

# The Efficiency of Negotiations with Uncertainty and Multi-Dimensional Deals\*

Matthew O. Jackson, Hugo F. Sonnenschein, Yiqing Xing, Christis G. Tombazos and Omar Al-Ubaydli<sup>†</sup>

Draft: February 2020

## Abstract

We develop a theory of negotiation in which deals have multiple dimensions that can be bundled together. We use theory and experiments to show that in such settings efficient trade is possible even with substantial asymmetric information. The benefits of identifying areas of mutual gain guide agents away from posturing and manipulating their share of the pie, and instead to finding the deals that maximize their mutual gains from trade.

**Keywords:** Negotiation, Bargaining, Exchange, Trade, Multiple Items, Linking, Contract Theory, Mechanism Design, Implementation, Efficiency, Sequential Equilibrium, Perfect Bayesian Equilibrium, Trembles.

**JEL Classification Numbers:** A13, C72, C78, D47, D82.

---

\*The theory part of this paper subsumes and extends results from the working paper “A Theory of Negotiation” by Jackson, Sonnenschein, and Xing (2015), and the experiments combine new treatments together with treatments that appeared in the working paper “Negotiation as the Art of the Deal” by Jackson, Sonnenschein, Xing, Tombazos, and Al-Ubaydli (2017). This paper replaces both of the previous papers.

<sup>†</sup>Jackson: Stanford University, CIFAR, Santa Fe Institute; Sonnenschein: University of Chicago; Xing: Johns Hopkins University; Tombazos: Monash University; Al-Ubaydli: Bahrain Center for Strategic, International, and Energy Studies (Derasat), George Mason University, and King Fahd University of Petroleum and Minerals. Financial support for the experiments from Monash University is gratefully acknowledged. We thank seminar participants at Chicago, UCCEMA, Keio, Monash, Oxford, Peking, Penn State, Princeton, Stanford, Tokyo, and Yale, as well as the Decentralization conference and the BEET workshop, for helpful comments and suggestions.

# 1 Introduction

The efficient exchange of goods and services is the cornerstone of economics, and, as such, the study of the design of mechanisms and markets has been a large and successful area of investigation over the last several decades. The respective design literatures have outlined when efficient outcomes are possible, and how interactions can be structured to lead to the best possible outcomes when asymmetries in information preclude fully efficient trade. However, many economic interactions and transactions occur outside of designed markets and are the result of relatively unstructured “negotiations”, more in line with what Leo Hurwicz (1972) referred to as the “natural game form”. When are the outcomes obtained by agents who interact in a free form manner as good as the best that can be obtained with a designed mechanism? Despite the importance of this question, it has not received the attention that it warrants and economists have only provided answers to different questions.<sup>1</sup>

In this paper we provide an answer in a context that covers many applications: that in which a set of goods or services is to be traded between two agents. These could involve a number of different goods, or a transaction that involves multiple dimensions. For instance, an agreement between a firm and a union would typically include: salaries, responsibilities, profit-sharing, work-hours, a pension, a medical plan, vacation-hours, safety-rules, seniority, promotion-schedules, etc. Treaties between countries can involve territorial agreements, trade concessions, mutual investments, military agreements, time limits, and so forth.

It is known that if two agents need to reach a multi-dimensional agreement that involves enough independent dimensions, then a mechanism designer, who knows a lot about the setting, is able to design a mechanism that constrains what agents can announce, and carefully links actions and outcomes across dimensions and produces approximately efficient outcomes (e.g., see Jackson and Sonnenschein 2007). We examine whether the same is true without the mechanism designer or any constraints on what agents can claim about their preferences or which offers they can make to one another.

We provide a positive answer and show that in a wide variety of settings with significant asymmetric information, natural forms of negotiation lead agents to efficient outcomes in all equilibria. Although neither agent alone has the knowledge needed to determine which deals are efficient, the availability of rich enough sets of offers guides agents to find the efficient allocation, without the need of forced mediation. Since payoffs depend on the overall outcome, it is natural that agents would consider offers and counteroffers that involve overall agreements as opposed to piecemeal offers that provide separate terms of trade for

---

<sup>1</sup>There are some examples of research showing that some “detail-free” mechanisms (those not designed based on distributions of beliefs or preferences) lead to efficient outcomes, such as the results of Satterthwaite and Williams (2002) showing that equilibria of double auctions with many buyers and sellers converge to Walrasian equilibria. However, a double auction is still an imposed mechanism where agents are highly constrained in how they interact (they submit bids); and, there is a single well-defined mapping between bids and outcomes. Here we are allowing for more open negotiations whereby agents can propose any packages of deals.

each dimension. This is consequential, since we prove that such overall offers are essential to the efficiency of negotiations. Furthermore, we show that there are a wide variety of situations, which include – but go beyond – the case in which the surplus is known, in which a relatively small set of offers crowd out the other offers and determine the equilibria. As long as agents are not precluded from using these powerful strategies, all equilibria are efficient. These powerful strategies involve (1) placing multiple offers on the table at once, and (2) demanding a fraction of the surplus from exchange. We refer to the availability of such powerful strategies as the “richness” of negotiations.

The bottom line is a theory of efficient negotiation for the case where efficient exchange is possible, and a demonstration that the set of situations in which negotiation does as well as forced mechanisms is significantly larger than might be conjectured based on the existing literature. This may help to explain the breadth of real world situations in which “haggling” appears to do rather well. Finally, we note that the techniques that we use to prove our main results bear more resemblance to arguments used in the bargaining literature than those of the mechanism design literature, thus showing that the same positive conclusions can be argued from very different perspectives.

We also examine how people actually negotiate in a controlled laboratory setting, shedding additional light on the theory. Our experiments illustrate the contrast between bargaining over one dimension and negotiating over several, and show how people pursue strategies of searching for the right deal. The experiments also shed light on the importance of ‘cheap talk’ and the structure of agents’ information.

This expands on the previous literature in several directions. The main models that economists have that offer insights into such negotiations are models of ‘bargaining’: splitting a pie. Bargaining models are fundamental to predicting how people split a known surplus (e.g., Nash 1953 and Rubinstein 1982), as well as to understanding why people may fail to reach an efficient agreement when bargaining over a price at which to potentially trade some object over which they have privately known values (e.g., Myerson and Satterthwaite 1983). However, bargaining models do not offer much insight into the multiple dimensional problems that practitioners consider particularly relevant. The importance of making multiple offers at once, that emphasize different aspects of a possible deal, has been emphasized as a key and productive negotiating strategy by business people and popular texts on negotiations. For instance, Fisher and Ury’s celebrated 1983 book on negotiating, “Getting to Yes,” which is extensively used by practitioners, is much more about how to find and craft the right deals, than about how to agree on a price. It directs readers to “Realize that each side has multiple interests,” “Broaden the options on the table rather than look for a single answer,” and “Search for mutual gain.” Existing bargaining models shed no light on this perceived wisdom, that offering multiple deals and searching for the right one is central to negotiations.

The gap between observed behaviors and what we might expect from bargaining models calls out for a model of negotiations. The results of Myerson and Satterthwaite (1983) suggest substantial inefficiency when two people bargain over a single item. This is consistent with

what Larsen (2014) finds with car sales: between 17 and 24 percent of pairs of bargainers fail to agree on the price of a single car (and thus trade) when there is positive surplus, which results in a loss of 12 to 23 percent of the overall available surplus. Interestingly, however, there is relatively little empirical evidence that such inefficiency extends to more complex multi-dimensional negotiations, such as labor contracts. In fact, between 1948 and 2005 “idleness due to strikes in the United States never exceeded one half of one percent of total working days in any year” (Kennan 2005). Since 1990 average lost time has been about twenty minutes per year per worker in the U.S.; and even in a more strike prone country, such as Spain, the number is less than 1/3 of a day per worker per year (again, Kennan 2005). This suggests that the efficiency loss associated from reaching an agreement with asymmetric information, when there are multiple dimensions at stake, is often avoided even without any forced mechanisms. It is important to provide an explanation.

Our model sheds light on this phenomenon by showing how negotiating over multiple dimensions provides strong incentives for people to reach efficient outcomes, while bargaining over a single object for trade can lead to inefficiency. In doing this, the theory offers insights behind the above quotes from practitioners and ideas of making multiple simultaneous offers.

There are three key aspects to our theory.

The first is that negotiations involve *multiple dimensions*, rather than just one. There is a deal to be crafted rather than just one good to be traded.

The second is that asymmetric information between the two negotiating parties is primarily about *which deals maximize their gains from trade*, rather than whether there exist gains from trade. Without knowing that there are gains from trade, added dimensions would just complicate the inefficiency. As dimensions expand, so do opportunities to find mutual gains from trade, and agents become increasingly certain that there are gains from trade, and increasingly confident of the size of such gains.<sup>2</sup> As a motivating example, consider a growing number of items for trade, each with independently realized preferences, so that although it is not clear which items should be traded, or how much the transaction of each item is worth, there is increasingly precise knowledge of the total potential gains from trade. The incentives of the agents turn from posturing to gain some price benefit on a particular item, to finding which items to trade in order to realize the maximal total gains.

The third aspect of our theory regards how people choose to negotiate. One might conjecture that the simple expansion of opportunities and known gains from trade make it obvious that outcomes will be efficient. An important preliminary result is that this intuition is wrong. We provide a simple example that strongly makes this point. There are four goods and it is common knowledge what the overall gains from the trade are. This example illustrates very starkly the contrast between ‘bargaining’ and ‘negotiating’. Here we show that if agents *bargain* over the goods, so that they offer and counter-offer prices for each of the goods in parallel, then despite the common knowledge of the gains from trade,

---

<sup>2</sup>Whether approximate gains from trade are known in labor-management negotiations is an empirical question. In Section 2.2 we comment on situations in which the gains from trade are known.

*all* equilibria are inefficient. In contrast, if agents negotiate, so that they can propose overall deals, then *all* equilibria are efficient. The distinction between negotiating and bargaining is very consequential.

A main insight that emerges from our analysis is that if people have multiple items or dimensions involved in a negotiation, and understand the richness of the set of deals that they can propose, then their incentives become better aligned. Equilibrium strategies involve agents searching for what an efficient deal should look like. They take advantage of the richness of the strategies that are available to them to discover which items should be prioritized, and they become less focused on posturing to extort value. It becomes in their interest to share information and find the right deal.

This previews our more general theoretical results. To understand why negotiation works, it is useful to distinguish between two varieties of asymmetric information. The first is knowledge of the overall possible gains from trade. The second is knowledge of the particular deals that realize these gains. Even with knowledge of the overall gains from trade, asymmetric information usually means that neither agent alone knows which deals are efficient. They must negotiate to find those deals. An essential insight behind our main results is that the knowledge that there exist efficient deals crowds out inefficient deals. Offering a deal that is inefficient is dominated by offering a set of deals that includes all the ones that could be efficient - even if one is not sure which one it is - since that leads to a higher total surplus to be split and the proposer can ask for a bit more than she was asking for with the inefficient deal in every one of the potentially efficient deals offered. Our result applies to a broad class of negotiation problems and ways in which agents communicate, but this intuition is at its heart.

In addition to our theoretical results, our experiments reiterate this contrast between bargaining and negotiations. Not only do we find predicted inefficiency in bargaining and efficiency in negotiations, but the ways in which agents negotiate and find their mutual gains of trade line up remarkably well with our theory.

Our paper proceeds as follows. In order to fix some main ideas, Section 2.1 presents an example of a negotiation problem over multiple goods that, as described above, illustrates the multi-dimensional negotiation setting and shows how it is the combination of knowledge of gains of trade, and understanding of the ability to offer a sufficiently rich set of deals that leads to efficiency; and that knowledge of the surplus alone with simple dimension-by-dimension bargaining leads to inefficiency.

We then move to our main theoretical analysis. In Section 2.2 we focus on the case of known surplus. This is a case in which agents know the value of the total utility maximizing agreement, even though they do not know which agreement it is. Here we show that if the manner in which agents negotiate is “rich”, then all equilibria result in fully efficient outcomes. Agents are able to propose overall deals and demand fractions of the overall surplus. Knowing the total surplus allows agents to negotiate over the total, and any misrepresentation of their private information can only lead to a reduction in that total surplus. This

aligns incentives and helps the agents find the right agreement quickly and efficiently.

In Section 4.1 we extend the discussions to the case in which the surplus is only approximately known. This introduces some substantial technical hurdles, since the mapping from games to their sets of sequential equilibria (or perfect Bayesian equilibria) is not upper-hemicontinuous. As we show, slight amounts of uncertainty lead to many equilibria that rely on extreme updating of beliefs (that survive the usual refinements) and are not close to the unique limiting equilibrium outcome. We show that introducing slight trembles eliminates those problems and restores continuity at the limit.<sup>3</sup> These results may be of independent theoretical interest, since it is not widely known that perfect Bayesian equilibria (and sequential equilibria) fail upper-hemi continuity in relatively simple settings. We then move on to show how the results extend to some more general settings with substantial uncertainty regarding the available surplus, but show that these efficiency results require some structure to the uncertainty.

In Section 5 we present results from a series of experiments. These experiments investigate various predictions of the theory. The basis for the experiments is exactly the examples from Section 2.1. Some subjects were put into treatments in which they bargained over individual items, based on the examples; while other subjects were put into treatments in which they negotiated over four items and were able to offer multi-item deals. The experimental results are very much in line with our theoretical analysis. Bargaining on one item is inefficient, exhibiting significant delay, posturing, and failure to trade. In contrast, negotiations over four items together lead to more efficient outcomes. We examine three different levels of uncertainty, all with substantial uncertainty about which deals should be reached, but with varying levels of uncertainty about the overall surplus and possible deals that could be reached.

We also contrast free-form versions of the treatments in which agents can chat and craft arbitrary offers until they reached mutual agreement, with more restricted forms of offers and counter offers. This free-form structure significantly increases the efficiency, and the discussion by many of the subjects is remarkably in line with the logic behind the theory.

## 1.1 Additional Comments on Related Literature

The bundling literature (e.g., Adams and Yellen 1976, McAfee, McMillan, and Whinston 1989, and Bakos and Brynjolfson 1999) has made the point that bundling goods together and selling them in packages can lead to more predictable valuations of a consumer and enhanced efficiency. Although there is some relation to our work in that bundling goods promotes efficiency, there are key differences. In our work, even though it is known that some deal or bundle can lead to an efficient outcome, the structure of that deal is unknown and there is two-sided uncertainty: both parties' information is involved in finding it and

---

<sup>3</sup>This is different from trembling hand perfection as we hold trembles constant and let the games converge, rather than the alternative.

realizing the surplus. The way in which deals are offered and splits of value are requested matters and is critical to our results, as we make precise below.

We emphasize that the game forms that we consider are “universal” in the sense that they are not based in any way on the utility functions or beliefs of the agents. The same negotiations work as the valuations and distribution over those valuations are varied. Thus, they respond to the critique of Wilson (1987), and Satterthwaite, Williams, and Zachariadis (2014), who view mechanisms that are based on details of the preferences and beliefs of the agents as impractical: “[the agents’] beliefs are not a datum that is practically available for defining economic institutions” (p.249).<sup>4</sup>

Beyond the motivating studies mentioned above, there are no experimental studies that examine the questions here, even though there are recent experiments that involve multi-dimensional negotiations. For instance, although Leonardelli et al. (2019) examine a negotiation that involves multiple dimensions between a sports player and a team owner, their setting does not involve uncertainty and so does not admit the issues considered here. New experiments that follow up on our work here, by Bochet, Khanna, and Siegenthaler (2019), consider take-it-or-leave-it bargaining games rather than the protocols in our theory. Such ultimatum negotiations make it impossible for responding agents to share their information and violate the richness condition that we show plays an important role in reaching efficiency.

## 2 Examples of Negotiation Problems and our Model

In order to preview some of the main intuition, we begin with a pair of negotiation problems between a buyer and a seller. These illustrate the theory and also serve as the foundation for our experiments.

### 2.1 Two Examples

*PROBLEM 1* *There is a seller who has cost \$0 or \$160, equally likely, for delivering a single unit of an indivisible object. The buyer has value \$40 or \$200, equally likely, for that object. The costs and values are private information and are determined independently. This is a standard bargaining problem, as in Myerson and Satterthwaite (1983), and the only decision is whether the agents can find a price at which both are willing to trade.*

*A “deal” consists of the transfer of the object from the seller to the buyer and a price paid to the seller. One calculates the profit from trade in the usual manner. Three quarters*

---

<sup>4</sup>This goes beyond some of the concepts of “robust” and “detail-free” that have been used in literature. Those address the more explicit aspect of the Wilson’s critique, namely the assumption of common knowledge among agents; e.g., see Bergemann and Morris (2005) and Roughgarden and Talgam-Cohen (2013). An exception is Matsushima (2008) who used “detail-free” with a meaning more similar to ours, but in an auction environment. See Carroll (2019) for a discussion of the literature. To avoid confusion in terminology, we use “universality” to capture the feature that a protocol/mechanism is not defined based on any knowledge of the prior distribution or utility functions of the agents.

*of the time there exist deals that will benefit both agents, while one quarter of the time there is no mutually beneficial deal.*

**PROBLEM 2** *The seller begins with four indivisible objects, each of which has cost either \$0 or \$40, equally likely, and the buyer has values for these objects that may be either \$10 or \$50, equally likely. The buyer can consume any subset of the goods, including the set of all four goods, and has a value for a set of the goods that is equal to the sum of the values of the goods.*

*Each of the four possible combinations of cost and value, (\$0 \$10); (\$0 \$50); (\$40 \$10) and (\$40 \$50), will appear in one the four goods. However, the order in which the four possible pairings occur is random. To be more precise, each of the four possible pairings of costs and values occurs exactly once, and the twenty-four 4-tuples of costs and values with this property are equally likely. Agents see their own values, but not the values of the other agent. So, the seller will know which two goods cost her \$0 and which two cost \$40, but does not know the buyer's values. Similarly, the buyer knows which two goods are worth \$10 to him and which two are worth \$50 to him, but does not know the seller's costs.*

*A "deal" prescribes the transfer of a subset of the goods from the seller to the buyer and a price paid to the seller for those goods.*

Note that the two problems have identical expected potential (maximal) gains from trade of \$70. In fact, each of the four objects in Problem 2 is exactly a one-fourth scaling of one of the possible outcomes of Problem 1. Essentially, each of the possible outcomes in Problem 1 appears exactly once in Problem 2, scaled by its probability.

