k,\delta}\right)$

Lemma 16. *For sufficiently high k and any belief μ, $e_B^{B,k,\delta}(\mu) \leq v\left(x^{B,k,\delta}\right)$.*

Proof. To shorten the notation, we denote $x^A = x^{A,k,\delta}, x^B = x^{B,k,\delta}$. Notice that $\lim_k \sup_{u\in\mathcal{U}} \left|e_B^{A,k,\delta}(u) - \delta v\left(x^B\right)\right| = 0$. Let k be high enough so that $\sup_{u\in\mathcal{U}} e_B^{A,k,\delta}(u) \leq v\left(x^B\right)$.

We are going to show that, if k is sufficiently large, then, for any offer $x \in \mathcal{S}_k$, $e_B(x) \leq v\left(x^B\right)$. First, suppose that $v(x) \leq v\left(x^B\right)$. Then, by choice of k,

$$e_B(x) \leq p^{x,k}(\mu)\,\delta e_B^{A,k,\delta}(\mu^x) + \left(1 - p^{x,k}(\mu)\right) v(x) \leq v\left(x^B\right).$$

Next, suppose that $v(x) > v\left(x^B\right)$. We are going to show below that, in such a case, $p^{x,k}(\mu) = 1$, or the offer x is rejected μ-almost surely. If so, $e_B(x) \leq \delta e_B^{A,k,\delta}(\mu^x) \leq v\left(x^B\right)$.

Recall that the offer is rejected by type u if $u(x) < \delta e_A^{A,k,\delta}(u) = \delta u\left(x^{A,k,\delta}(u)\right)$. Let $\rho^A = \frac{x^A_{A,c}}{x^A_{A,s}}$ be the ratio of chocolate to strawberry in Alice's portion of allocation x^A.

- Because $u^*(x) < \delta u^*\left(x^A\right)$ (recall that x^B was Bob's largest payoff from an allocation that led to the payoff of at least $\delta u^*\left(x^A\right)$ for type u^*), offer x is rejected by type u^*.
- Take any u such that $u^*_c \leq u_c$ (which implies $u^*_s \geq u_s$) and suppose that x hass less chocolate for Alice than x^A:

$$\rho = \frac{x_{A,c}}{x_{A,s}} \leq \rho^A. \tag{21}$$

Then, we check that

$$\frac{u(x)}{\delta e_A^{A,k,\delta}(u)} \leq \frac{u(x)}{\delta u(x^A)} \leq \frac{u^*(x)}{\delta u^*(x^A)} < 1.$$

(Indeed, the first inequality comes from the fact that $e_A^{A,k,\delta}(u) \geq u\left(x^A\right)$. The second inequality is a consequence of the fact that u^* likes chocolate less than u: after some algebra, it is equivalent to

$$\frac{u^*_c}{u^*_s}\rho + \frac{u_c}{u_s}\rho^A + \left(\frac{u_c}{u_s}\frac{u^*_c}{u^*_s}\rho\rho^A + 1\right) \geq \frac{u^*_c}{u^*_s}\rho' + \frac{u_c}{u_s}\rho + \left(\frac{u_c}{u_s}\frac{u^*_c}{u^*_s}\rho\rho^A + 1\right).$$

After subtracting the terms in the bracket, we obtain $\frac{u_c^*}{u_s^*}\rho + \frac{u_c}{u_s}\rho^A \geq \frac{u_c^*}{u_s^*}\rho^A + \frac{u_c}{u_s}\rho$ which holds due to the fact that $0 \leq \frac{u_c^*}{u_s^*} < \frac{u_c}{u_s}$.) It follows that type u rejects x.

- From now on, we assume that x is strictly more chocolatey than x^A. Consider the case $u_c \leq v_c$, i.e., Alice likes chocolate less than Bob. Let x^0 be the solution to equations $v(x^0) = v\left(x^B\right)$ and $\frac{x^0_{A,c}}{x^0_{A,s}} = \rho^A$. Then, $v(x) \geq v\left(x^B\right)$ and $\frac{x_{A,c}}{x_{A,s}} > \rho^A$, and, by Lemma 15, we have $u(x) \leq u(x_0)$. However, $u(x) \leq u(x_0) \leq \delta e_A^{A,k,\delta}(u)$ due to the fact that x_0 satisfies (21), and it falls under the previous case.
- Finally, consider the case $u_c \geq v_c$, Alice likes chocolate more than Bob. Then, for all sufficiently high k, $\delta v\left(x^B\right) = v\left(x^A\right) + O\left(\eta^k\right) > \frac{1}{2}$. At the same time, because Alice types receive substantially less than their Rubinstein's payoffs, $\dot{\lim}_{k\to\infty}\frac{\max_{x:v(x)\geq v(x^B)} u(x)}{\max_{x:v(x)\geq v(x^A)} u(x)} < \delta$. It follows that for sufficiently high k, $u(x) < \delta e_A^{A,k,\delta}(u)$ (uniformly across u st. $u_c \geq v_c$).