Both Problems involve substantial asymmetric information. In the first problem, which is one-dimensional, agents are unsure which of the four possible joint outcomes have occurred. Half of the time, an agent knows that there is a mutually-beneficial deal possible (e.g., a \$0 cost seller is certain that there are positive gains from trade), and the other half of the time an agent is uncertain as to whether such a deal even exists (e.g., a \$160 cost seller anticipates a 50 percent chance that there are gains from trade). Inefficiency is unavoidable under the requirements of (interim) incentive compatibility and individual rationality (the price lies between the two valuations). This follows from a well-known argument and is a finite adaption of the results of Myerson and Satterthwaite (1983).<sup>5</sup> Regardless of how agents bargain in Problem 1, as long as an agent is never forced to accept a deal that gives her a negative utility, the outcome will be inefficient.

The second problem also involves substantial asymmetric information; however, the agents know that there are gains from trade. It is, in fact, common knowledge that there exists a combined deal with a surplus of \$70: the seller transfers the three goods to the buyer for which her cost is less than the buyer's value. The (substantial) uncertainty is about which deal generates that surplus.

---

<sup>5</sup>See Segal and Whinston (2016) for a discussion of extensions of Myerson and Satterthwaite's results.

There is no theory to guide us about this second problem, and so that is developed below. In order to preview our main results, and to set up the experiments, we present results about this example.

### 2.1.1 Inefficiency with Bargaining on Problem 2

We begin by showing that the challenges of reaching efficient outcomes with multiple aspects requires not only having knowledge that there are gains from trade, but also that agents are not overly restricted in their ability to propose deals. This is related to the notion of rich strategy spaces, mentioned above and defined more formally below. To preview the necessity of rich strategy spaces, we show that, in the context of Problem 2, if agents can only negotiate by quoting separate prices for each of the four items, then all equilibria are necessarily inefficient.

For instance, consider the following alternating-offers bargaining protocol. One of the agents quotes separate prices for each of the four items. The other agent can accept any of the offers. If some items remain untraded, then the roles reverse and the other agent offers prices on those items. If some of those prices are accepted, then those items are traded. Agents can continue to alternate making offers on the remaining items, but consume current gains from trade when they are realized. The outcome to this bargaining protocol is inefficient *in all equilibria* for some discount factor ( $\delta < 1$ ). In Appendix B, we prove that all equilibria satisfying a refinement (that rules out fully incorrect beliefs off the equilibrium path) are inefficient regardless of the discount factor.<sup>6</sup>

Myerson and Satterthwaite’s fundamental inefficiency result for exchange with asymmetric information demonstrates that there is a robust set of cases in which efficiency cannot be realized no matter how people bargain over a single item. Here, we show that when considering multiple items, efficiency is possible, but the manner in which agents negotiate is consequential. When bargaining on each item individually, without knowing which one the buyer values more, the seller is willing to delay trade of the low-surplus item to try to screen the buyer and grab more of the surplus on the higher item. Our proof shows that the full gains from trade can only be realized if the agents are not overly restricted in the manner in which they communicate and negotiate.

### 2.1.2 Efficiency with Negotiations on Problem 2

The key to our efficiency results is that agents are able to suggest menus of deals to the other agent, in which they can include all the outcomes that are possibly efficient, and then let the other agent choose among them. The availability of such strategies is formalized in the definition of “richness”.

---

<sup>6</sup>We are not sure whether there exist efficient equilibria for some high discount factors, but know that if they do exist, then they must have beliefs that completely rule out the true state, and are thus quite extreme.

To see how this works, consider Problem 2 in which the seller has realized values of 0, 0, 40, 40 for the four items in order, while the buyer has realized values 10, 50, 50, 10. Although the seller does not know which goods should trade, she knows which trades can possibly lead to efficient outcomes. The seller offers two deals: deal A is to trade items 1, 2, 3 at a combined price of  $\$70/(1 + \delta) + 40$ , and deal B is to trade items 1, 2, 4 at a combined price of  $\$70/(1 + \delta) + 40$ . She could include more deals, but they would be inefficient and so would generate less surplus than the better of these two. Given the realization of the buyer's type, he accepts deal 1 in the first period and the game concludes. There is no way that the agents can earn a higher combined utility, and as we show below, this deal is the outcome in all equilibria of the game.

Observe how the richness of the offers is used in the preceding argument. The seller is able to place two deals on the table, each of which involves the transfer of more than one good. The seller uses her knowledge to identify which deals are potentially efficient, and then the seller allows the buyer to use his knowledge to select the efficient deal. In general, the richness of the set of offers means that the seller knows that she can offer deals that realize the full surplus and then ask for her split of that surplus. Any offers that get less than the total surplus can only lead to lower utility for one or both agents. Any offer that is clearly inefficient is dominated by one (or a set that is sure to include one) that is efficient, which can offer better total utilities for both agents. This is the force that pushes agents to efficiency. The proof is longer than this, of course. It is also more involved, since it covers negotiation games that are not direct offerings of menus of deals, but might have other forms of communication; for example, declaring values on different dimensions, or being free-form. We show that the knowledge of the total surplus, and the ability to communicate in ways that allow agents to find that surplus, aligns incentives.

## 2.2 Multi-Item/Aspect Negotiations: The Model

### 2.2.1 Multiple Aspects and Decisions

A multi-aspect negotiation problem consists of:<sup>7</sup>

- two agents, Alice  $a$  and Bob  $b$ ,
- a finite number,  $n$ , of *items*, initially belong to Alice,
- lists of sets of items that trade at different times:  $N^t \subset \{1, \dots, n\}$ , where items trade at most once so that  $N^t \cap N^s = \emptyset, \forall s \neq t$ ; and

---

<sup>7</sup>For simplicity, our language applies to the case of multi-item exchange; however, the results encompass more general negotiations. For instance, a faculty member may negotiate over a contract which could include whether they teach a particular course, whether they are paid for summer months, whether they get a sabbatical, the number of courses they teach, and so forth, which could be viewed as items. More generally, a contract between a union and a firm, or two firms, or two countries, may include decisions on a number of different provisions, each of which can either stay at the status quo or change, which can then be thought of as trading an item.

- monetary transfers may be made from Bob to Alice and a transfer in period  $t$  is denoted  $p^t \in \mathbb{R}$ .

Each item begins with a default decision that is in place if no agreement is reached on that decision. We allow different items to be traded at different times. For instance, a seller may first sell some rugs to a buyer, then sell some of the rest in a later period. Transfers may also be made in multiple periods.

### 2.2.2 Timing, Uncertainty, and Preferences

Time advances in discrete periods  $t = 0, 1, 2, \dots$

Uncertainty and information about preferences are captured via:

- *finite valuation or type spaces*  $\Theta_i \subset \mathbb{R}$ ,  $i \in \{a, b\}$ , for each individual item,
- a joint type space  $\Theta \subset (\Theta_a)^n \times (\Theta_b)^n$ ,
- a probability distribution  $f$  over types  $\Theta$ , with  $f_i$  denoting the marginal of  $f$  on  $\Theta_i^n$ , and
- a common discount factor  $\delta \in (0, 1)$  that is known to both agents.

In the beginning of period 0, the types are drawn according to  $f$  and agent  $i$  observes  $\theta_i = (\theta_{i1}, \dots, \theta_{ik}, \dots, \theta_{in})$ , with  $\theta_{ik}$  being agent  $i$ 's type for aspect/item  $k$ .

The generality of  $f$  allows for correlated values and also allows for different distributions over various classes of items (say some big, some small).

For now, we assume that the agents' payoffs across items are additively separable, but this is not essential to the analysis (see Section 3.3).

Agents' time-0 utilities from some sequence of trades are

- for Alice:  $U_a = \sum_t \delta^t (p^t - \sum_{k \in N^t} \theta_{ak})$ ;
- for Bob:  $U_b = \sum_t \delta^t (\sum_{k \in N^t} \theta_{bk} - p^t)$ .

Utilities are thus relative to the default decision on each item (i.e., no trade) which is in place if no agreement is reached on that item.

Discounting captures that the seller holds the items, each of which generates a flow payoff in every period up to period  $t$ , when she forgoes the future flow payoffs for those traded items, i.e.  $\theta_{ak}$  is the time- $t$  value of flow payoffs the seller could get from item  $k$ .

The welfare from an item  $k$  is  $\theta_{bk} - \theta_{ak}$  if it is traded, and 0 if not. The social surplus from efficient trade is denoted

$$S(\theta_{ak}, \theta_{bk}) \equiv \max\{\theta_{bk} - \theta_{ak}, 0\}.$$

With an abuse of notation, the surplus from an efficient joint decision over all items is

$$S(\theta_a, \theta_b) \equiv \sum_k \max\{\theta_{bk} - \theta_{ak}, 0\}.$$

We assume that  $S(\theta_{ak}, \theta_{bk}) > 0$  for some  $(\theta_{ak}, \theta_{bk}) \in \Theta_a \times \Theta_b$ , so there are potential gains from trade. In addition, we note that  $S(\theta_a, \theta_b) \geq 0$  for every  $(\theta_{ak}, \theta_{bk}) \in \Theta_a \times \Theta_b$ , since the agents can always choose not to trade.

Ex ante, interim, and ex post efficiency require that the items for which  $\theta_{bk} > \theta_{ak}$  trade at  $t = 0$ .

### 2.2.3 A General Definition of Alternating-Offer Negotiations

We now provide a definition of an *alternating-offer negotiation*,  $\Gamma$ , with  $n$  items. This serves as a foundation for defining three reference protocols in 2.2.4, which are distinguished by their spaces of offers and reactions.

We focus on the “alternating-offer” protocols, in which Alice offers at  $t = 0, 2, 4, \dots$  and Bob offers at  $t = 1, 3, 5, \dots$ . The results extend directly for other alternation patterns, except with changes to the expressions for the split of the surplus.<sup>8</sup>

- One of the agents,  $i$  (the offerer), announces from a finite set of possible announcements (‘offers’)  $A_i^0$ , with a generic offer denoted  $a_i^0$ .
- The other agent,  $j$  (the responder), responds by choosing a subset of items to trade  $N^0 \in N(a_i^0)$  (where  $N(a_i^0)$  is a list of possible sets of items that can trade as a function of the offer); we say the responder “rejects” the offer if he chooses the empty set.
- As a function of  $(a_i^0, N^0)$ , a transfer  $p^0$  is made from Bob to Alice.
- If the set of items  $N \setminus N^0$  is non-empty, we start again with the roles of the agents reversed, and one period of discounting ensues.

...

Inductively, at the beginning of period  $t$ , let  $N(h^{t-1}) = N \setminus (\bigcup_{s < t} N^s)$  be the set of remaining items, which could depend on  $h^{t-1} \equiv (a_{i(0)}^0, N^0, \dots, a_{i(t-1)}^{t-1}, N^{t-1})$  – the full history of negotiations through the last period.

- In period  $t$ , agent  $i(t)$  makes offers from a set  $A_{i(t)}(N(h^{t-1}))$  for the items remaining.
- The responder  $j(t)$  responds by choosing which items to trade  $N^t \in N(a_i^t)$  (where  $N(a_i^t)$  is the list of possible sets of items that can be traded as a function of the offer).
- As a function of  $(a_i^t, N^t)$ , a transfer  $p^t$  is made from Bob to Alice.
- This continues as long as there are goods remaining to be traded.

### 2.2.4 Reference Negotiation Protocols

We present several examples of ways in which agents may negotiate that are special cases of the general framework introduced in 2.2.3. These are prominent examples and so we refer to them as “references”.

---

<sup>8</sup>For instance, if the seller makes all of the offers then the seller will get all of the surplus. The pattern of alternation must be either known in advance or random, but not depend on the history of the game.

These reference protocols are differentiated by the spaces of offers and the rules for transfers. They are all ‘universal’ in the sense that the game forms are independent of the type space and distributions over types.

**Item-by-Item Negotiations** This is a formal description of the item-by-item negotiations (or “bargaining”) introduced in 2.1.1, where each item is independently negotiated via Rubinstein-Stahl alternating offer bargaining. In particular:

- In period  $t$  with remaining items  $N(h^{t-1})$ , a feasible offer  $a_i^t : N(h^{t-1}) \rightarrow \mathbb{R}$  specifies a price  $a_i^t(k)$  for each remaining item  $k$ , and all remaining items are considered to be offered for trade (the offerer can offer extreme prices for ones that she does not wish to trade).
- The transfer, when the responder chooses to trade  $N^t \in N(h^{t-1})$ , is  $p^t = \sum_{k \in N^t} a_i^t(k)$ .

As an example, a consumer thinking about buying several rugs, might bargain with a seller on a item-by-item basis, as in this protocol, or for a set of carpets as in the protocols that we introduce next. This includes the case in which different goods can be invested in and consumed at different times.

Note that if the parties are negotiating over a contract with many aspects, then there are applications in which agents cannot consume any until all of the aspects are agreed upon. For instance, an employment contract would have to specify wages, a pension plan, hours, holidays, etc., and employment might not be feasible until all of the aspects are agreed upon. In such cases, if it is impossible to implement different decisions at different times, then item-by-item negotiations are not possible, and only combinatorial and other holistic negotiations can take place, as described below.

**Combinatorial Negotiations** At the other extreme in terms of universal negotiations, instead of negotiating item-by-item, people can negotiate in ways that allow them to “price” all possible subsets of items.

- In period  $t$  with remaining items  $N(h^{t-1})$ , a feasible offer  $a_i^t : 2^{N(h^{t-1})} \rightarrow \mathbb{R}$  specifies a price for every subset of the items that remain, and all possible subsets are considered up for trade.
- The transfer, when the responder chooses to trade  $N^t \in N(h^{t-1})$ , is  $p^t = a_i^t(N^t)$ .

This protocol provides an important theoretical benchmark: it allows for the richest offer space.

The richness in choice space of the combinatorial protocol can be a disadvantage in practice, given the exponential number of deals that are offered. Agents, instead, tend to use “reduced forms”, such as the one we introduce next showing that a much smaller message space suffices to convey the essential information.

## Value-Announcing Negotiations

We now present an intermediate universal form of negotiation, in which agents negotiate in terms of announcing their values for items and demanding a net payoff. This reduces the amount of information that needs to be communicated. For instance, in the case of trading rugs, this reduces the dimension of the initial offer space from  $2^n$  to  $n + 1$ . Agents announce how much they value each rug and then a net gain in utility from the transfer price that they demand. They then allow the other agent to choose which items to trade.

- In period  $t$  with remaining items  $N(h^{t-1})$ , a feasible offer  $a_i^t = \left( (\widehat{\theta}_{ik})_{k \in N(h^{t-1})}, v_i \right) \in (\Theta_i)^{|N(h^{t-1})|} \times V$  announces (not necessarily truthfully) her types  $\widehat{\theta}_{ik}$ 's for the remaining items and demands a payoff of  $v_i \in \mathbb{R}$ .
- The transfer, when the responder chooses to trade  $N^t \in N(h^{t-1})$ , is such that the offerer gets a net payoff equal to  $v_i$  based on her announced types; that is,

$$p^t = \begin{cases} v_i + \sum_{k \in N^t} \widehat{\theta}_{ik}, & \text{if } i = a; \\ \sum_{k \in N^t} \widehat{\theta}_{ik} - v_i, & \text{if } i = b. \end{cases}$$

In practice, one might think of the following. A seller (Alice) claims her costs for the goods and demands an additional margin. The buyer (Bob), if accepting an offer, chooses which goods to buy, at a price equal to the sum of the seller's declared costs, plus the demanded margin.

### 2.2.5 Universality

The protocols introduced in Section 2.2.4 are “universal”, in the sense that the same game forms will result in efficient equilibria across many environments (distributions over types), and are not tailored to the particular setting. In contrast, the “linking mechanisms” in Jackson and Sonnenschein (2007), for instance, restrict the announcements of types and must be changed with the setting in order to reach efficient outcomes.

In most, if not all, applications there may be no one who would know all the relevant statistical details of the setting and also be able to impose a mechanism that the agents must play. Here we find that in a wide variety of situations, including many of those covered by the above linking mechanism, agents can reach efficiency without the need of mediation.

### 2.2.6 Equilibrium

We work with the following variant of (weak) perfect Bayesian equilibrium adapted directly to our setting, because here beliefs can be defined over types instead of nodes in information sets.

At the beginning of any period  $t$  agents share a common history of observed actions  $h^{t-1} \equiv (a_{i(0)}^0, N^0, \dots, a_{i(t-1)}^{t-1}, N^{t-1})$  (and additionally each privately know their types). After

the offerer moves the common history becomes  $(h^{t-1}, a_{i(t)}^t)$ . We denote the set of all possible histories by  $H$ , including  $h^{-1} \equiv \emptyset$  which is the initial node.

A *belief system* for agent  $i$  is a function  $\tilde{f}_i : H \times \Theta_i^n \rightarrow \Delta(\Theta_{-i}^n)$  that maps each history and own type to a distribution over the other agent's type space. In particular,  $\tilde{f}_i(E_{-i} | h, \theta_i)$  denotes  $i$ 's belief over an event (i.e., a collection of the opponent's types)  $E_{-i}$ , conditional on a history  $h$  and the agent's own type  $\theta_i$ . To capture the idea that these beliefs apply to nodes in the game, we require that a belief system only places positive probability on those  $\theta_{-i}$  for which  $f(\theta_i, \theta_{-i}) > 0$ .

Let  $H_i \subset H$  be the set of histories at which agent  $i$  chooses an action. Then, an agent  $i$ 's *strategy*,  $\sigma_i$ , specifies a distribution over the current action space,  $\sigma_i(h, \theta_i) \in \Delta(A_i(h))$ , for each  $(h, \theta_i) \in H_i \times \Theta_i^n$ .

Let  $U_i(\sigma, \tilde{f}_i, h, \theta_i)$  denote  $i$ 's expected utility under the strategies  $\sigma$ , conditional on being of type  $\theta_i$  and history  $h$  given the belief system  $\tilde{f}_i$ .

Beliefs are *consistent* if for each  $i$  and  $\theta_i$  they correspond to a conditional distribution (relative to the common prior  $f$ ) at almost every  $h$  in the support of  $\sigma_{-i}, \sigma_i(\theta_i)$ .<sup>9</sup>

A strategy profile  $\sigma$  satisfies sequential rationality (relative to a belief system  $\tilde{f}$ ) if  $\sigma_i$  maximizes  $U_i(\sigma_i, \sigma_{-i}, \tilde{f}_i, h, \theta_i)$  for each  $i, \theta_i$  in the support of  $f$ , and every  $h \in H$  at which  $i$  chooses an action.

A *weak perfect Bayesian equilibrium* is a profile  $(\sigma_a, \sigma_b, \tilde{f}_a, \tilde{f}_b)$  of a strategy profile and a consistent belief system for which the strategy satisfies sequential rationality.

We work with weak perfect Bayesian equilibria, adapted to continuum games, rather than sequential equilibria, since the latter are difficult to define for games with a continuum of actions. Moreover, given that we are proving results that hold for all equilibria, this strengthens our results.

### 3 Multi-Aspect Negotiations with Commonly Known Surplus

We first focus on the case in which the surplus is commonly known and treat situations with unknown surplus in later sections.

#### 3.1 Known Surplus

A negotiation problem  $(n, \Theta, f)$  (as defined above) has a known total surplus  $\bar{S}$  if

$$\bar{S} = S(\theta_a, \theta_b) \equiv \sum_k \max\{\theta_{bk} - \theta_{ak}, 0\}, \quad \forall (\theta_a, \theta_b) \in \Theta.$$

---

<sup>9</sup>The usual definitions of consistency apply to finite action spaces. Here, we allow for games with a continuum of actions. Conditional probability measures are then defined by Radon-Nikodym derivatives, and are only tied down up to sets of measure 0.

Throughout, we assume that  $\bar{S}$  is positive.

The limiting case in which the overall surplus is commonly known serves as a proxy. There are many justifications for such a known surplus, and we mention one of them. Suppose there are enough items so that the law of large numbers applies. Working at the limit, where the average surplus is known, rather than along the limit, provides clear intuition. Short of the limit, the growing strategy spaces as the number of items gets large make the arguments more complex. That case is handled separately below.

### 3.2 First Efficiency Results

Our first efficiency results, with known surplus, deal with situations in which there is substantial asymmetric information. Specifically, neither agent knows the values of their counterparty for the items, nor which items need to be traded to reach efficiency. Nevertheless, we show that under both the combinatorial and value-announcing negotiations, *all* weak perfect Bayesian equilibria are efficient and lead to the same division of the known surplus. This is a prelude to showing that the same holds for all “rich negotiations,” of which the above negotiations are examples (see section 3.3).

**THEOREM 1** *If a negotiation problem  $(n, \Theta, f)$  has a known surplus  $\bar{S} > 0$ , then in all weak perfect Bayesian equilibria of the combinatorial negotiations and value-announcing negotiations introduced in Sections 2.2.4:*

- *the agreement is reached immediately,*
- *the full surplus is realized, and*
- *agents’ expected net payoffs are uniquely determined. In particular, they are the Rubinstein shares; i.e.,  $\frac{\bar{S}}{1+\delta}$  for Alice, and  $\frac{\delta\bar{S}}{1+\delta}$  for Bob.*

This follows from Theorem 2 below, and all proofs appear in the appendix.

The intuition behind Theorem 1 is as follows. If there were inefficiency on the anticipated equilibrium path, then because the agents know the potential surplus and can make demands for fractions of that total surplus, there is an offer that they each know makes them strictly better off if it is immediately accepted. The existence of such an offer rules out inefficient equilibria. The argument for the precise Rubinstein shares is based on an adaptation of Shaked and Sutton (1984) to accommodate many dimensions as well as asymmetric information regarding valuations.

### 3.3 Rich Negotiations and Efficiency

Combinatorial negotiations and value-announcing negotiations allow the agents to negotiate over all dimensions in an integrated manner, which takes advantage of their knowledge of the total surplus, while item-by-item negotiation does not. A general condition that

captures such integration is now presented, formalizing the richness concept mentioned in the introduction and 2.1.2.

Again, consider a negotiation problem  $(n, \Theta, f)$  with known surplus  $\bar{S}$ .

### Rich Negotiations

An alternating offer negotiation  $\Gamma$  includes a *fraction- $\kappa$  demanding offer* in some period  $t$  for some offerer  $i(t)$ ,  $\theta_{i(t)} \in \Theta_i^n$ ,  $\kappa \in [0, 1]$ , and history  $h^{t-1}$  such that all items are still available ( $N(h^{t-1}) = N$ ), if there exists  $a_{i(t)} \in A_{i(t)}(N)$  such that<sup>10</sup>

- for every non-empty subset of items that the other agent picks,  $a_{j(t)} \in 2^N \setminus \{\emptyset\}$ , the realized payoff for  $i(t)$  in the current period is at least a fraction  $\kappa$  of the remaining surplus, and
- for any  $\theta_{j(t)}$  for which  $f(\theta_{i(t)}, \theta_{j(t)}) > 0$ : there exists an  $a_{j(t)} \in 2^N \setminus \{\emptyset\}$  for which the remaining surplus is realized, and the payoffs in the current period are a fraction  $\kappa$  of the remaining surplus for  $\theta_{i(t)}$ , and  $(1 - \kappa)$  for  $\theta_{j(t)}$ .

An alternating offer negotiation is *rich* if the current offerer  $i(t)$  has a fraction- $\kappa$  demanding offer, for each type  $\theta_{i(t)}$  for which  $f_{i(t)}(\theta_{i(t)}) > 0$  and each  $\kappa \in [0, 1]$ , at every point of the negotiation such that all items are available.

### Efficiency under Rich Negotiations

**THEOREM 2** *If a negotiation problem with  $n$  items has a known surplus  $\bar{S} > 0$  and the alternating offer negotiation  $\Gamma$  is rich, then in all weak perfect Bayesian equilibria:*

- *agreement is reached immediately,*
- *the full surplus is realized, and*
- *the agents' expected payoffs equal their Rubinstein shares; i.e.,  $\frac{\bar{S}}{1+\delta}$  for Alice, and  $\frac{\delta\bar{S}}{1+\delta}$  for Bob.*

Theorem 1 is a corollary to Theorem 2, since both the combinatorial negotiations and the value-announcing negotiations are rich.

In particular, for value-announcing negotiations, at any point of the game and for  $\kappa \in [0, 1]$ , a fraction- $\kappa$  demanding offer is such that the current offerer lists the types truthfully and demands a total net payoff of  $\kappa\bar{S}$ . Such an offer, once accepted, gives the offerer exactly a net payoff of  $\kappa\bar{S}$  regardless of the responder's decisions. It gives the responder  $(1 - \kappa)\bar{S}$  if she chooses her most preferred option. As long as these rather straightforward value-announcing offers are not ruled out, then a negotiation is rich. This observation is the basis for believing

---

<sup>10</sup>Payoffs expressed here are not discounted; i.e., they are evaluated in the current period  $t$ .

that negotiations should be expected to be rich in practice, provided that agents are free to interact rather freely.

The combinatorial negotiation is similarly rich since every offer available in the negotiation from the value-announcing negotiations has an equivalent offer in the combinatorial negotiations.<sup>11</sup> In this sense, combinatorial negotiations have a “richer” message space and hence are also rich. More generally, expansions in the offer space can only enhance richness.

We note that the combinatorial negotiation has the advantage of allowing for rather general payoff structures. In particular, the agents’ utilities can be non-additively separable across items, but quasi-linear in money; i.e. (assuming all trades are made in the same period)

$$U_a^0 = \delta^t(u_a(N^t, \theta_a) + p),$$

$$U_b^0 = \delta^t(u_b(N^t, \theta_b) - p),$$

$$S(\theta_a, \theta_b) = \max_{N^0 \subset N} (u_a(N^0, \theta_a) + u_b(N^0, \theta_b)),$$

where  $\theta_i \in \Theta_i^n$  is agent  $i$ ’s joint type. In such an environment, when the surplus is known, the efficiency is achieved in all equilibria under combinatorial negotiations.

We note that item-by-item negotiations do not include fraction-demanding offers: the offerer’s payoff depends on which items the responder accepts, and yet the offerer cannot request an overall surplus that must be taken as a whole rather than in part.

Finally, our model of negotiation allows agents to continue negotiation over not-yet-traded items and our efficiency results (Theorem 2) are robust to the reopening of trade. In auction theory, for example, reopening can alter agents’ behavior, since they anticipate future interactions (for instance, if they anticipate that an unsold item in an auction might be put up for sale again with a different reserve price). In contrast, even though there is substantial uncertainty and the potential to reopen discussions usually distorts incentives, here all equilibria are efficient provided the negotiation is rich.

## 4 Multi-Item Negotiations with Unknown Surplus

Our analysis so far illustrates that with known surplus, the richness of negotiations leads to efficient equilibria. The ability to bargain over a full bundle enables the known surplus to dominate the screening of particular items and the uncertainty about which items should trade. We now examine the extent to which the results extend when there is uncertainty about the overall surplus. We first examine vanishing uncertainty about the overall surplus, and then examine non-vanishing uncertainty.

---

<sup>11</sup>Any offer (e.g., from the seller)  $(\hat{\theta}_a, v_a)$  in the value-announcing negotiations has an equivalent offer  $p(N^t) = \sum_{k \in N^t} \hat{\theta}_{ak} + v_a, \forall N^t$  in the combinatorial negotiations.

## 4.1 A ‘Nearly-Known’ Surplus

The exact knowledge of the full surplus is an expository device. With large numbers of items and sufficient independence, the law of large numbers ensures that agents will have a good idea of the total surplus possible, but still have substantial uncertainty about which items should trade. Thus, it is useful to verify that there is not a substantial discontinuity between having the total surplus being ‘nearly-known’ versus exactly-known. In addition, given that *all* equilibria are efficient in the limit, it is enough to establish the upper hemi-continuity of the map from negotiation problems to their equilibria.

There are several technical difficulties that must be addressed.

First, incomplete information game theory is still not well-understood in the case of a continua of types and actions, since measurability issues and the issue of updating beliefs conditional on atomless events are not easily overcome (e.g., sequential equilibria are not well defined for such settings, see Myerson & Reny 2015). In order to overcome the measurability issues which would distract us from our questions, we require that the transfers between the agents can only be selected from some arbitrarily large but finite grid, so that the games are discrete.

Second, as is well-known, the freedom of updating beliefs in incomplete information games can sometimes lead to equilibria that are quite unintuitive. Importantly, this also leads to a problem that is not as well-known and is particularly problematic here. Notice that in the previous section we did not impose any restriction on belief updating off the equilibrium path. Once there is uncertainty of the surplus, however, the freedom of off-path belief updating precludes any hope for upper hemi-continuity. We illustrate this point via examples. It is important to note that this is a general problem with incomplete information games and not unique to our setting. In particular, under standard equilibrium notions including sequential equilibria (even when well-defined) and even when using stronger refinements, the upper hemi-continuity of the set of equilibria can fail at the limit (when uncertainty diminishes). Therefore, a new refinement, or restriction on beliefs, is needed.

Third, in order to apply the law of large numbers as support of ‘nearly known’ surplus, it is necessary to work with a large number of items. However, the action spaces of some rich negotiations explode exponentially as the number of items increases. This leads to challenges in characterizing how beliefs evolve in equilibria.

To handle these issues, we work with a fixed number of items, possibly large, with uncertainty of the surplus that converges to full knowledge. We solve the upper-hemi continuity issue by having agents tremble in small but non-vanishing ways, so that beliefs are precisely determined. We introduce a new reference negotiation protocol in which the strategy space satisfies a size restriction, but still allows for the richness of fraction-demanding offers. These approaches, in combination, allow us to bound beliefs and characterize the equilibrium correspondence for a class of negotiation games, and prove that all equilibria are efficient. In the appendix we show that similar results hold in more general negotiation games if one

directly bounds the rate at which beliefs update.

We begin by illustrating the failure of upper hemi-continuity.

#### 4.1.1 Multi-Item Negotiation with Converging Surplus

Consider a sequence of negotiation games indexed by  $m$ . The  $m$ -th economy has  $n_m$  items.

A sequence of negotiation problems with priors  $f^m \in \Delta(\Theta_a^{n_m} \times \Theta_b^{n_m})$  have surpluses *converging to a per-item surplus*  $\bar{s} > 0$ <sup>12</sup> if

$$s^m \equiv \frac{S^m}{n_m} \rightarrow_p \bar{s}, \text{ as } m \rightarrow \infty$$

where  $S^m$  is the random total surplus in the  $m$ -th problem; i.e.,

$$S^m(\theta_a, \theta_b) \equiv \sum_{k \in \{1, \dots, n_m\}} \max\{\theta_{bk} - \theta_{ak}, 0\},$$

and  $s^m$  is the corresponding per-item surplus.

#### 4.1.2 A Challenge: Failure of Upper-Hemicontinuity of Perfect Bayesian and Sequential Equilibria at the Limit of Certainty

We first illustrate the substantial challenge that sequential equilibria fail a fundamental upper hemi-continuity condition. We view this as a shortcoming of the concepts of perfect Bayesian and sequential equilibrium and the current tool-box of game theory. It is not solved by existing refinements.

Games with arbitrarily small uncertainty are very different from their limit, in the sense that some sequences of sequential equilibria of games with vanishing uncertainty have no limit in the set of sequential equilibria (subgame-perfect equilibria) of the limit game that involves certainty. This is the failure of upper hemi-continuity. This occurs because the notion of sequential equilibrium allows for great freedom in off-path beliefs, and as a result too many outcomes can be supported as part of a sequential equilibrium by extreme off-path beliefs.<sup>13</sup> This challenge is not specific to our negotiation games. It applies to many simple games.<sup>14</sup> Thus, there is a fundamental discontinuity between equilibrium concepts with slight amounts of incomplete information and the limit of full information.<sup>15</sup>

---

<sup>12</sup>We use  $S$  to represent the total surplus, and  $s$  for the per-item surplus.

<sup>13</sup>Upper hemi-continuity generally holds for Bayesian equilibrium (e.g., see Jackson, Simon, Swinkels and Zame (2002)), but fails for sequential equilibria and perfect Bayesian equilibria.

<sup>14</sup>The problem that we are pointing out here is endemic: the example still works with perturbations in the payoffs and/or how the small uncertainty is introduced, as it is freedom in specifying beliefs that cause problems, and not exact indifferences (which lead to lower hemi-continuity problems).

<sup>15</sup>This does not contradict the fact that when a sequence of priors, and its limit, are in the interior of the distribution space, the set of sequential equilibria satisfies upper hemi-continuity (Kreps and Wilson (1982), Proposition 2, p.876). Here upper-hemi continuity fails because we are converging to complete information. Given the importance of the complete information case in the theory and (its approximation) in practice, the failure of upper hemi-continuity is important and troubling.

We show that upper hemi-continuity even fails in a simple single-item Rubinstein bargaining with the most basic forms of uncertainty.

Consider a Rubinstein bargaining game with one item and one-sided uncertainty (let  $\delta = 0.8$ ): Bob's value is commonly known as 50, and Alice's cost is either 0 or 40, so that it is commonly known that the agents should always trade immediately to get efficiency. In addition, suppose that Alice's cost has increasing probability on 40, converging to 1 along the sequence. One might conjecture that all sequential equilibria in this game converge to the unique and efficient equilibrium in the limiting complete information bargaining game in which Alice's cost is 40. However, this is not the case.

In particular, in order to have a finite-action game, consider a price grid  $P^5 = \{0, 5, 10, \dots, 45, 50\}$ . The unique subgame perfect equilibrium of the limiting game (i.e., a complete information game with  $\theta_a = 40$  and  $\theta_b = 50$ ) is an immediate trade at a price of 45. Below we show that with arbitrarily small uncertainty, so that  $f_a(40) = 1 - \varepsilon$  for any tiny  $\varepsilon$ , sequential equilibria allow for substantial inefficiency, and a wide range of prices at which the agents trade. We illustrate this point with the following example.

**EXAMPLE 1** *With the parameters given above, there exists a sequential equilibrium with no trade in the first period. In particular, the following occurs on equilibrium path: At  $t = 0$ , both types of the seller offer  $p = 50$  and are rejected. At  $t = 1$ , the buyer offers a  $p = 45$ , which is accepted by both types of the seller.*

The key to supporting this as an equilibrium is as follows. The on-path behavior is supported by the buyer's belief that  $\Pr(\theta_a = 40) = 0$  upon seeing any off-path offer  $p \neq 50$  at  $t = 0$ .<sup>16</sup> Given this belief, the buyer plays as if in a complete information Rubinstein bargaining game, with "0 meets 50". That is, the buyer always offers  $p = 25$ , and rejects any offer with  $p > 25$ . It is then easy to verify that given the buyer's off-path behavior, both types of the seller prefer to stay on path. ■

The possibility of substantial inefficiency that is illustrated in Example 1 does not depend on the discount factor or on the fineness of the grid of transfers. One can still construct sequential equilibria with no trade in the first several periods, and for which the efficiency loss from delay is at least as big as in the above example.

The existence of substantially inefficient equilibria when one introduces uncertainty is related to the multiplicity of sequential equilibria. The multiplicity is in turn due to the dramatic off-path belief updating, which is not sufficiently ruled out by standard refinements.<sup>17</sup> In Example 1, upon seeing a single off-path offer, the buyer believes the seller is of a 0 type

<sup>16</sup>It is direct to verify that this satisfies the consistency conditions of sequential equilibrium, as one can have a sequence of mixed strategies where the 0 types are arbitrarily more likely to play strategies other than 50 compared to the 40 types.

<sup>17</sup>Beyond Kreps and Wilson (1982), see, for instance, Rubinstein (1985), Banks and Sobel (1987), Grossman and Perry (1986), and Cho and Kreps (1987).

with probability 1, completely discarding the prior belief which puts almost all weight on the 40 type. This leads to beliefs that are dramatically different than those that the agents started with, a different continuation of the game, and a different predicted outcome. It is this “explosion” in the set of equilibria, not just in their quantity but also their nature, that accompany the introduction of slight uncertainty that leads to the failure of upper hemicontinuity of the equilibrium correspondence. In the limiting complete information game with no uncertainty, there is no room for belief updating and equilibrium is unique.

In order to extend the efficiency arguments to small amounts of uncertainty regarding the surplus, we need some way of reigning in the belief updating. Our approach is to introduce trembles, which can be small but are not forced to zero. This places all actions on the equilibrium path with some minimal weight for all types, and avoids the need to make ad-hoc restrictions concerning off-path beliefs. Alternative approaches and results are discussed in Appendix B.1.

#### 4.1.3 A Remedy: Approximate Efficiency Results via Trembles

We introduce trembles that bound the rate at which beliefs are updated. For the trembles not to vanish on each possible action, we discretize the action space by introducing price/transfer grids. We now illustrate the role of trembles in a single-item bargaining game.