$\square$

E.1.3. *Equilibrium.* Let Δ^* be the set of beliefs μ such that $u_s^* = \arg\max_{u\in \text{supp}\mu} u_s$. Then, $\delta_{u^*} \in \Delta^*$. For each $\mu \in \Delta^*$, we are going to construct an equilibrium with payoffs $(e_A, e_B) \in E^B(\delta, \mathcal{M}_k, \mu)$ such that $e_B(u) \geq \delta v\left(x^{B,k,\delta}\right)$ and $e_A(u) \leq \max_{x\in\mathcal{S}_k : v(x)\geq \delta v\left(x^{B,k,\delta}\right)} u(x)$.

For any belief $\mu \in \Delta^*$, we construct the following strategies and belief-updating, and we verify that they form an equilibrium using the one-shot deviation strategy:

- We say that a history is *good* if Alice has always proposed in set $A^{k,\delta}$. After a history that is not good, Bob's beliefs are fixed at δ_{u^*}. The continuation behavior is expected to be as in the equilibrium of the complete information game against type u^*:
  - If it is Alice's turn to make an offer, the expected payoffs are $\left(e_A^{A,k,\delta}(.), v\left(x^{A,k,\delta}\right)\right)$.

- If it is Bob's turn to make an offer, the expected payoffs are

$$\left(\left(\max\left(x^{B,k,\delta},\delta e_A^{A,k,\delta}(.)\right)\right)_{u\in\mathcal{U}},v\left(x^{B,k,\delta}\right)\right).$$

- Alice always proposes in $A^{k,\delta}$. Any other offer is rejected. Type u rejects Bob's offer only if she strictly prefers to wait until the next period to receive $\delta e_A^{A,k,\delta}(u)$. (In particular, type u^* accepts the offer). Bob always accepts an offer in $A^{k,\delta}$.
- Clearly, the behavior in the continuation game after a not-good history is an equilibrium.

- Let $\mu(h)$ be a belief after some good history h.
  - If it is Bob's turn, he offers $x^{B,k,\delta}(\mu(h))$.
    - The expected payoffs from offer x are

$$p^{x,k}(\mu)\,\delta e_B^{A,k,\delta}(\mu^x)+\left(1-p^{x,k}(\mu)\right)v\left(x_B^{B,k,\delta}\right).$$

      Hence, the choice of $x^{B,k,\delta}(\mu(h))$ as the maximum is a best response.
  - After Bob's offer x, Alice's type u accepts it only if $u(x)\leq\delta e_B^{A,k,\delta}(u)$. If she rejects the offer, the beliefs are updated to $\mu^x(h)$.
    - The expected payoffs after she accepts are $u(x)$. If she rejects, the payoffs are $\delta e_B^{A,k,\delta}(u)$. Hence, her choice is a best response.
  - If it is Alice's turn, type u makes an offer $x^{A,k,\delta}(u)$.
    - the expected payoff from offer $x\in A^{k,\delta}$ is $u(x)$.
    - the expected payoff from offer $x\notin A^{k,\delta}$ is $\delta u\left(x^{B,k,\delta}\right)<u\left(x^{B,k,\delta}\right)\leq\max_{x\in A^{k,\delta}}u(x)$. Hence, it is a best response for any Alice type to choose (one of) the best outcome(s) in $A^{k,\delta}$.

    * If Alice chooses $x \notin A^{k,\delta}$, the beliefs change to δ_{u^*} (this can be justified in a sequential equilibrium by appropriately chosen sequence of full-support strategies).
  - Bob accepts any offer in $A^{k,\delta}$ and rejects any other offer.
    * The expected payoff from accepting offer $x \in A^{k,\delta}$ is $v(x) \geq \delta v\left(x^{B,k,\delta}\right)$, and the beliefs are potentially updated to $\mu(h,x)$ following Alice's choice. The expected payoff from rejecting the offer is $\delta e_B^{B,k,\delta}(\mu(h,x))$. By Lemma 16, the latter is smaller than the former, and accepting is a best response.
    * The expected payoffs from accepting an offer $x \notin A^{k,\delta}$ are $v(x) < \delta v\left(x^{B,k,\delta}\right)$. Because any such offer leads to beliefs δ_{u^*}, the left-hand side of the inequality is equal to the expected and discounted payoff from waiting until the next period. Thus, rejecting x is a best response.