Consider some grid of transfers, so that  $P^\Delta$  is finite with a grid structure <sup>18</sup>

$$P^\Delta = \{0, \Delta, 2\Delta, \dots, S_{\max}\},$$

in which  $S_{\max} = \max S(\theta_a, \theta_b) = \max \Theta_b - \min \Theta_a$  is the maximum realized surplus on one item.

Consider a variation of Example 1 and some small  $0 < \gamma < 1$  such that at every node in the game, each type of the player who moves at that node places probability at least  $\gamma/|P^\Delta|$  on each possible action and, subject to that constraint, chooses the remaining probability according to a best response under the agent’s beliefs.<sup>19</sup> It is as if an agent best responds with probability  $1 - \gamma$  and then trembles with the remaining probability  $\gamma$ . With trembles, all nodes are reached and so beliefs are completely tied-down by Bayes’ Rule. Thus, we can work with a trembling version of Bayesian equilibrium in which agents’ update beliefs via Bayes’ rule at all nodes.

---

<sup>18</sup>The increment  $\Delta$  can be viewed as the smallest currency unit (e.g., van Damme, Selten, and Winter (1990)), also the grids can be as fine as possible simply by renormalizing  $S_{\max}$ .

<sup>19</sup>Kreps and Wilson (1982) also use trembles when defining sequential equilibria, but consider a sequence of vanishing trembles, so the size of trembles become eventually negligible, whereas we consider a limit theorem where the size of trembles is fixed (although they can be arbitrarily small) and then there is vanishing uncertainty about overall surplus. Our motivation is quite different from the literature on bargaining with “reputational” types, where each agent has some type(s) being fully rational and some being irrational (e.g., Compte and Jehiel 2002, Abreu and Pearce 2007, Wolitzky 2012, and the papers cited therein).

**EXAMPLE 2** Consider  $P^5 = \{0, 5, 10, \dots, 50\}$  and  $\gamma = 0.11$ , so that the probability of trembles to each possible price is  $\gamma/|P^5| = .01$ . If the prior is .999 on some type, then the posterior, after one-period of belief updating, is at least .9 on that type.

The claim in this example follows easily from the bounds on Bayesian updating (see Lemma 1 in the Appendix). In particular,  $1 - \Pr(a_i | \theta'_i) \leq (.01)^{-1} \times .999 \leq .1$ , hence  $\Pr(a_i | \theta'_i) \geq .9$ . ■

The message from the above example is that even very small trembles limit the extent to which posteriors can deviate from the priors. We next show that the introduction of small, but non-vanishing trembles, is enough to restore continuity of the equilibrium correspondence.

**PROPOSITION 1** Consider a single-item alternating-offer (Rubinstein) bargaining game. Given any  $\varepsilon > 0$ , there exist  $\Delta > 0$ ,  $\gamma > 0$ , and  $\bar{f} < 1$ , such that if the prior  $f$  places at least  $\bar{f}$  on a single pair of types  $(\bar{\theta}_a, \bar{\theta}_b)$ , then in all Bayesian equilibria of the game with grid  $P^\Delta$  and  $\gamma$ -trembles, with probability at least  $1 - \varepsilon$ .<sup>20</sup>

- if  $\bar{\theta}_b > \bar{\theta}_a$ , the price offered in the initial period is in  $((1 - \varepsilon)\frac{\delta\theta_a + \theta_b}{1 + \delta}, (1 + \varepsilon)\frac{\delta\theta_a + \theta_b}{1 + \delta})$  and is accepted; that is, if the buyer's value is higher than the seller's cost, then approximately the Rubinstein price is offered in the initial period and is accepted.
- if  $\bar{\theta}_b < \bar{\theta}_a$ , then trade does not occur.

The proof of this proposition is a variation on that of Theorem 3, and so we omit it.

Trembles tie down beliefs and avoid the problems of updating off the path that can drive discontinuities in sequential equilibria. Here, we get a continuity result at the limit (a technique that could also be helpful in other settings, beyond negotiations).

#### 4.1.4 Approximate Efficiency Results with Trembles: Multi-Item Negotiations under a Frequency Negotiation

Next, we illustrate how the near efficiency result with trembles applies to negotiations with multiple items. We first introduce a “frequency negotiation”, which has a “small” strategy space so that as the number of items become larger, the trembles do not shrink too quickly on each possible action. We extend the results to other negotiation protocols in the appendix.

---

<sup>20</sup>One can strengthen the result by introducing the quantifiers in order. Given any  $\varepsilon$  it is clear that we need a fine enough grid and small enough trembles to be sure that an  $\varepsilon$ -approximate efficiency is possible. If the grid is too coarse then the right prices could not be chosen, and if the trembles are too likely then bargaining breakdown because of random behavior becomes too likely. Once these are fine enough so as not to get in the way of efficiency, they can be as small as we like. However, as trembles become smaller, we need to have closer to complete information, and thus  $\bar{f}$  needs to be chosen as a function of  $\gamma$ , so that belief on ‘very unlikely types’ cannot become too large under trembles.

**Frequency Negotiations** Agents first negotiate over frequencies of valuations that then govern their admissible announcements in a second stage game. This fits with many settings in which people first negotiate over “basic terms”, reach a tentative agreement, and then fill in details. In particular, a “negotiation over frequencies” consists of two phases:

Phase 1 (alternating offers of games characterized by frequencies): in each period  $t$

- The offerer,  $i \equiv i(t) \in \{a, b\}$ , quotes a frequency distribution  $\hat{\phi}_i^n \in \Phi_i^n$ ,<sup>21</sup> and a target payoff  $v_i \in V^{n,\Delta}$ .
- The responder  $j(t)$  accepts or rejects.
- If accepted we move to Phase 2.
- If rejected we move to  $t + 1$  (and one period of discounting ensues), in which agents play Phase 1 again with the roles reversed.

Phase 2 (the game is played):

- The offerer in Phase 1 announces  $\hat{\theta}_i \in \Theta_i^n$  that has a frequency distribution  $\hat{\phi}_i^n$  quoted in Phase 1.
- The responder either chooses to trade a subset of items  $N^t \subset N$  or not to trade. The transfer, when  $N^t$  is not empty, is

$$p^t = \begin{cases} v_i + \sum_{k \in N^t} \hat{\theta}_{ik}, & \text{if } i = a; \\ \sum_{k \in N^t} \hat{\theta}_{ik} - v_i, & \text{if } i = b. \end{cases}$$

- The game ends.

The strategy space in the frequency negotiation is the smallest among the reference protocols of negotiations we have introduced: the space of frequencies has a size of less than  $n^{|\Theta_i|-1}$ , which is typically much smaller than the space of types or the space of all possible trades.

The small strategy space makes the handling of beliefs under trembles tractable: as the number of items becomes larger, the size of a tremble to each possible action does not shrink to zero too quickly. Next, we formally introduce trembles.

**Trembles.** Let  $\Delta = 1$  so that the grids of target payoffs become<sup>22</sup>

$$V^{n,1} = \{0, 1, 2, \dots, nS_{\max}\}.$$

<sup>21</sup> $\Phi_i^n$  is the set of possible frequencies (with  $n$  items). For instance, one such example could be  $(\frac{1}{3}, \frac{2}{3})$  which represents “ $\frac{1}{3}$  of 0’s and  $\frac{2}{3}$  of 8’s”. Note that the quoted frequency  $\hat{\phi}_i^n$  may differ from  $i$ ’s true frequency.

<sup>22</sup>The results presented in Theorem 3 generally hold for any finite  $\Delta > 0$ . This is because the gap of per-item transfers is  $\Delta/n$ , which is negligible when  $n$  becomes large enough.

Again, consider trembles in *Phase 1* such that in any period each type of any player trembles with probability  $\gamma$ , spread uniformly across all actions.<sup>23,24</sup>

The following distributional assumptions help put the frequency negotiation to work. For simplicity, suppose agent  $i$ 's ( $i = a, b$ ) valuations  $(\theta_{i1}, \dots, \theta_{ik}, \dots, \theta_{in})$  are i.i.d distributed according to a frequency  $\bar{\phi}_i$  over  $\Theta_i$ , and independent across agents.<sup>25</sup> Without loss of generality, let  $\bar{\phi}_i(\theta_{ik}) > 0, \forall \theta_{ik} \in \Theta_i$  (i.e., defining  $\Theta_i$  to be the support).

In addition, let  $\theta_i^{\max} = \max \Theta_i$  and  $\theta_i^{\min} = \min \Theta_i$  be the extreme types. We assume  $\theta_a^{\min} < \theta_b^{\min} < \theta_a^{\max} < \theta_b^{\max}$  so that there is non-trivial uncertainty and there is non-trivial expected gain from trade.

With such distributions, there is an expected surplus (per item) of

$$\bar{s} = \sum_{\theta_a} \sum_{\theta_b} \bar{\phi}_a(\theta_a) \bar{\phi}_b(\theta_b) \sum_k \max\{\theta_{bk} - \theta_{ak}, 0\}.$$

In this setting of exchangeable distributions it is natural to restrict attentions to *exchangeable* strategies and equilibria: each agent adopts the same strategy (in *Phase 1*) for each of her types that have the same frequency. We can then prove the following approximate efficiency result for the frequency negotiation.

**THEOREM 3** *Consider the frequency negotiations with the above grid and fix any  $\delta < 1$ . Consider a sequence of negotiation problems  $\{n, \Theta_a, \Theta_b, \bar{\phi}\}$ , indexed by the number of items  $n$ , such that in every problem agent  $i$ 's valuations are i.i.d distributed according to  $\bar{\phi}_i$  over  $\Theta_i$ , and independent across agents.*

*For any  $\varepsilon > 0$ , there exists a small enough tremble probability  $\gamma$  and large enough  $\underline{n}$ , such that if  $n > \underline{n}$ :*

1. *There exist exchangeable (weak) perfect Bayesian equilibria, subject to the trembles.*<sup>26</sup>
2. *In any such equilibrium, with probability at least  $1 - \varepsilon$ :*
  - *agreement is reached in the initial period,*
  - *the realized surplus is at least  $(1 - \varepsilon)n\bar{s}$ ; and*

<sup>23</sup>Trembles are not needed in *Phase 2*. Nonetheless, Theorem 3 is robust to adding similar trembles to *Phase 2* as well.

<sup>24</sup>We assume uniformity of the trembles for expository convenience. It suffices that trembles are distributed in a manner such that the relative probability of trembling to any two different actions is bounded above (and hence below).

<sup>25</sup>These assumptions are stronger than needed. All that is needed is that the distribution over types is exchangeable ( $f^m$  remains the same under any permutation of  $\theta$ : if  $\pi$  is a bijection, then  $f^m(\theta^\pi) = f^m(\theta)$  for all  $\theta$ , where  $\theta_k^\pi = \theta_{\pi(k)}$ ); and there is an exponential rate of precision improvements, sufficient conditions for which are stationarity and summable covariance (cf. <https://stat.duke.edu/courses/Fall11/sta205/lec/wk-07.pdf> Section 7.2).

<sup>26</sup>In this game with trembles all nodes are reached by all types with positive probability. Therefore perfect Bayesian equilibria, sequential equilibria, and weak perfect Bayesian equilibria coincide.

- *Expected payoff / ‘Full-Information Rubinstein share’ for each agent lies in  $(1 - \varepsilon, 1 + \varepsilon)$ .*

It is important to note that frequency negotiations are “approximately” rich, subject to the discreteness of the payoff grids. In particular, by quoting her frequency truthfully and demanding some payoff  $v$ , an offerer can “guarantee” herself a payoff of  $v$  if the responder accepts the offer.<sup>27</sup>

Although we do not provide rates of convergence, they are easy to calculate for Theorem 3. The inefficiency per-item is of order  $O(n^{-0.5-\tau})$ ,  $\forall \tau > 0$ ; that is, the inefficiency vanishes at a rate arbitrarily close to the square-root of  $n$ . In particular, inefficiency comes from the following sources: First, there is inefficiency due to the increment in price grids. Although the increment is a constant ( $\Delta = 1$ ), the corresponding per-item inefficiency is of the order  $O(\frac{1}{n})$ . Second, there is inefficiency due to the trembles, which is proportional to the total probability  $\gamma$ , and again is of order  $O(\frac{1}{n})$  per-item. Third, the realized surplus can be different from the limit surplus. Such a difference induces inefficiency that is of order  $O(n^{-0.5-\tau})$ ,  $\forall \tau > 0$ , a rate similar to those provided by standard central limit theorems.<sup>28</sup>

The theorem is stated for exchangeable equilibria. We suspect that the result also holds for non-exchangeable equilibria, but in those cases the second phase of the protocol becomes more difficult to analyze, as now an agent may have a posterior that places more weight on some types with a given frequency than others. We conjecture that a similar result holds when extending to those equilibria, since the rate at which helpful information is gained is bounded by the trembles.

## 4.2 Substantial Uncertainty

Our attention so far has been on the case of (nearly) known overall surplus. Clearly getting approximate efficiency in equilibria with arbitrary unknown surplus is not possible, since that would violate the Myerson-Satterthwaite Theorem. Nonetheless, the result that full efficiency

---

<sup>27</sup>One can see how this happens when agents play the following equilibrium continuation in Phase 2: The responder has a unique strict best reply to trade the subset of items whose valuations exceed the costs at a transfer that exactly delivers the target payoff demanded by the offerer – provided doing so leads to positive payoffs. As for the offerer, if she was truthful in Phase 1 on the quoted frequency, then her strategy in Phase 2 is to truthfully list her valuations. Even if the responder did not play an optimal strategy, by being truthful, the surplus delivered would still have to be at least the target amount.

<sup>28</sup>Fix any rate of trembles  $\gamma > 0$  and pick  $T$  such that the time-0 continuation value after period  $T$  is negligible. The rate at which time- $T$  posterior (of any event) may differ from time-0 prior is of order  $O(n^{|\Theta_i|T})$  due to the trembles, in which  $O(n^{|\Theta_i|})$  captures the size of  $i$ 's action space. Let  $d$  be a distance allowed between the realized and limit surpluses. By standard concentration inequalities, e.g. Hoeffding (1963), the probability that the actual surplus and realized surplus differ by more than  $d$  (under any allowable posterior) is at most  $\alpha^T \sim O(e^{-2d^2n} \times n^{(|\Theta_i|)T})$  which is still of the order  $O(e^{-2d^2n})$  since the exponential term dominates. Finally, the amount of inefficiency,  $\max\{O(d), O(e^{-2d^2n})\}$ , is minimized with an optimal selected distance  $d \sim O(n^{-0.5-\tau})$ . The resulting overall inefficiency is of order  $O(n^{-0.5-\tau})$ . It is not possible to achieve exact square-root rate of convergence because we need to simultaneously control the difference allowed and the likelihood of the tails.

can be obtained in settings far beyond known surplus is true. While a full characterization of all settings for which universal negotiations lead to efficient outcomes is a challenging open question, we provide an example showing when efficiency is and is not feasible, and present some results on sufficient conditions for efficiency in an appendix.

**EXAMPLE 3 (2 goods, known frequencies, unknown surplus)** *There are  $n = 2$  items. Alice’s costs are either  $(0, 40)$  or  $(40, 0)$ , and Bob’s values are either  $(10, 50)$  or  $(50, 10)$ , both equally likely and independent across the agents. The surplus under optimal trade is either 20 or 50, equally likely.*

*Under the combinatorial negotiation, there exists an equilibrium that achieves full efficiency. On path, at  $t = 0$ , Alice with  $(0, 40)$  offers the following two options (symmetrically for Alice with  $(40, 0)$ ):*

- (1) trade the first item only at  $p = \frac{50}{1+\delta}$ ; or*
- (2) trade both items at  $p = 40 + \frac{20}{1+\delta}$ ;*

*(and price the second item only at a very high price, say,  $p = 50$ ).*

*Bob accepts, picking option #1 if he has  $(50, 10)$  and #2 if he has  $(10, 50)$ , game ends.*

In the above example, there is substantial uncertainty about the total surplus and yet efficiency is obtained. The key to efficiency is richness, including the possibility of offering multiple deals and demanding a fraction of the surplus: the offerer is able to provide a rich enough set of options for the responder to find the efficient trade, while guaranteeing herself a fraction of the surplus. On the equilibrium path, Alice, as the initial offerer, offers two options, (1) and (2), both of which, if accepted, guarantee Alice her Rubinstein share which is  $\kappa = \frac{1}{1+\delta}$  of the total realized surplus.

Appendix B.2 generalizes this example: Proposition 2 shows that full efficiency is obtainable under rich negotiations as long as overall frequencies are known and take on two values (for each agent). Once one moves to three or more substantially different values, one needs additional structure beyond known frequencies, which appears to be a challenging problem as we show via Example 5 in that appendix.

In our experiments (below), one of our treatments is a 4-good version of the above example (formally presented as Example 4 in the Appendix). The experimental finding is consistent with the theoretical prediction: despite substantial uncertainty about the total surplus, a high efficiency is obtained by the subjects (see Table 3).

## 5 Experiments

We now examine how people actually negotiate when they are put in some of the key situations that we have analyzed theoretically. There is no “market design”, so that the agents

come together without a third party, who has knowledge of the statistical structure of the problem, and who can force the agents to play a particular game. The experiments provide a controlled exploration of the extent to which unmediated negotiation results in efficient exchange. To make this point most forcefully, we also include free-form versions of treatments in which negotiations are quite open.

## 5.1 Experiment Design

There are two basic facets to our experimental investigation.

The first question is whether the complexities of multidimensional negotiations in terms of both the information structure and potential surplus, and the offers that need to be made in order to achieve efficiency, will be beyond the capabilities of subjects to grasp. This is a direct examination of whether the theory predicts how subjects actually behave: do they fail to reach efficiency when they bargain over a single good and then reach efficient outcomes when they have a rich set of deals available that allows them to negotiate over a set of goods? To address this question we compare two treatments: a single good with unknown values that is based on Problem 1, and four goods with unknown values but a known surplus that is based on Problem 2. Beyond this basic comparison, to better understand how the details of the knowledge of the surplus matters, we also include a four-good information structure that matches that of Proposition 2 that has unknown surplus but known frequencies; as well as a benchmark four-good information structure with i.i.d. draws.

A second question concerns whether improvements in efficiency with multidimensional negotiation are really due to our theory, or whether instead they are coming simply because a richer space of offers also provides richer opportunities for communication. A literature has documented that cheap talk can enhance efficiency in coordination settings (e.g., Charness 2000 and Charness and Dufwenberg 2006). This may mean that simply by enriching the communication structure by having negotiations with four dimensions one gets enhanced efficiency not because of the richness of actions and our theory, but instead because people have richer ways to “talk” to each other. To understand the role of enhanced communication, we supplement the basic alternating offer treatments with ones that we call “free form”, in which the subjects are free to openly chat however they like and they simply have to confirm an agreement if they come to one. By examining how these free-form treatments compare with the more structured alternating offers negotiations, both with one and four items, we are able to separately see the roles of the multiple dimensions from the opportunities for communication.

Thus, altogether, we present six experimental treatments - as listed in Table 1. The comparison between the first two treatments examines the basic theory with structured alternating offers negotiations, comparing one good to four goods with known surplus. The comparison of the third and fourth treatments with the first two treatments helps us see how enhanced opportunities for communication affect the efficiency of negotiations, separately

from moving from one to four goods with known surplus. The fifth and sixth treatments when compared to the fourth treatment, then helps us further understand the role of different levels of knowledge about the surplus.

Table 1: The Six Treatments

Problem	Format	Number of Subjects
1 good - unknown values	Structured	94
4 goods - known surplus	Structured	96
1 good - unknown values	Free-form	88
4 goods - known surplus	Free-form	86
4 goods - unknown surplus, fully independent values	Free-form	82
4 goods - unknown surplus, known frequencies	Free-form	82

A short preview is that the results are very much in line with the theory: single item bargaining leads to substantial inefficiencies while four-item known surplus negotiation leads to very high efficiency. More flexible (“open chat”) negotiations lead to the highest surplus, but having four goods significantly enhances the fraction of surplus realized regardless of the richness of the communication protocol. Having substantial uncertainty in terms of total surplus leads to lower efficiency than known surplus, but we still find nearly efficient trade when there are known frequencies even with substantial uncertainty over the level of surplus.

In terms of details, in the structured treatments agents take turns in making offers, and at most one offer is made in each period. Discounting applies after each period. In the one good structured protocol (Problem 1), the agents make offers via a standard alternating-offer bargaining game. In the four good structured protocol (Problem 2) we extend alternating offers to allow the agents to offer a menu of deals. In particular, an offer consists of a list of deals: which items are to be traded and a total price to be transferred from buyer to seller. For instance, the seller might choose to list say three deals: trade items 1, 3, and 4, for a price of \$50, trade items 2, 3, and 4, for a price of \$44, trade items 1 and 4, for a price of \$65. In the case of four goods, the agent making offers could include as many deals as the agent wished. The buyer could then choose to accept one of these deals or to reject them all. If they were all rejected, then the period ended and it then became the buyer’s turn to make offers. Typically, two deals were offered - both potentially efficient ones. For a precise statement of the rules and screenshots, see the Experimental Supplement (available at <http://www.stanford.edu/~jacksonm/BargainingExperimentalSupplement.pdf>)

Discounting occurred after each period. The stakes shrunk by 10 percent after each alternation. There was a cap on 8 periods, and if they did not agree by then the game ended with no trade.

The free-form treatments were similar, except that there was no alternation. At each instant, each subject could propose an offer that would show up on the counter-party’s

screen, or accept an offer that was currently on the screen from the counter-party; each subject can also make any edit(s) to the current offer as long as it is not yet accepted. In addition, the subjects can chat by typing text in a chat box. Negotiation ended when an offer was accepted. Also, the same as above discounting applied after each “period”. The initial period lasted 40 seconds, and each of the remaining lasted 20 seconds,<sup>29</sup> with the same timing across all treatments. Again, for precise statements of free-form negotiations, see the Experimental Supplement.

## 5.2 Administrative Details

We conducted the experiment at the Monash University Laboratory for Experimental Economics (MonLEE), using z-Tree (Fischbacher 2007). Subjects were predominantly undergraduate students from Monash University, recruited from a database maintained by MonLEE.

We conducted 30 sessions employing a total of 528 subjects. A detailed summary of the sessions is in the Experimental Supplement. Each session lasted approximately 90 minutes.

Subjects were paid an attendance bonus of Australian \$10,<sup>30</sup> in addition to their positive or negative earnings from the experiment. Subjects made \$37.06 on average, ranging from \$3 to \$190, including the attendance bonus.

The experiment design is between-subjects. No subject participated in more than one session, and each session was devoted to one treatment. Thus all subjects in a given session did just one treatment, and were randomly rematched within their session to play that same treatment with a series of different counter-parties. In any given session, subjects experienced 10 rounds (matchings) of negotiating in pairs. The first four matchings were “practice” and the last six rounds were “real”. Earnings were based on an ex post random selection of one of the last six rounds. Our approach allows us to work with a very conservative statistical comparison of behavior across treatments, with standard errors clustered at the session level.

## 5.3 Experimental Results

We begin by comparing the efficiency of outcomes, comparing the one and four-good structured treatments, as well as the corresponding free-form treatments. As we have suggested, the fact that the overall gains are known in Problem 2 with 4 goods, and that agents can negotiate over aspects simultaneously, allow the subjects to find deals that will lead to mutual gain and greater efficiency than in Problem 1. This is the case. As presented in Table 2, regardless of whether we look at the structured or free-form formats, the negotiation (four-good) treatments lead to significantly more efficient outcomes than the bargaining (one-good)

---

<sup>29</sup>We set the timing based on pilots to allow enough time for a subject to construct an offer of a menu of deals and then for his or her partner to respond, but not to allow too much idle time.

<sup>30</sup>\$3 out of the \$10 is guaranteed, so that if a subject made a loss in excess of \$7, that subject would walk away with a show-up fee of \$3.

treatments. The results are significant at above the 98 percent level.

The  $p$ -values are from a most-conservative statistical analysis in which we consider each session as a single observation. We take this extreme caution since rematching of agents could lead to dependent outcomes across pairs of individuals. This gives us, for instance, only four observations of the one good structured treatment. The results are still highly significant since the variance across sessions of the same treatment was very low (see the Experimental Supplement). The results are similar if we do regressions clustered at the session level or if we do a Mann-Whitney test instead of a  $t$ -test (again, see the Experimental Supplement).

Table 2: Efficiency (Percent of Total Surplus Realized)

	1 good	4 goods known surplus	$p$ -value
Structured	67.2%	85.7%	.013
Free-form	73.7%	93.5%	.001
$p$ -value	0.260	0.021	

We see that the free-form (open chat) versions had higher levels of efficiency regardless of whether there is one good or four goods. The direction of the comparisons is consistent with how cheap talk has helped in other settings, for instance, Charness (2000), Valley, Thompson, Gibbons, and Bazerman (2002), Ellingsen and Johannesson (2004ab), Charness and Dufwenberg (2006), and Feltovich and Swierzbinski (2011). Regardless, there is a larger difference between bargaining and negotiation, as with one good the forces that push towards efficiency are still not sufficiently present even with a free-format and open chat, and the four good treatments with structured negotiations still outperform the one good treatment with free-form.

We also see an increasing level of efficiency as we vary the information structure, in the direction that we would expect, as presented in Table 3. Four goods always dominates one good, and with full independence doing significantly worse than known frequencies and known surplus.

Table 3: Efficiency Comparisons across Free-Form Treatments

	Efficiency
1 Good	73.7%
4 Goods, Unknown Surplus Full Independence	82.9%
4 Goods, Unknown Surplus Known Frequencies	90.1%
4 Goods, Known Surplus	93.5%

There are different ways to measure efficiency. Tables 2 and 3 track the total surplus that each pair of subjects realized as a percentage of the total surplus available (so pairs are effectively weighted by how much surplus they could have generated, and pairs who had no

possible surplus are ignored, as negative surplus was not observed in the experiment). The results are even more pronounced when we just track the percent of surplus that each pair loses and then average those percentages across all pairs unweighted by their size of surplus. Those numbers are more pronounced since the pairs who have the most delay in the one good case are those with the least surplus (again, see the Experimental Supplement).

The inefficiencies in the 1 good case come from both delay and failure to ever trade even though there are gains from trade. In contrast, for the 4 good negotiation case, there is almost always eventual trade - especially in the free-form treatments in which almost all pairs traded - inefficiency only comes from delay when agents try to find the right deal. This is illustrated in Table 4. The differences across treatments is again significant, even with a most conservative test that treats each session as a single observation.<sup>31</sup>

Table 4: Efficiency and Fraction Trading

	Efficiency	$p$ -value rel. to 1 good treat.	Fraction Trading	$p$ -value rel. to 1 good treat.
1 good structured	67.2%		87.1%	
1 good free-form	73.7%		86.3%	
4 goods known-surplus struct.	85.7%	.013	97.6%	.041
4 goods independent free-form	82.9%	.043	97.0%	.008
4 goods known-frequencies free.	90.1%	.006	100.0%	.005
4 goods known-surplus free.	93.5%	.001	99.7%	.005

Table 5 provides the breakdown of how the trades vary by period in the free-form treatments. We see earlier and higher trading in the four good negotiations compared to the one-good bargaining.<sup>32</sup>

As discussed above, with the four-good known surplus negotiation setting, agents no longer have incentives to misrepresent their preferences, while in the one-good bargaining setting they do. The experiments shed some light on this question. In the free-form treatment, out of the subjects who make any claim about their values in the chat, we can track whether those claims are true. There is some subjectivity in categorizing when people are declaring a value, but most cases are fairly clear and we describe the precise rules we followed

<sup>31</sup>Again, we are very conservative and treat each session as an observation, and so the entry is the average of session averages and the  $p$ -values are from a t-test across these averages. The  $p$ -values are lower if we work with other less conservative tests - such as regressions with clustered standard errors. The high accuracy in spite of having few observations in each cell again comes from the very low variance across sessions of the same treatment.

<sup>32</sup>The percent not trading is of all pairs, and so that is why, for instance, the 13.3% number for the 1 good free-form treatment is 0.4% different from the 86.3% eventually trading from Table 4. That previous table considers each session as an observation, and then averages across sessions.

Table 5: Percent of Pairs with Positive Surplus Trading by Rounds, Free Form

	In period #								No Trade
	1	2	3	4	5	6	7	8	
1 good	22.1	19.5	11.8	10.8	7.2	3.6	3.1	8.7	13.3%
4 goods Independent	39.4	27.8	13.3	5.8	4.1	2.1	1.7	2.5	3.3%
4 goods Known Frequencies	63.8	19.3	9.1	4.9	1.2	0.0	0.8	0.8	0.0%
4 goods Known Surplus	82.6	10.9	1.9	1.6	1.6	1.2	0.0	0.0	0.4%

in the Experimental Supplement.<sup>33</sup> The results are presented in Table 6. Misrepresentations routinely occur in the 1 good case throughout the game, but rarely in the four good case (only four times and in the first period in the known surplus case).

Table 6: Fraction of Declarations that are Untrue in Free-Form Treatments

Goods	In period #								Avg.
	1	2	3	4	5	6	7	8	
1	97/385	24/85	13/52	14/37	6/24	2/22	1/10	2/15	25.2%
4 Independent	83/431	2/6	3/4	0/1	0/1	n.a.	n.a.	n.a.	19.9%
4 Known Frequencies	8/439	0/4	n.a.	n.a.	0/1	0/1	n.a.	n.a.	1.8%
4 Known Surplus	4/199	0/2	0/1	n.a.	n.a.	n.a.	n.a.	n.a.	2.0%

The independent four good case involves some nontrivial misrepresentation in the first period, but at a lower frequency than in the one good case, and that quickly disappears with four goods while posturing continues throughout all periods in the one good case.

To put the 25 percent misrepresentations in the bargaining treatment in context, note that half of the time, when a buyer has a low value or a seller has a high value then there are no real gains to misrepresentation. Thus, misrepresentations should only occur for half of the types, and equilibrium is in mixed strategies. So, 25.2 percent represents a rate of more than one half in the cases where the agents should be mixing.

The fact that there was only four misrepresentations out of almost two hundred declarations in the known surplus negotiations treatment (and, at least proportionately, not many more in the known frequency case), suggests that the force that aligns subjects incentives and makes the issue about finding the right deal rather than posturing is not subtle: there is no heterogeneity here and so the force is strong and not one that requires high levels of sophistication among the subjects.

<sup>33</sup>For example, we categorize a declaration to be any instance in which a person communicates one of the possible numbers that they might have on an item or items, and a misrepresentation to be when they communicate a number that differs from their actual value. In a number of instances declarations are indirect. As an example, consider the following exchange, from session 11, round 5, group 2, Seller: “What do you have?”, Buyer: “Whats not 200?”, Seller: “40?”, Buyer: “Yeah”. The buyer never says 40 explicitly, but the message seems clear. Further details appear in the Experimental Supplement.

## Examples of how Subjects Bargain and Negotiate

We present some typical examples of how things work in the free-form bargaining and negotiation treatments, as they give an idea of how the experiments worked, and also dovetail with the theory. Of course, they are anecdotal as we have chosen only a few of hundreds of pairings; but these provide insight into Table 6 and the theory to follow.

We start with the free-form one-good bargaining treatment.

Here is an example of a seller with value \$0 and buyer with value \$40.<sup>34</sup> We see the attempt of the seller to get a high price, but the seller eventually gives in and 50 percent of the surplus is lost to discounting.

- Seller: (Offers price of \$180)
- Buyer: (Offers price of \$0)
- Seller: “If you have 200, please accept it. Quickly.”
- Buyer: “I dont”
- Seller: “Split 20 20. What do you have?”
- Buyer : “40”
- Seller : “Make an offer”
- Buyer : (Makes an offer of \$20)
- Seller: (Makes a counteroffer of \$179)
- Buyer: “Why not make a leap of faith”
- Seller: “Accept mine then”
- Buyer: “Trust that Im 40. I cant make negative profit”
- Seller : “Why should I? ha ha”
- Buyer : “Guess we profit 0 then. Your call.”
- Seller : “Ok I trust you if you offer 21”
- Buyer: (Offers a price of \$21)
- Seller: (Accepts in bargain period 6 and profits are subjected to a 50% discount.)

Here is another example of a seller with value 0 and buyer with value \$40. Here, the seller explicitly misrepresents, and in this case the good never trades.

- Buyer: “Hi whats your cost”
- Seller: “Im guessing yours is 40”
- Buyer “Stage 1 lets go. Yeah mine is 40 - hahah good guess”
- Seller: “This time I got 160 - Not joking.”
- Buyer: “Oh no. Lets not do anything then.”
- Seller: (Makes an offer of \$180 and remains firm throughout bargaining periods. Good never traded.)

Here is an example of a seller with value 0 and buyer with value \$200 who both misrepresent. Eventually the seller gives in and they trade in the last period but lose 70 percent of

---

<sup>34</sup>In some cases, subjects sent messages at or near the same time. The ordering is set to make the chats as clear as possible. Chats are edited for clarity and not all chatter is reported in these transcripts. For example in the fifth line where the buyer reports “40” - this is edited from Buyer: “Whats not 200?”, Seller: “40?”, Buyer: “Yeah”.

the surplus:

- Seller: “Im 160, u?”
- Buyer: “40”
- Seller: “Damn”
- Seller: (Offers price \$180)
- Buyer: “Lets be honest [...]”
- Seller: “Whats ur offer”
- Buyer: (Offers price of \$20)
- Seller: “Ofc”
- Buyer: “Because its 40 for me - if you have 0 - then this is even”
- Seller: “Well its 160 for me - so yeah”
- Buyer: “We are in stage 5 - no time to waste - if you have 0 - go for it”
- Seller: “Go for mine”
- Buyer: “I sure would have - but I cannot - stage 8 - go for mine”
- Seller: (Accepts the price of \$20 in period 8 and profits are subjected to a 70% discount)

In contrast, chat and negotiation in the four-item known-surplus case tend to be very short and to the point - either offering the possibly efficient deals or expressing valuations truthfully (as in Table 6) and then reaching an efficient deal in the first period. Here are typical examples.

In the first example, the buyer offers what “he” thinks could be efficient deals and they are quickly accepted.

- Buyer: (Offers the two deals that could be efficient given his information, with prices that split surplus evenly)
- Buyer: “Its half-half. Pick one. Quick”
- Seller: (Accepts the offer that maximises joint surplus. 35 seconds elapse in total)

In the next one, instead of starting with the deals that could be efficient, the agents begin by declaring their private information truthfully and then constructing the efficient deal.

- Seller: “ 0 40 0 40 ”
- Buyer: “My 10s are in 1 n 4 ”
- Buyer: “50/50 ” (Offers a deal that maximises joint surplus and splits it evenly)
- Seller: (Accepts. 21 seconds elapsed in total)

Here is another example.

- Buyer: “10s in 2 and 4”
- Seller: “40s in 1 and 4”
- Buyer: (Offers a deal that maximises joint surplus and splits it evenly)
- Seller: (Accepts. 18 seconds elapsed in total)

**Experiment Summary** To summarize, the efficiency of the four-good negotiation is consistent with the predictions of our theory, and contrasts significantly with the inefficiency of

the one-good bargaining setting. This holds even in the free-form setting, where there is no real structure on the ways in which subjects can negotiate: effectively there is no “mechanism design”. Moreover, the dialogs of the experiments are suggestive of the anatomy of win-win negotiation - subjects quickly exchange information truthfully and move to an efficient deal. In addition, the extent of the surplus in the four-good case is greater with known frequencies or surplus than with fully independent types.

## 6 Concluding Remarks

Although negotiations frequently involve several aspects of a contract or deal, traditional bargaining theory focuses on a situation in which there is a single aspect to be determined. We extend that theory to encompass negotiations, in which deals have many aspects. Our model is descriptive. Agents freely negotiate the terms of a deal with offers and counteroffers, and they do so in the absence of any mediation. Despite the fact that they intend to serve only their own self-interest, we define a robust class of meaningful situations in which outcomes are always socially efficient. This leads to a new perspective, which would appear to have some empirical relevance regarding the costs of asymmetric information. It is a tale about the reach of the invisible hand.

In both structure and technique, our theoretical analysis is an extension of Rubinstein (1982) to allow for deals with multiple aspects and asymmetric information. The new ideas concern the way in which we decompose the knowledge structure when deals are multi-aspect, as well as the manner in which we model strategic possibilities when the interactions between agents are more complex than in bargaining theory. The decomposition of knowledge into two parts: knowledge of the possible gains from trade and knowledge of where these gains are to be found, is demonstrated to be productive. Even when the gains from trade are not approximately known, we establish that the distinction between these two forms of knowledge is useful.