E.2. **Proof of Theorem 2.** The lower bound on Alice's payoff is a consequence of Lemma 7.

We show the bound on Bob's payoff. For each Alice's type u, 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 $Y(v) = \{x : v(x) \geq v\}$.

The approximation $\mathcal{M}_k \to \mathcal{M}$ ensures that

$$\eta_k := \sup_{Y \in \mathcal{M}_{\text{menu}}} \min_{Y_k \in \mathcal{M}_{\text{menu}} \cap \mathcal{M}_k} d(Y, Y_k) \to 0.$$

Let $Y_k \in \mathcal{M}_k \cap \mathcal{M}_{\text{menu}}$ be the sequences of menus chosen in the minimum part of the above expression.

Let

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

be the lowest equilibrium payoff across all possible beliefs in the game in which Alice makes the first offer. We are going to show that for each $\varepsilon > 0$, there is sufficiently high k_0 so that for all $k \geq k_0$, $e_B^{A,k,\delta} \geq \frac{1}{1+\delta} - \varepsilon$.

In any equilibrium of the game where Alice makes the first offer, Bob's expected payoff is not lower than $e_B^{A,k,\delta}$. Since it is not possible for all Alice types u to receive payoffs larger than $u\left(e_B^{A,k,\delta}\right)$, a positive-measure fraction of them must accept any offer from Bob that is strictly larger than $\delta u\left(e_B^{A,k,\delta}\right)$. But, in a fashion similar to the proof of Lemma 4, we can show that all Alice types should accept any menu with payoffs described by menu function $y(u) > \delta u\left(e_B^{A,k,\delta}\right)$. (If some types reject, then a positive fraction of them would receive tomorrow's payoffs that are lower than $u\left(e_B^{A,k,\delta}\right)$. But then, a rejection would not be a best response.) Due to linearity,

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

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

On the contrary, suppose that $e_B^{A,k,\delta} < \frac{1}{1+\delta} - \varepsilon$. Then, there exists an equilibrium of the game where Alice makes the first offer with Bob's expected payoffs $e_B \leq e_B^{A,k,\delta} + \eta_k$. Consider a deviation where Bob rejects any offer from Alice and, instead, proposes menu $Y_k(x)$. The above paragraph implies that such menu is accepted for sure if

$x \geq \delta e_B^{A,k,\delta} + 1 - \delta - 2\eta_k$. Bob's deviation is profitable if $\delta\left(x - \eta_k\right) \geq e_B \geq e_B^{A,k,\delta} + \eta_k$. The two inequalities can be satisfied simultaneously if $e_B^{A,k,\delta} \leq \frac{\delta}{1+\delta} - 3\frac{1}{1-\delta^2}\eta_k$. Take k_0 such that for all $k \geq k_0$, $\eta_k\left(1-\delta\right)^2 \leq \varepsilon$.

E.3. **Proof of Proposition 2.**

E.3.1. *Preliminary observations.*

Lemma 17. *For each $y \in \mathcal{Y}$ and $\mu \in \Delta\left\{\tau_1, \tau_2\right\}$, if $\Pi\left(y, \mu\right) > \frac{2}{3}$, then there exists i such that $y\left(\tau_i\right) \leq \frac{2}{3}$, and if $\Pi\left(y, \mu\right) > \delta\frac{2}{3}$, then there exists i such that $y\left(\tau_i\right) \leq \frac{1}{\delta}\frac{2}{3}$.*

Proof. Let $b\left(a, u\right) = \max_{x \in X : u(x) \geq a} v\left(x\right)$ be the maximal payoff of Bob given that Alice type y obtains payoff a. Then, $b\left(., u\right)$ is concave and decreasing and $b\left(\frac{2}{3}, \tau_i\right) = \frac{2}{3}$ for each i and $b\left(\frac{2}{3}, \tau\right) = 1 - \frac{4}{9}$. Also, $\Pi\left(y, \mu\right) \leq \sum_u \mu\left(u\right) b\left(y\left(u\right), u\right)$ .The first claim follows from the fact that, if $\mu \in \Delta\left\{\tau_1, \tau_2\right\}$, and $y\left(\tau_i\right) \geq \frac{2}{3}$ for both i, then the above implies that $\Pi\left(y, \mu\right) \leq \frac{2}{3}$. The second claim is analogous. □

Lemma 18. *For each $y \in \mathcal{Y}$, if $y\left(\tau_i\right) \leq \frac{2}{3}$ for $i = 1, 2$, then $y\left(\tau\right) \leq \frac{2}{3}$. Similarly, if $y\left(\tau_i\right) \leq \frac{1}{\delta}\frac{2}{3}$ for $i = 1, 2$, then $y\left(\tau\right) \leq \frac{1}{\delta}\frac{2}{3}$.*