When the gains from trade are known, the manner in which agents negotiate is determined by the presence of powerful strategies, which we argue are available to thoughtful players. These strategies, in a sense, crowd out less efficient ones. They lead the parties to honestly reveal their private information and, when they possess the private information of a counterparty, to use it in a manner that promotes mutual gain. As a consequence, information is shared truthfully and an efficient deal is reached without delay.

Our experiments complement our theoretical treatment. Beyond providing modest tests of the theory, we regard them as an integral part of our analysis, in that the dialogs which we observe provide some considerable comfort regarding the manner in which we have argued the agents negotiate. In particular, the dialogs suggest the relevance of fraction demanding strategies. Also, when the gains from trade are not approximately common knowledge and fraction demanding offers do not exist, the dialogs demonstrate “posturing” and the paths that lead to inefficiency. We believe that the experiments are also noteworthy in the manner

in which they allow for behaviors that are both free-form as well as structured and compare the resulting outcomes.

## References

- Abreu, D., D. Pearce, and E. Stacchetti (1991): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* 58 (5), 1041-1065.
- Adams, W.J., and Yellen, J.L. (1976): "Commodity bundling and the burden of monopoly," *The quarterly journal of economics*, 90 (3), 475-498.
- Armstrong, M. (1999): "Price Discrimination by a Many-Product Firm," *Review of Economic Studies*, 66 (1), 151-168.
- Ausubel, L.M. and R.J. Deneckere. (1989): "Reputation in bargaining and durable goods monopoly." *Econometrica*, 57 (3), 511-531.
- Ausubel, L.M., P. Cramton and R.J. Deneckere. (2002): "Bargaining with incomplete information." *Handbook of game theory with economic applications* 3: 1897-1945.
- Bakos, Y., and Brynjolfsson, E. (1999): "Bundling information goods: Pricing, profits, and efficiency," *Management science*, 45(12), 1613-1630.
- Banks, J.S. and J. Sobel (1987), "Equilibrium Selection in Signaling Games," *Econometrica*, 55 (3), 647-661.
- Bergemann, D. and S. Morris (2005): "Robust mechanism design," *Econometrica*, 73 (6), 1771-1813.
- Billingsley, P. (1968): *Convergence of Probability Measures*, Wiley, New York, New York.
- Charness, G. (2000): "Self-Serving Cheap Talk: A Test of Aumann's Conjecture," *Games and Economic Behavior*, 33(2), 177-94.
- Charness, G. and M. Dufwenberg (2006) "Promises and Partnership," *Econometrica*, 74(6), 1579-601.
- Cohn, Z. (2010): "A Note on Linked Bargaining," *Journal of Mathematical Economics*, 46 (2), 238-247.
- Carroll, G. (2019) "Robustness in Mechanism Design and Contracting," *Annual Review of Economics*, 11, 139-166.
- Dvoretzky, A., Kiefer, J., and Wolfowitz, J. (1956): "Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator," *The Annals of Mathematical Statistics*, 27 (3) 642-669.
- van Damme, E., Selten, R., and Winter, E. (1990): "Alternating bid bargaining with a smallest money unit," *Games and Economic Behavior*, 2 (2), 188-201.
- Fang, H. and P. Norman (2003): "An Efficiency Rationale for Bundling of Public Goods," Cowles Foundation Discussion Paper no 1441.
- Fudenberg, D. and D. Levine (1983): "Subgame-perfect equilibria of finite-and infinite-horizon games," *Journal of Economic Theory*, 31 (2), 251-268.
- Fudenberg, D. and D. Levine (1994): "Efficiency and Observability in Games with Long-Run and Short-Run Players," *Journal of Economic Theory*, 62 (1), 103-135.
- Fudenberg, D., D. Levine and J. Tirole (1985), "Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information," in Alvin Roth (ed), *Game Theoretic Models of Bargaining*, Cambridge, England: Cambridge University Press.

Green, E.J. and R. Porter (1984): "Noncooperative Collusion Under Imperfect Price Information," *Econometrica*, 52 (1), 87-100.

Grossman, S. and M. Perry (1986): "Sequential Bargaining under Asymmetric Information," *Journal of Economic Theory*, 39 (1), 120-154

Hurwicz, L. (1972): "On Informationally Decentralized Systems,"

Leonardelli, G. J., J. Gub, G. McRuer, and V. Husted Medvec (2019) "Multiple equivalent simultaneous offers (MESOs) reduce the negotiator dilemma: How a choice of first offers increases economic and relational outcomes," *Organizational Behavior and Human Decision Processes* 152, 64-83.

Inderst, R. (2003): "Alternating-offer bargaining over menus under incomplete information," *Economic theory*, 22 (2), 419-429.

Jackson, M.O. (2001) "A Crash Course in Implementation Theory," *Social Choice and Welfare*, Vol. 18, No. 4, pp 655-708.

Jackson, M.O. (2003) "Mechanism Theory," in the *Encyclopedia of Life Sciences*, EOLSS Publishers: Oxford UK. <http://ssrn.com/abstract=2542983>

Jackson, M.O., L.K. Simon, J.M. Swinkels and W.R. Zame (2002): "Communication and equilibrium in discontinuous games of incomplete information," *Econometrica*, 70 (5), 1711-1740.

Jackson, M.O. and H.F. Sonnenschein (2007): "Overcoming Incentive Constraints by Linking Decisions," *Econometrica*, 75 (1), 241 - 258.

Kennan, J. and R. Wilson (1993): "Bargaining with private information," *Journal of Economic Literature*, 31 (1), 45-104.

Krutz, G.S. (2001): "Hitching a Ride: Omnibus Legislating in the US Congress," Ohio State University Press.

Larsen, B. (2014) "The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions," National Bureau of Economic Research working paper 20431.

Leonardelli, G.J., Gu, J., McRuer, G., Medvec, V.H., and Galinsky, A.D. (2019): "Multiple equivalent simultaneous offers (MESOs) reduce the negotiator dilemma: How a choice of first offers increases economic and relational outcomes," *Organizational Behavior and Human Decision Processes*, 152, 64-83.

Matsushima, H. (2008): "Detail-free mechanism design in twice iterative dominance: Large economies," *Journal of Economic Theory*, 141 (1), 134-151.

McAfee, R.P., (1992): "Amicable Divorce: Dissolving a Partnership with Simple Mechanisms," *Journal of Economic Theory*, 56 (2), 266-293.

McAfee, R.P., J. McMillan and M.D. Whinston (1989): "Multiproduct Monopoly, Commodity Bundling and the Correlation of Values," *Quarterly Journal of Economics*, 104 (2), 371-383.

Mnookin, R. *Bargaining with the devil: When to negotiate, when to fight*, Simon and Schuster, 2010.

Myerson, R. and M. Satterthwaite (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29, 265-281.

Radner, R. (1981): "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship," *Econometrica* 49, 1127-1148.

Roughgarden, T. and I. Talgam-Cohen (2013): "Optimal and near-optimal mechanism design

with interdependent values,” *Proceedings of the fourteenth ACM conference on Electronic commerce* 767-784.

Rubinstein, A. (1982): “Perfect Equilibrium in a Bargaining Model,” *Econometrica* 50 (1): 971-109.

Rubinstein, A. (1985): “A Bargaining Model with Incomplete Information about Time Preferences,” *Econometrica*, 53 (5), 1151-1172.

Satterthwaite, M. and S. Williams (2002) “The optimality of a simple market mechanism,” *Econometrica*, 70 (5), 1841-1863.

Satterthwaite, M., S. Williams and K. Zachariadis (2014): “Optimality versus practicality in market design: a comparison of two double auctions,” *Games and Economic Behavior*, 86, 248-263.

Sen, A. (2000): “Multidimensional bargaining under asymmetric information.” *International Economic Review*, 41 (2): 425-450.

Stahl, I. (1972): “Bargaining theory,” Doctoral thesis, Stockholm School of Economics.

Skrzypacz, A. and J. Toikka (2013): “Mechanisms for Repeated Bargaining,” mimeo.

Townsend, R. (1982): “Optimal Multiperiod Contracts and the Gain from Enduring Relationships under Private Information,” *Journal of Political Economy*, 90 (6), 1166-1186.

Wilson, R. (1987) “Game theoretic analysis of trading processes,” in T. Bewley (Ed.), *Advances in Economic Theory*, Cambridge University Press (1987)

## A Appendix: Proofs

### Proof of Theorem 1:

This Theorem is a corollary to Theorem 2 given the fact that the negotiation protocols 2.2.4 are rich: they include fraction-demanding offers. In particular, for any  $v$ , a fraction- $v$  demanding offer is to announce the truth  $\theta_i^n$  and demand a payoff of  $v$ . ■

### Proof of Theorem 2:

We begin with some notation. Let  $\Theta^n(f) \equiv \{(\theta_a, \theta_b) \mid f(\theta_a, \theta_b) > 0\}$  be the set of (profiles of) types that are “possible” under the joint prior distribution  $f$ , and  $\Theta_i^n(f) \equiv \{\theta_i \mid f_i(\theta_i) > 0\}$  is similarly defined for agent  $i$ .

Note that the assumption of known surplus implies that  $S(\theta_a, \theta_b) = \bar{S}$  for all  $(\theta_a, \theta_b) \in \Theta^n(f)$ . In addition, in a wPBE, after any history, the joint posterior distribution  $\tilde{f}$ ’s support is a subset of  $\Theta^n(f)$ , and similarly the posterior over  $i$ ’s type has a support as a subset of  $\Theta_i^n(f)$ . This is true both on and off the equilibrium path since consistent beliefs must have a support that is a subset of the prior’s support.

We now establish the upper and lower bounds of the seller’s utility in any equilibrium, as well as the buyer’s utility, and show that they all correspond to a unique equilibrium payoff that corresponds to immediate and efficient trade, and the Rubinstein shares.

In particular, let  $M_i^t [L_i^t]$  be the supremum [infimum] of the expected continuation payoff for agent  $i$ , starting at the beginning of period  $t$  over all histories such that  $N(h^{t-1}) = N$  (i.e. all items are still remaining), in all wPBE, and all  $i$ ’s types in  $\tilde{\theta}_i \in \Theta_i^n(f)$ .

We first show that  $M_a^0 \leq \frac{1}{1+\delta}\bar{S}$ .

At  $t+1$  ( $k$  even), the buyer makes the offers. Suppose all items are remaining, we argue that any buyer with  $\theta_b \in \Theta_b^n(f)$  can guarantee a payoff arbitrarily close to

$$L_b^{t+1} \equiv \bar{S} - \delta M_a^{t+2}.$$

The buyer does so by offering a fraction- $\kappa_{t+1}$  demanding offer with  $\kappa_{t+1} \equiv (L_b^{t+1} - \eta)/\bar{S}$  for  $\eta > 0$  arbitrarily small. Such an offer is accepted for sure for a seller with any type  $\hat{\theta}_a$  s.t.  $(\hat{\theta}_a, \theta_b) \in \Theta^n(f)$ : notice that  $\hat{\theta}_a \in \Theta_a^n(f)$  by construction, hence  $S(\hat{\theta}_a, \theta_b) = \bar{S}$ ; therefore (recall the second part of the definition of “fraction-demanding” in 3.3) by accepting the offer the seller can find a subset of items to trade such that she gets  $\delta M_a^{t+2} + \eta$ , which exceeds  $\delta M_a^{t+2}$ , the present value of the payoff from the continuation of the game if rejecting the offer (which means all items are still remaining).

By the first part of the definition of “fraction-demanding”, the buyer gets a payoff of  $L_b^{t+1} - \eta$  if the above fraction- $\kappa_{t+1}$  demanding offer is accepted, regardless which items the seller picks to trade.

At  $t$ , the seller makes offer. Suppose all items are remaining, we argue that a seller with any type  $\theta_a \in \Theta_a^n(f)$  can get a payoff at most  $\bar{S} - \delta L_b^{t+1}$ : With any type  $\theta_b \in \Theta_b^n(f)$ , by rejecting an offer at  $t$  (therefore all items are still remaining), the buyer’s payoff from the continuation of the game has a present value of at least  $\delta L_b^{t+1} - \delta\eta$  for  $\forall \eta > 0$ . Hence the payoff left to the seller with  $\theta_a \in \Theta_a^n(f)$  is at most  $\bar{S} - \delta L_b^{t+1}$ , as the (expected) surplus is  $\bar{S}$  by construction.

By definition of  $M_a^t$ , we have  $M_a^t \leq \bar{S} - \delta L_b^{t+1} \leq (1 - \delta)\bar{S} + \delta^2 M_a^{t+2}$

The above is true for any  $k = 0, 2, 4, \dots$ . Iteratively applying the above leads to

$$M_a^0 \leq \frac{1}{1+\delta}\bar{S}.$$

By a similar argument, it follows that

$$L_a^0 \geq \frac{1}{1+\delta}\bar{S}.$$

Therefore, the payoff for the seller with any possible type  $\theta_a \in \Theta_a^n(f)$  in any wPBE is  $U_a \leq \frac{1}{1+\delta}\bar{S}$ .

From the above, we also know that  $M_a^2 = L_a^2 = \frac{1}{1+\delta}\bar{S}$ , hence  $M_b^1 = L_b^1 = \frac{1-\delta}{1+\delta^2}\bar{S}$  (both in terms of the present value then), i.e. the total surplus realized is at least  $\frac{1}{1+\delta}\bar{S} + \delta\frac{1}{1+\delta}\bar{S} = \bar{S}$  which is the surplus from efficient trade. Hence the negotiation outcome must be efficient, which means immediate trade with the efficient set of items being exchanged. The utility terms correspond to the Rubinstein shares. ■

Next, we turn to the case with uncertainty. We begin by a lemma that establishes a rate of updating in a protocol with trembles.

LEMMA 1 For any event  $E \subset \Theta_i$ , let  $\Pr(E)$  be its prior in some period and  $\Pr(E | a_i)$  be the posterior one-period after conditional an action  $a_i$ . It follows that

$$\Pr(E | a_i) \leq \Pr(E)/\underline{\gamma},$$

where  $\underline{\gamma} > 0$  is the lower bound of the size of trembles (from any type) to  $a_i$ .

**Proof of Lemma 1:**

Giving updating according to Bayes' rule:

$$\Pr(E | a_i) = \frac{\Pr(a_i | E) \Pr(E)}{\Pr(a_i | E) \Pr(E) + \Pr(a_i | E^c) \Pr(E^c)} \leq \Pr(E)/\underline{\gamma},$$

where  $E^c$  is the complement of  $E$ , and the inequality comes from  $\Pr(a_i | E) \leq 1$  and  $\Pr(a_i | \cdot) \geq \underline{\gamma}$  due to trembles. ■

**Proof of Theorem 3:**

**Notation:**

Let  $\Phi_i^n \subset \Delta(\Theta_i)$  be the collection of all possible frequencies of  $n$  items with valuations picked from  $\Theta_i$ .

Let  $\phi_i^n[\theta_i] \in \Delta(\Theta_i)$  denote the frequency of a valuation type  $\theta_i$ . So  $\phi_i^n : \{\Theta_i^n\}_{i,n} \rightarrow \{\Phi_i^n\}_{i,n}$ , and the notation  $\phi_i^n[\theta_i](\theta_{ik})$  denotes the fraction of items having a specific value  $\theta_{ik}$ .<sup>35</sup>

The expected frequency is the prior distribution for each item, which is  $\bar{\phi}_i$  over  $\Theta_i$ .

When there is no ambiguity, we drop  $[\theta_i]$  and use  $\phi_i^n \in \Phi_i^n$  for  $i$ 's true realized frequency (given her type  $\theta_i$ ), and  $\hat{\phi}_i^n$  for a feasible frequency that can be announced.

Recall that  $S(\theta_a, \theta_b)$  is the total surplus, and  $s(\theta_a, \theta_b)$  is the per-item surplus, with the corresponding pair of valuation types.

We work with the per-item surplus  $s$ , unless otherwise noted.

With a slight abuse of notation, we extend the definition of this function to capture the expected surplus as a function of a frequencies:

- $s(\theta_a, \phi_b^n) = \frac{1}{n} \sum_k \sum_{\theta_{bk}} \phi_b^n(\theta_{bk}) \cdot (\theta_{bk} - \theta_{ak})_+$
- $s(\phi_a^n, \theta_b) = \frac{1}{n} \sum_k \sum_{\theta_{ak}} \phi_a^n(\theta_{ak}) \cdot (\theta_{bk} - \theta_{ak})_+$
- $s(\phi_a^n, \phi_b^n) = \frac{1}{n} \sum_k \sum_{\theta_{ak}, \theta_{bk}} \phi_a^n(\theta_{ak}) \phi_b^n(\theta_{bk}) \cdot (\theta_{bk} - \theta_{ak})_+$

Note that  $s(\phi_i^n[\theta_i], \phi_j^n) = s(\theta_i, \phi_j^n)$ ,  $\forall \theta_i, \phi_j^n$ ; i.e., the expected surplus (given a frequency of the other agent) depends only on one's true frequency  $\phi_i^n[\theta_i]$ , due to the independence across agents' valuations.

---

<sup>35</sup>Notice that a term in the square brackets is the valuation type, i.e. a  $(n)$ -vector, whereas a term in the parentheses is a number. For instance, if the seller  $s$ 's valuations are drawn from  $\{0,8\}$  for each of the  $n = 5$  items, then  $\Phi_s^5 = \{(x, y) \in \{0, \frac{1}{5}, \dots, 1\}^2 \mid x + y = 1\}$ ; and with a type  $\theta_a = (0, 0, 8, 8, 8)$ , agent  $a$ 's realized frequency is  $\phi[\theta_a] = (\frac{2}{5}, \frac{3}{5})$ , where  $\phi[\theta_a](0) = \frac{2}{5}$  and  $\phi[\theta_a](8) = \frac{3}{5}$ .

Although the agents' beliefs are defined over each other's valuation types  $\theta_i$ , when both agents use exchangeable strategies in Phase 1, the above observation implies that it suffices to focus on each other's frequencies (when analyzing beliefs at any nodes except the last part of Phase 2 at which point beliefs are not longer relevant).

We next define sets of frequencies that are less than some pre-specified distance  $d$  from the expected frequency (i.e., the prior)  $\bar{\phi}_i$ :  $\Phi_i^n(d) \equiv \{\phi_i^n : |\phi_i^n - \bar{\phi}_i| < d\}$  and  $\Phi^n(d) \equiv \{(\phi_a^n, \phi_b^n) : |\phi_i^n - \bar{\phi}_i| < d, i = a, b\}$ , where  $|\cdot|$  is sup norm. Note that the sets naturally depend on  $\bar{\phi}_i$ 's, but we omit them in the notation since they are fixed throughout the statement and proof of the theorem.

Let

$$\alpha_i^0(n, d) \equiv \Pr(|\phi_i^n[\theta_i] - \bar{\phi}_i| \geq d)$$

be the time-0 prior probability of frequencies that differ by at least  $d$  from the expected frequency. Let  $\alpha_i^t(n, d, \gamma) \equiv (\gamma_n)^{-t} \alpha_i^0(n, d)$ , where  $\gamma$  is the total rate of trembles and  $\gamma_n = \frac{\gamma}{\max_i |\Phi_i^n| \cdot |V^n|}$  is the (minimal) rate of trembles to each action when there are  $n$  items. Then, by Lemma 1, conditional on *any* history  $h^t$  up to time- $t$ , the likelihood of frequencies that are at least distance  $d$  from the expected frequency is bounded above by  $\alpha_i^t(n, d, \gamma)$ , i.e.

$$\Pr(|\phi_i^n[\theta_i] - \bar{\phi}_i| \geq d \mid h^t) \leq \alpha_i^t(n, d, \gamma) \quad (1)$$

When there is no ambiguity, we simplify notation by using  $\alpha_i^0$  and  $\alpha_i^t$ 's, for a given set of parameters  $(n, d, \gamma)$ .

The proof proceeds as follows: We first show the existence of exchangeable equilibria. Then we extend the idea in the proof of Theorem 2, providing expected payoff bounds. Due to the uncertainty about overall surplus, we are no longer able to provide useful bounds over all types nor in an ex-post sense; however we can focus on the types that are less than distance  $d$  from the expected frequency and bound their expected payoffs. The bounds do not exactly pin down one's payoffs, but approximately so - thanks to the bounds on posteriors. In addition, we show that the overall errors brought by types that are at least distance  $d$  from the expected frequency and approximate bounds vanish as  $n$  becomes large. Finally we illustrate that of the vanishing errors imply our main statements, i.e. the approximate efficiency and uniqueness of divisions.

### Existence of exchangeable equilibria.

For any original negotiation game  $\mathcal{G}$  with  $n$  items, construct an induced game  $\tilde{\mathcal{G}}$  as follows: Suppose an agents  $i$  only observes  $i$ 's frequency type  $\phi_i^n[\theta_i]$ , instead of the valuation type  $\theta_i$ , until the beginning of Phase 2. (The phase 1 of) such a game  $\tilde{\mathcal{G}}$  has finite type spaces, finite action in each period, and "continuity at infinity", thus has an sequential equilibrium.<sup>36</sup>

Given a sequential equilibrium of  $\tilde{\mathcal{G}}$ , we construct the following exchangeable strategy/belief profile for  $\mathcal{G}$  which is also an equilibrium. In Phase 1, an agent  $i$  with type

---

<sup>36</sup> "Continuity at infinity" means the (time-0) continuation value of the game after period  $T$  vanishes as  $T$  goes to infinity. See, Fudenberg and Levine (1983), p.258 for the definition, and p.267 Theorem 6.1 for an existence result.

$\theta_i^n$  copies the type  $\phi_i^n[\theta_i]$ 's strategy in the sequential equilibrium for  $\tilde{\mathcal{G}}$ . The belief system induced by the belief system of in  $\tilde{\mathcal{G}}$ , so that in each information set an agent/type  $\theta_i^n$  shares the beliefs that the type  $\phi_i^n[\theta_i]$  has in game  $\tilde{\mathcal{G}}$  over the frequency space; and over the valuation type spaces, the beliefs are equally assigned to types corresponding to a same frequency.

The strategy is exchangeable by construction. In addition, such a strategy/belief profile is a Bayesian equilibrium of the original game  $\mathcal{G}$ : In Phase 1, for any  $i$ , given that the other agent always assign the same beliefs over  $i$ 's types with a same frequency,  $i$  cannot be strictly better off by deviating to a non-exchangeable strategy. Therefore we have shown the existence of exchangeable Bayesian equilibria in the original game  $\mathcal{G}$ .

### Expected payoff from an offer.

Exchangeability of strategies, together with the iid distributions and exchangeable trembles, implies the following strategies as part of an equilibrium continuation in Phase 2: The responder has a unique strict best reply to trade a subset of items  $\{k \mid \theta_{bk} > \theta_{ak}\}$  at a transfer that exactly delivers the payoff demanded by the offerer - provided these lead to positive payoffs, and otherwise to say 'No' that is to trade nothing (and do either if there is indifference),<sup>37</sup> If the offerer was truthful in Phase 1 on the announced frequency, then this offerer's strategy is to truthfully list valuations in Phase 2. (What happens in other subgames will not be important for the argument below.)

The Phase 2 strategies imply the following expected payoff (again, per-item) from offering/accepting an offer in Phase 1: An offerer of type  $\theta_i^n$  gets a (non-discounted, per-item) payoff of  $\hat{v}_i$  with a "truthful" offer  $(\phi_i^n[\theta_i], \hat{v}_i)$  if the offer is accepted; recall that  $\phi_i^n[\theta_i]$  is the true frequency of  $\theta_i$ . A responder's (non-discounted) expected per-item payoff from accepting an offer  $(\hat{\phi}_i^n, \hat{v}_i)$  is  $\mathbb{E}v_j = s(\theta_j, \hat{\phi}_i^n) - \hat{v}_i$ . This is true regardless of whether  $\hat{\phi}_i^n$  is  $i$ 's true frequency, since that is the constraint subject to which  $i$  has to list valuations in Phase 2.

### Payoff bounds for agents with frequencies that are less than distance $d$ from the expected frequency.

Next we bound the expected payoffs of agents whose frequencies are less than distance  $d$  from the expected frequency in any Bayesian equilibrium, and then show the upper and lower bounds to an agent's payoffs are close to each other. Formally, define the following payoff bound(s), for  $i \in \{a, b\}$ :

- $M_i(\alpha_a, \alpha_b; d, n)$  [ $L_i(\alpha_a, \alpha_b; d, n)$ ] is the sup [inf] of expected *per-item* payoff from the continuation of the game (discounted to the current point of the game) that agent  $i$  can obtain in any equilibrium, with any  $\phi_i^n \in \Phi_i^n(d)$ , and at any decision node of the game such that the current posteriors satisfy  $\Pr(|\phi_a^n[\theta_a] - \bar{\phi}_a| < d \mid h^t) \geq 1 - \alpha_a$  and  $\Pr(|\phi_b^n[\theta_b] - \bar{\phi}_b| < d \mid h^t) \geq 1 - \alpha_b$ .

When there is no ambiguity we write them as  $M_i(\alpha_a, \alpha_b)$  and  $L_i(\alpha_a, \alpha_b)$ , but notice that the payoff bounds do depend on  $(d, n)$ .

---

<sup>37</sup>Generally, if items can have the same value for buyers and sellers then the strategy specification on whether those particular items trade is undetermined and does not influence the argument.

Note that the expected surplus with any type whose frequency is less than distance  $d$  from the expected frequency is close to the limit surplus  $\bar{s}$ :

$$|\mathbb{E}s(\phi_a^n, \phi_b^n) - \bar{s}| < 2ds_{max}, \forall (\phi_a^n, \phi_b^n) \in \Phi^n(d), \quad (2)$$

in which

$$s_{max} \equiv \max\{\Theta_b\} - \min\{\Theta_a\}.$$

1. We now show that when  $i$  makes an offer at  $t$  and  $j$  is the responder:

$$M_i(\alpha_a^t, \alpha_b^t) \leq \bar{s} + 2ds_{max} - \delta(1 - \alpha_j^{t+1})L_j(\alpha_a^{t+1}, \alpha_b^{t+1}) + \frac{1}{n} \quad (3)$$

where the  $\alpha_i^t$ 's are the previously defined bounds on posteriors (of frequencies that at least distance  $d$  from the expected frequency).

$$L_i(\alpha_a^t, \alpha_b^t) \geq (1 - \gamma)(1 - \alpha_j^t) \left[ \bar{s} - 2ds_{max} - \delta M_j(\alpha_a^{t+1}, \alpha_b^{t+1}) - \frac{1}{n} \right] \quad (4)$$

Proof of (3) and (4):

(3) is straightforward, by noting that  $\bar{s} + 2ds_{max}$  is an upper bound on the expected total surplus that remains by (2), and  $\delta L_j(\alpha_a^{t+1}, \alpha_b^{t+1})$  is a lower bound of  $j$ 's expected present value of rejecting  $i$ 's current offer, with the extra  $\frac{1}{n}$  being the largest possible (per-item) loss due to the unit gap of payoff grids.

(4): Noting that  $\bar{s} + 2ds_{max}$  is a lower bound on the expected total surplus that remains by (2) Consider an offer from  $i$  with her true frequency  $\phi_i^n$  and any demanded payoff of nor more than  $S - 2ds_{max} - \delta M_j(\alpha_a^{t+1}, \alpha_b^{t+1})$ . Such an offer will be accepted by  $j$  with any frequency  $\phi_j^n \in \Phi_j^n(d)$ , since  $j$ 's payoff from this offer exceeds the present value of  $M_j(\alpha_a^{t+1}, \alpha_b^{t+1})$ , the upper bound of what she can get when rejecting the offer. Hence the probability of acceptance is at least  $(1 - \alpha_j^t)(1 - \gamma)$ , with  $1 - \gamma$  being the likelihood that trembles do not apply.

2. Iteratively applying Equations (3) and (4) lead the following time-0 bounds on payoffs (assuming  $S$  is the offerer in the initial period, the other case is analogous).

Let  $error^{2t}$  be a bound on ‘‘error terms’’ that will bound how far expected payoffs can differ from the Rubinstein shares, which is defined by

$$error^{2t} = (\alpha_a^{2t} + \alpha_b^{2t} + \delta(\alpha_a^{t+1} + \alpha_b^{t+1}))s_{max} + (1 + \delta)2ds_{max} + (1 + \delta)\frac{1}{n} + \delta\gamma s_{max}.$$

Then it follows that

$$\sum_{t=0}^{T-1} error^{2t} = \sum_{t=0}^{2T-1} [\delta^t(\alpha_a^t + \alpha_b^t)] s_{max} + \frac{1 - \delta^{2T-2}}{1 - \delta} (2ds_{max} + \frac{1}{n} + \delta\gamma s_{max}).$$

Then, from an iterative application of (3) and (4):

$$\begin{aligned} M_a(\alpha_a^0, \alpha_b^0) &\leq \bar{s}(1 - \delta) + \delta^2 M_a(\alpha_a^2, \alpha_b^2) + error^0 \\ &= \bar{s}(1 - \delta)(1 + \delta^2 \dots + \delta^{2T}) + \delta^{2T} M_a(\alpha_a^{2T}, \alpha_b^{2T}) + \sum_{t=0}^{T-1} error^{2t} \\ &\leq \frac{1}{1 + \delta} \bar{s} + \left( \frac{\delta^{2T}}{1 - \delta^2} s_{max} + \sum_{t=0}^{T-1} error^{2t} \right), \end{aligned} \quad (5)$$

$$\begin{aligned}
L_a(\alpha_a^0, \alpha_b^0) &\geq \bar{s}(1 - \delta) + \delta^2 L_a(\alpha_a^2, \alpha_b^2) - error^0 \\
&= \bar{s}(1 - \delta)(1 + \delta^2 \dots + \delta^{2T}) + \delta^{2T} L_a(\alpha_a^{2T}, \alpha_b^{2T}) - \sum_{t=0}^{T-1} error^{2t} \\
&\geq \frac{1}{1+\delta} \bar{s} - \sum_{t=0}^{T-1} error^{2t}.
\end{aligned} \tag{6}$$

This implies that

$$\frac{1}{1+\delta} \bar{s} - \sum_{t=0}^{T-1} error^{2t} \leq L_a(\alpha_a^0, \alpha_b^0) \leq M_a(\alpha_a^0, \alpha_b^0) \leq \frac{1}{1+\delta} \bar{s} + (\delta^{2T} s_{max} + \sum_{t=0}^{T-1} error^{2t}), \tag{7}$$

3. Next, we show that all “error” terms go to 0 as  $n \rightarrow \infty$ . In particular, for  $\forall \varepsilon > 0$ ,  $\forall \delta, \delta < 1$ , first pick  $\eta > 0$  such that  $\max\{4\eta, \frac{4\eta}{\delta(\bar{s}-\eta)} + \frac{\eta}{5s_{max}}, \frac{6(1+\delta)\eta}{5\delta\bar{s}}, \frac{(1-\delta)\delta\eta}{1-\delta\bar{s}}\} < \varepsilon$  – this is the tolerance level of errors in payoff bounds that we allow for. Then in turn:

- pick  $T \in \mathbb{Z}_+$  such that  $\delta^{2T} s_{max} < \eta/5$
- pick  $d > 0$  such that  $\frac{2ds_{max}}{1-\delta} < \eta/5$
- pick  $\gamma_\varepsilon < \frac{\varepsilon}{5(1-\delta_{max})s_{max}}$ , hence  $\frac{\gamma(\eta)s_{max}}{1-\delta} < \eta/5$
- for any  $\gamma \in (0, \gamma_\varepsilon)$  and  $d$  (already picked), pick the threshold number of items,  $n_\gamma$ , so that for  $\forall n > n_\gamma$  we have  $\sum_{t=0}^{2T-1} [\delta^t (\alpha_a^t(d, n, \gamma) + \alpha_b^t(d, n, \gamma))] s_{max} < \eta/5$ . To do so, recall  $\alpha_i^t(d, n, \gamma) = \gamma_n^{-t} \alpha_i^0(d, n)$ , where
  - ◊  $\gamma_n^{-t} \sim O(n^{\max_i |\Theta_i| t})$ , where  $|\Theta_i|$  is the number of feasible valuations (for each item);
  - ◊  $\alpha_i^0 = \Pr^n \left( |\widehat{\phi}_i^n - \bar{\phi}_i| \geq d \right) \leq 2e^{-2d^2 n}$ , according to the Dvoretzky-Kiefer-Wolfowitz (1956) inequality,
  - ◊ hence fix any  $t$ ,  $\alpha_i^t = \gamma_n^{-t} \alpha_i^0 \rightarrow 0$  as  $n \rightarrow \infty$ ; so does their discounted sum (up to  $T - 1$ ),
- finally, to control the errors directly introduced by the price grids, if  $\frac{1}{n_\gamma(1-\delta)} \geq \eta/5$ , replace  $n_\gamma$  by  $\frac{5}{\eta(1-\delta)}$  so that  $\frac{1}{n(1-\delta)} < \eta/5$  for  $\forall n > n_\gamma$ .

4. We now put the pieces together to obtain tight equilibrium payoff bounds. In any equilibrium, the seller’s expected time-0 payoff

$$\mathbb{E}U_a^0(\phi_a^n) \in \left( \frac{1}{1+\delta} \bar{s} - \eta, \frac{1}{1+\delta} \bar{s} + \eta \right), \forall n > \bar{n}, \forall \phi_a^n \in \Phi_a^n(d) \tag{8}$$

Similarly, the buyer’s expected time-1 payoff (whenever time-1 is reached)

$$\mathbb{E}U_b^1(\phi_b^n) \in \left( \frac{1}{1+\delta} \bar{s} - \eta, \frac{1}{1+\delta} \bar{s} + \eta \right), \forall n > \bar{n}, \forall \phi_b^n \in \Phi_b^n(d) \tag{9}$$

Hence the buyer’s expected time-0 payoff, in any equilibrium, is at least

$$\mathbb{E}U_b^0(\phi_b^n) \geq \delta(1 - \gamma) \left( \frac{1}{1+\delta} \bar{s} - \eta \right) > \frac{\delta}{1+\delta} \bar{s} - \frac{6}{5} \eta, \quad \forall n > \bar{n}, \forall \phi_b^n \in \Phi_b^n(d) \tag{10}$$

By the construction of  $\eta$ , we have  $\mathbb{E}U_a^0(\phi_a^n)$  and  $\mathbb{E}U_b^0(\phi_b^n)$  are both in the region of  $(1 - \varepsilon, 1 + \varepsilon)$  times the corresponding Rubinstein shares with the limit surplus.

**Realized surplus and likelihood of immediate trade.**

From Equations (8) and (10),  $\forall n > \bar{n}$ , in any equilibrium, the realized surplus is at least

$$\mathbb{E}U_a^0(\phi_a^n) + \mathbb{E}U_b^0(\phi_b^n) > \bar{s} - 3\eta, \quad \forall (\phi_a^n, \phi_b^n) \in \Phi^n(d) \quad (11)$$

In expectation, the surplus realized is at least:

$$\Pr(\Phi^n(d))(\bar{s} - 3\eta) \geq \left(1 - \frac{\eta}{5s_{max}}\right)(\bar{s} - 3\eta) > \bar{s} - 4\eta > \bar{s} - \varepsilon \quad (12)$$

recall  $\Pr(\Phi^n(d)) = (1 - \alpha_a^0)(1 - \alpha_b^0) \geq 1 - 2d > \frac{\eta}{5s_{max}}$  for  $\forall n \geq \bar{n}$ .

Now we turn to the likelihood of immediate trade:

With any pair of types whose frequencies are less than distance  $d$  from the expected frequency, the maximal surplus is at most  $\bar{s} + \eta$ , and the total cost of delay for one period is at least  $\delta(\bar{s} - \eta)$ . Hence with such types, the likelihood of delay is at most  $\Pr(\text{delay} | \Phi_{Emp}^n(d)) = 4\eta / [\delta(\bar{s} - \eta)]$ . This gives a bound on overall delay:

$$\Pr(\text{delay}) \leq \Pr(\text{delay} | \Phi^n(d)) \Pr(\Phi^n(d)) + 1 - \Pr(\Phi^n(d)) \leq \frac{4\eta}{\delta(\bar{s} - \eta)} + \frac{\eta}{5s_{max}} < \varepsilon \quad (13)$$

## B Supplementary Appendices

**Proof of Example 3** We construct such an equilibrium. We begin with the following simple situation of one-sided uncertainty.

*Situation I: one-sided uncertainty in the posteriors*

For instance, (0, 40) meets (10, 50) or (50, 10). The following is an equilibrium.