Proof. The claims follow from the fact that, due to the concavity of menu function y, we have $y\left(\tau\right) \leq \frac{1}{2}\left(y\left(\tau_1\right) + y\left(\tau_2\right)\right)$. □

Recall that $E\left(m, .\right) : \Delta\mathcal{U} \rightrightarrows \mathcal{Y} \times [0, 1]$ is the correspondence of equilibrium outcomes $\left(e_A, e_B\right) \in E\left(\mu; m\right)$ of m with initial beliefs μ. Standard arguments show that $E\left(m, .\right)$ is a non-empty-valued, and u.h.c. correspondence. Additionally, because of the public randomization, $E\left(m, \mu\right)$ is convex. Hence, $E\left(m, .\right)$ is a Kakutani correspondence.

The next two results present two binary divisions of the space of all mechanisms.

Lemma 19. *For each mechanism m, there exist $\mu^A(m) \in \Delta\{\tau_1, \tau_2\}$ and $\left(e_A^A(m), e_B^A(m)\right) \in E(m, \mu(m))$ such that either*

(1) $e_B^A(m) \leq \delta\frac{2}{3}$, *or*
(2) $e_A^A(\tau_1; m), e_A^A(\tau; m), e_A^A(\tau_2; m) \leq \frac{1}{\delta}\frac{2}{3}$.

Let $\mathcal{M}_1^A \subseteq \mathcal{M}$ denote the set of mechanisms that satisfy the first condition.

Proof. Take mechanism $m \notin \mathcal{M}_1^A$, and define set $E \subseteq \Delta\{\tau_1, \tau_2\} \times \mathcal{Y} \times [0,1]$ of tuples (μ, e_A, e_B) such that $(e_A, e_B) \in E(m, \mu)$. Set E is compact and connected as a graph of a Kakutani correspondence. Let $P_i \subseteq E$ be the set of all tuples $(\mu, e_A, e_B) \in E$ such that $e_A(\tau_i) \leq \frac{1}{\delta}\frac{2}{3}$. Set P_i is a closed subset of a compact set, and hence is compact. Moreover, by the first part of Lemma 17, $E = P_1 \cup P_2$. Because E is connected, the intersection of the two sets is non-empty. Take $(\mu, e_A, e_B) \in P_1 \cap P_2$. By construction, $e_A(\tau_1), e_A(\tau_2) \leq \frac{1}{\delta}\frac{2}{3}$. Lemma 18 implies that $e(\tau) \leq \frac{1}{\delta}\frac{2}{3}$. □

Lemma 20. *For each mechanism m and each belief $\mu \in \Delta\{\tau_1, \tau, \tau_2\}$, there is an acceptance probability $\alpha^B(m, \mu)$, beliefs $\mu_\alpha^B(m, \mu), \mu_r^B(m, \mu) \in \Delta\{\tau_1, \tau, \tau_2\}$ and payoffs $\left(e_A^B(m, \mu), e_B^B(m, \mu)\right) \in E(m, \mu_\alpha(m, \mu))$ such that $\alpha^B\mu_\alpha^B + \left(1 - \alpha^B\right)\mu_r^B = \mu$, and for each $u \in \{\tau_1, \tau, \tau_2\}$,*

(1) *if $e_A^B(u; m, \mu) > \frac{2}{3}$, then $\left(1 - \alpha^B(m, \mu)\right)\mu_r(u; m, \mu) = 0$,*
(2) *if $e_A^B(u; m, \mu) < \frac{2}{3}$, then $\alpha^B(m, \mu)\mu_\alpha(u; m, \mu) = 0$.*

Proof. Consider a mechanism m' in which, first, Alice chooses whether to play mechanism m or menu Y^B, and second, the chosen mechanism is implemented. Such a mechanism has equilibrium strategies. Denote the probability of choosing $m$as $\alpha^B(m, \mu)$. Let μ_α^B denote the conditional beliefs after choosing m, and let μ_r^B denote the conditional beliefs after choosing Y^B. Finally, let $\left(e_A^B(.), e_B^B\right)$ denote the payoffs in the (sub-) mechanism m. □

E.3.2. *Proof of Proposition 2.* We describe the equilibrium, and beliefs. We need to consider four type of histories h. In each case, we denote the beliefs as $\mu(h)$:

- Alice's turn to make an offer:
  - In equilibrium, Alice offers Y^A. The continuation payoffs are $\frac{1}{\delta}\frac{2}{3}$ for each type $u \in \{t_1, \tau, \tau_2\}$, and $\left(\mu(\tau_1, \tau_2|h)\frac{1}{\delta}\frac{2}{3} + \mu(\tau|h)\left(1 - \frac{1}{\delta}\frac{4}{9}\right)\right)$ for Bob.
  - If Alice offers a mechanism $m \neq Y^A, m \notin \mathcal{M}_1^A$, the beliefs are updated to $\mu(m)$. The continuation payoffs are $e_A^A(.|m) \leq \frac{1}{\delta}\frac{2}{3}$ for all types of Alice.
  - In all other cases, the beliefs are updated to $\mu(m)$ and Alice's continuation payoffs are $\delta\frac{2}{3} < \frac{1}{\delta}\frac{2}{3}$.
  - Hence, offeringY^A is Alice's best response.
- Bob's turn to accept:
  - If Alice proposed Y^A, Bob accepts resulting in continuation payoffs $\frac{1}{\delta}\frac{2}{3}$ for each Alice type $u \in \{t_1, \tau, \tau_2\}$, and $\left(\mu(\tau_1, \tau_2|h)\frac{1}{\delta}\frac{2}{3} + \mu(\tau|h)\left(1 - \frac{1}{\delta}\frac{4}{9}\right)\right)$ for Bob.
    * If Bob rejects, his discounted continuation payoff is equal to $\mu(\tau_1, \tau_2|h)\frac{2}{3} + \mu(\tau|h)\delta\left(1 - \frac{4}{9}\right)$. Notice that $\delta\left(1 - \frac{4}{9}\right) < 1 - \frac{1}{\delta}\frac{4}{9}$ for sufficiently high δ. Hence, Bob's behavior is a best response.
  - If Alice proposes a mechanism $m \neq Y^A, m \notin \mathcal{M}_1^A$, Bob accepts and the continuation equilibrium has payoffs $e_A^A(.|m)$ for Alice and $e_B^A \geq \delta\frac{2}{3}$ for Bob.
    * If Bob rejects, his discounted continuation payoff in the game with beliefs $\mu(m)$ is equal to $\delta\frac{2}{3}$. Hence, Bob's behavior is a best response.
  - In all other cases, Bob rejects the offer, which leads to a discounted payoff of $\delta\frac{2}{3}$.

        * If Bob accepts m, an arbitrary equilibrium of m is played, which leads to payoff $e_B \leq \delta \frac{2}{3}$. Thus, accepting is a best response.

- Bob's turn to make an offer:
  - In equilibrium, Bob offers Y^B. The continuation payoffs are $\frac{2}{3}$ for each Alice $u \in \{t_1, \tau, \tau_2\}$ and $\left(\mu\left(\tau_1, \tau_2 | h\right) \frac{2}{3} + \mu\left(\tau | h\right)\left(1 - \frac{4}{9}\right)\right)$ for Bob.
  - If Bob offers a mechanism $m \neq Y^B$, Alice's continuation payoffs are $y\left(u\right) = \max\left(\frac{2}{3}, e_A^B\left(u; m, \mu\left(h\right)\right)\right)$ for each $u \in \{\tau_1, \tau, \tau_2\}$ and not more than $\Pi\left(y, \mu\right)$ for Bob. By the first part of Lemma 17, $\Pi\left(y, \mu\right) \leq \frac{2}{3}$.
  - Hence, offering Y^B is Bob's best response.
- Alice's turn to accept:
  - If Bob proposes Y^B, Alice accepts and the continuation payoffs are $\frac{2}{3}$ for each type $u \in \{t_1, \tau, \tau_2\}$ and $\left(\mu\left(\tau_1, \tau_2 | h\right) \frac{2}{3} + \mu\left(\tau | h\right)\left(1 - \frac{4}{9}\right)\right)$ for Bob.
    * If Alice rejects, the discounted continuation payoffs are $\frac{2}{3}$ for each of her type. Hence, accepting is a best response.
  - If Bob proposes a mechanism $m \neq Y^B$, each Alice type decides whether to accept or reject, based on which choice maximizes her payoffs. If she accepts, beliefs are updated to $\mu\left(m, \mu\left(h\right)\right)$ and an equilibrium of m with payoffs $\left(e_A^B\left(m, \mu\left(h\right)\right), e_B^B\left(m, \mu\left(h\right)\right)\right)$ is played. If she rejects, her discounted continuation payoffs are $\frac{2}{3}$ for each type. Thus, her best response payoffs are equal to $\max\left(\frac{2}{3}, e_A^B\left(u; m, \mu\left(h\right)\right)\right)$.