In any even period, including the initial one ( $t = 0$ ), Alice, the seller, offers the following options: (1) trade the first item only at  $p_1^0 = \frac{50}{1+\delta}$ , (2) trade both items at  $p_{12}^0 = 40 + \frac{20}{1+\delta}$ , or (3)  $p_2^0 = 50$  (which will never be chosen). This is a fraction- $\frac{1}{1-\delta}$  demanding offer.

For Bob, the buyer: with type (10, 50) he accepts any offer that includes an option from which he gets at least the Rubinstein share of  $\frac{20\delta}{1+\delta}$ , otherwise he rejects and counter-offers (in the next period) with a fraction- $\frac{1}{1-\delta}$  demanding offer: “both items at  $p_{12}^1 = 40 + \frac{20}{1+\delta}$ , or 2nd item at  $p_2^1 = \frac{50}{1+\delta}$ , or 1st item at  $p_1^1 = 0$ ”.

With type (50, 10) he accepts any offer that includes an option from which he gets at least the Rubinstein share of  $\frac{50\delta}{1+\delta}$ , otherwise he rejects and counter-offers with a fraction- $\frac{1}{1-\delta}$  demanding offer: “1st item at  $\frac{50}{1+\delta}$ , or both items at  $40 + \frac{20}{1+\delta}$ , or 2nd item at 0”.

The seller updates beliefs only when receiving an offer. (No updating upon a rejection of her offer, that is, when the responder chooses to trade nothing).

The seller’s posterior (after receiving an offer) is “the buyer has (10, 50) for sure” if receiving the offer described above that corresponds to the (10, 50) buyer; otherwise her posterior is “(10, 50) for sure” which implies that her belief is certain about the total surplus being 50. The seller accepts any offer that gives her a payoff of at least  $\frac{20\delta}{1-\delta}$  if she believes the buyer has (10, 50); otherwise she accepts any offer that gives her at least  $\frac{50\delta}{1-\delta}$ .

Note that if  $p_1^1 = p_2^1$  the buyer is believed to have (50, 10).

*Situation II: two-sided uncertainty.*

The (0, 40)-seller always acts (including offers, and decides whether to accept an offer) the same way as she does in Situation I. Similar for the (40, 0) seller, adjusting the item id’s.

Belief updating is also similar to situation I: upon receiving the first offer from the buyer, she believes the buyer has a certain type for sure if that offer corresponds the one (described in above) that shall be used by that buyer’s type; otherwise, she believes for sure that the buyer has the type that leads to the high surplus of 50 when matched with her own type.

The buyer offers the same way as he does in Situation I. He updates his belief in a similar manner as the seller: on-path, upon receiving the initial offer his posterior is certain about the seller’s type since the two types of the seller use separate offers; off-path, he believes for sure that the seller has the type that leads to the high surplus of 50 when matched with his own type.

We conclude the proof by verifying the optimality of the above-mentioned strategy profiles.

To begin with, for either agent, neither type can gain from mimicking the strategy of the other type. To illustrate this, the following table summarizes the outcomes on path. Consider the seller's type  $(0, 40)$ : by mimicking the other type  $(40, 0)$  her payoff would be  $\frac{50}{1+\delta} - 40$  and  $\frac{20}{1+\delta}$ , respectively, given the two types of the buyer; both are strictly lower than her on-path payoff ( $\frac{20}{1+\delta}$  and  $\frac{50}{1+\delta}$ , respectively) for any  $\delta \in (0, 1)$ . The same holds true for the buyer.

	(10, 50)	(50, 10)
(0, 40)	both at $p = 40 + \frac{20}{1+\delta}$	1st at $p = \frac{50}{1+\delta}$
(40, 0)	2nd at $p = \frac{50}{1+\delta}$	both at $p = 40 + \frac{20}{1+\delta}$

In addition, no agent/type gains from deviating to an off-path action. It follows from the construction of the belief system that doing so leads the other agent to believe that the realized surplus is high (50) for sure. That only makes the other agent more demanding.

Finally, a similar argument as in the proof of Theorem 2 establishes the payoff bounds, which justifies the the division of surplus  $\frac{1}{1+\delta}$ . ■

**Proof behind the Example from Section 2.1.1:**

To understand why this is the case, let us consider a simplified version of the Problem 2 in which things become quite transparent. Without loss of generality, presume that the seller makes the first offer.

Let us focus on the two items that the seller knows should trade - the ones for which she has value \$0. So, the seller is known to have value \$0 on the two items, while the buyer has value \$10 for one item and \$50 for the other item, with equal probability for which item is preferred by the buyer, and that realization is known only to the buyer. This problem has a known surplus of \$60, which is common knowledge along with the fact that both items should trade immediately. Showing that inefficiency results here where both items should trade, makes it easy to extend the argument to the situation in which the agents are not sure which goods should trade.

The basic logic behind inefficiency is as follows (the proof is below). Effectively, in order to get an efficient outcome, a price of no more than \$10 must be offered on each item in the first period. In doing this, the seller gives much of the surplus to the buyer. By instead offering higher prices and attempting to screen to get more on the \$50 item, the seller improves her expected utility. One can use upper and lower bounds on the continuation values to show that the buyer will accept a price of more than \$10 on the \$50 item. Thus, the seller is better off giving up some efficiency on trade - delaying trade on the \$10 item - to get more surplus. With more items, this loss of efficiency is substantial.

Let  $L_a$  be the seller's worst continuation payoff in any seller-offer period in any wPBE with both items remaining. This means when the buyer makes an offer, he gets a continuation payoff of at most  $\delta(60 - \delta L_a)$  since the seller can always reject on both items and counteroffers in the subsequent period.

Consider the seller's offer  $(p, p)$  with some  $p > 2$ . The buyer rejects  $p$  on the value-10 item, and accepts  $p$  on the value-50 item for sure if  $p < \tilde{p}$ , s.t.

$$50 - \tilde{p} + \frac{10\delta}{1 + \delta} = (60 - \delta L_a)\delta,$$

where on the left-hand side  $50 - p$  is the payoff from the value-50 item and  $\frac{10\delta}{1+\delta}$  from the value-10 item (the corresponding Rubinstein share, since it is commonly known that the item left is of value-10).

Therefore, with an offer of  $(\tilde{p} - \epsilon, \tilde{p} - \epsilon) \forall \epsilon > 0$ , the seller can always get an acceptance on the value-50 and a discounted Rubinstein share on the value-10, i.e. a payoff of  $\tilde{p} + \frac{10\delta^2}{1+\delta} - \epsilon$ .

On the other hand, since  $L_a$  is the seller's payoff in some SE, it must exceed the payoff from the above deviation  $(\tilde{p} - \epsilon, \tilde{p} - \epsilon)$ . This requires

$$L_a \geq \tilde{p} + \frac{10\delta^2}{1 + \delta} - \epsilon.$$

This, combined with the definition of  $\tilde{p}$ , gives  $(1 - \delta^2)L_a \geq 50(1 - \delta) - \epsilon$ , i.e. (since  $\epsilon$  can be arbitrarily small)

$$L_a \geq \frac{50}{1 + \delta}$$

Finally, for both items to be traded in the initial period the seller's expected payoff is at most 20, which is not possible in any wPBE: In order to have both items traded with a positive probability, the seller's strategy in the first period must put positive weight (if mixing) on an offer that has prices at most 10 on each item. The seller gets a payoff of at most 20 from such an offer, and hence an expected payoff of at most 20 from the game since the seller must be indifferent among any strategies used with positive probability. ■

## B.1 Additional Results with a Nearly-Known Surplus

Subsection 4.1.2 discussed a (technical) challenge associated with the multiplicity of sequential equilibria due to dramatic updating in posteriors. This challenge was handled in subsection 4.1.4 by introducing trembles and working with the frequency protocol. Here, we present two other approaches. The first approach imposes a restriction on how fast beliefs can be updated. The second approach considers trembles, but with a fixed number of items. Thus, in the second approach the convergence of beliefs is not derived from the law of large numbers but must come from some justification based on the knowledge the agents have about their environment.

The main advantage of these approaches is that they work with all protocols that have fraction-demanding offers (see section 3.3, including the value-announcing negotiations and combinatorial negotiations from Section 2.2.4, as well as their variations that allow for continued negotiation over not-yet-decided-upon items).

Both approaches consider a sequence of negotiation problems  $\{(n_m, \Theta, f^m)\}_{m=1,2,\dots}$  whose (per-item) surpluses converges to  $\bar{s} > 0$ , i.e.

$$\frac{S^m}{n^m} \rightarrow_p \bar{s}, \text{ as } m \rightarrow \infty$$

where recall that  $S^m$  is the surplus in the  $m$ -th problem.

We do not impose further additional assumptions, hence allow for, for instance:

- correlations across players (e.g., common shocks),
- correlations across items, and
- asymmetries across items.

We use  $m$  as an index since the numbering of the sequence may differ from the number of items. For example, the second approach has  $n_m = n, \forall m$ . Thus, the sequence applies to settings in which agents have increasingly accurate knowledge of the surplus based on some fundamental economic reason - e.g., having good information about the environment - rather than just relying on the law of large numbers.

Let the  $m$ -th problem have a surplus grid  $V^{n_m, \Delta} = \{0, \Delta, 2\Delta, \dots, n_m S_{\max}\}$  from which an agent can demand a total surplus.

### B.1.1 Approximate Efficiency Results with Bounded Belief Updating

For simplicity we consider protocols for which all agents' past actions are commonly observed; e.g., our first protocol from Section 2.2.4 and the combinatorial protocol from Section 2.2.4. In such protocols, at the beginning any period  $t$  agents share a common history  $h^{t-1} \equiv (a_{i(0)}^0, N^0, \dots, a_{i(t-1)}^{t-1}, N^{t-1})$ , and after the offerer moves the common history becomes  $(h^{t-1}, a_{i(t)}^t)$ . We denote the set of all possible histories by  $H$ , including  $h^0 \equiv \emptyset$  being the initial decision node of the negotiation game.

An agent  $i$ 's beliefs  $\tilde{f}_j : H \rightarrow \Delta(\Theta_j^n)$  map histories to a distribution over the other agent's type space. In particular, we let  $\tilde{f}_j(E, h^{t-1}, \theta_i)$  denote  $i$ 's belief over  $E$  conditional history  $h^{t-1}$ . Note that an agent  $i$ 's posterior belief can depend on  $i$ 's own type  $\theta_i$ .

The beliefs at the initial node (before types are drawn) are the common prior, i.e.  $\tilde{f}(\cdot | \emptyset) = f(\cdot)$ .

We require agents' initial beliefs conditional upon their types to be consistent with the common prior, in particular,  $\tilde{f}_j(\cdot, \emptyset, \theta_i) = f_j(\cdot, \theta_i), \forall \theta_i$ , where  $f_j(\cdot, \theta_i)$  is the marginal distribution (of  $f$ ) over  $\Theta_j^n$ .

Now we introduce a restriction on how fast beliefs can be updated. We say a belief system  $\tilde{f}_j(\cdot)$  satisfies *bounded updating at rate*  $\beta \geq 1$  if for  $\forall E \subset \Theta^n, h^{t-1} \in H, a_{i(t)}^t \in A_{i(t)}(h^{t-1})$ , and  $\theta_i \in \Theta_i^n$ :

$$\tilde{f}_j(E, (h^{t-1}, a_{i(t)}^t), \theta_i) \leq \beta \tilde{f}_j(E, h^{t-1}, \theta_i);$$

and for  $\forall E \subset \Theta^n$ ,  $(h^{t-1}, a_{i(t)}^t) \in H$ ,  $N^t \in N(h^{t-1})$ , and  $\theta_i \in \Theta_i^n$ :

$$\tilde{f}_j(E, (h^{t-1}, a_{i(t)}^t), N^t, \theta_i) \leq \beta \tilde{f}_j(E, (h^{t-1}, a_{i(t)}^t), \theta_i).$$

We consider (arbitrarily) large but bounded  $\beta$ s. This means that the restriction binds only for events that are very unlikely under the prior. In particular, when  $\tilde{f}(E | h^{t-1}) > \frac{1}{\beta}$ , the restriction is not binding.

Next, we introduce an equilibrium notion with bounded belief updating.

**DEFINITION 1 (EQUILIBRIUM WITH BOUNDED BELIEF UPDATING)** *An equilibrium with bounded belief updating at rate  $\beta$  is a profile of the agents' (mixed) strategies and posterior systems  $\tilde{f}_j, j = a, b$ , such that*

1. *At any decision node, the mover  $i$ 's strategy maximizes his/her expected payoff given the other's strategies and his/her posterior system about the other's types  $\tilde{f}_j$ ;*
2. *Both agents' posterior systems  $\tilde{f}_a$  and  $\tilde{f}_b$  satisfy bounded updating at rate  $\beta$ .*

Definition 1 imposes only minimal requirements on beliefs other than the bounded-updating requirement. In particular, we do not require that agents' posterior systems be induced by some joint posterior system, nor do we require that the agents' posterior systems are common knowledge, nor do they even have to be consistent with Bayes' rule. So this can be viewed as a notion that allows for the most possible outcomes as equilibria, under some (arbitrarily large) bound on updating. We show that even with such a minimal restriction, all equilibria are approximately efficient with vanishing uncertainty about overall surplus.<sup>38</sup>

**THEOREM 4** *Consider a sequence of negotiation problems  $\{(n_m, \Theta, f^m)\}_{m=1,2,\dots}$  such that the distributions  $\{f^m\}$  have a converging per-item surplus  $\bar{s} > 0$ , and the protocol includes fraction-demanding offers. For any  $\varepsilon > 0$ ,  $\forall \beta \in [1, \infty)$ ,  $\forall \delta, \delta < 1$ , there is  $\Delta(\varepsilon) > 0$  such that for any  $\Delta \in (0, \Delta(\varepsilon))$  there exists  $m_\Delta$  such that if  $m > m_\Delta$  then:*

1. *There exist equilibria with bounded belief updating at rate  $\beta$ ;*
2. *In any such equilibrium, with probability at least  $1 - \varepsilon$ :*
  - *agreement is reached in the initial period;*
  - *the realized surplus is at least  $(1 - \varepsilon)\bar{s}$ ; and*
  - *Expected payoff / 'Full-Information Rubinstein share for  $\bar{s}$ ' for each agent lies in  $(1 - \varepsilon, 1 + \varepsilon)$ .*

---

<sup>38</sup> An alternative notion would apply a variation on sequential equilibria, but with the modification that "beliefs are rounded to the boundary"; i.e., if the posterior on any event exceeds  $\beta$  times the prior, that belief on that event is replaced by  $\beta$  times the prior. Same theorem holds under that alternative notion.

#### A sketch of the proof of Theorem 4:

Most parts of the proof are similar to Proof of Theorem 3 and so are not repeated here.

One difference is that we require the distance between prices, measured by  $\Delta$ , to go to zero. This is because, in the current analysis, we do not impose the assumption that the number of items goes to infinity. The key assumption is again that the *per-item* grids are small enough.

Another new feature of the current theorem is that it does not impose the assumption of independence across agents' types, so that an agent's beliefs about the other's types may depend on his or her own type. This brings extra steps in proving the theorem, which are our main focus here.

First, for any distance  $d > 0$  and any type  $\theta_i$ , let  $\Theta_j^{n_m}(\theta_i, d) = \{\theta_j \mid |s(\theta_i, \theta_j) - \bar{s}| < d\}$  be the set of other's types for which the per-item surplus is close enough to the limit  $\bar{s}$  (within a distance of  $d$ ).

Recursively, construct the following sequences of subsets of the agents' types (for  $i = a, b$ ), given some  $\iota_0, \iota_1, \dots > 0$

$$\begin{aligned}
 0. \quad & \tilde{\Theta}_i^{n_m}(d, 0, \iota_0) = \{\theta_i \mid f_b(\Theta_j^{n_m}(\theta_i, d), \theta_i) > 1 - \iota_0\}, i = a, b; \\
 1. \quad & \tilde{\Theta}_i^{n_m}(d, 1, \iota_1) = \{\theta_i \mid f_b(\Theta_j^{n_m}(\theta_i, d) \cap \tilde{\Theta}_j^{n_m}(d, 0, \iota_0), \theta_i) > 1 - \iota_1\}, i = a, b; \\
 & \dots \\
 t. \quad & \tilde{\Theta}_i^{n_m}(d, t, \iota_t) = \{\theta_i \mid f_b(\Theta_j^{n_m}(\theta_i, d) \cap \tilde{\Theta}_j^{n_m}(d, t-1, \iota_{t-1}), \theta_i) > 1 - \iota_t\}, i = a, b; \\
 & \dots
 \end{aligned}$$

Intuitively, for an agent  $i$  with any type in the  $t$ -th set above,  $i$ 's prior is such that with a probability of at least  $1 - \iota_t$  the surplus is close to the limit  $\bar{s}$  (within a distance of  $d$ ) and that the other has a type in the  $(t-1)$ -th set. In terms of posteriors: for any period  $t'$ ,  $i$ 's posterior (following *any* history up to that period) is at least  $\beta^{t'}(1 - \iota_t)$  for the above events. To simplify notation we omit the superscript  $n_m$  when there is no ambiguity.

We provide payoff bounds for those sets. In particular, for some fixed  $T \in \mathbb{N}$  and  $\{\iota_0, \dots, \iota_{2T}\}$  (we discuss how to pick these below), for  $i \in \{a, b\}$ , let  $M_i^t$  [ $L_i^t$ ] be the sup [inf] of the expected payoff, per item, from the continuation of the game (discounted to the current point of the game) that agent  $i$  can obtain in any equilibrium, with any  $\theta_i \in \tilde{\Theta}_i(d, 2T - t, \iota_{2T-t})$ .

We derive the bounds recursively:

- Backward from period  $2T$ :

$$M_i^{2T} < \beta^{2T}(1 - \iota_{2T-t})(1 - \iota_0)(\bar{s} + d) + (1 - \beta^{2T}(1 - \iota_0))s_{\max}; \text{ and}$$

$$L_i^{2T} \geq 0.$$

...

- In period  $t < 2T$ , agent  $it$  makes the offer:

$M_{i(t)}^t < \beta^t(1 - \iota_{2T-t})(\bar{S} + d - \delta_{i(t+1)}L_{i(t+1)}^{t+1}) + (1 - \beta^t(1 - \iota_{2T-t}))\iota_{2T-t}s_{\max}$ , which is the maximum surplus left minus the share that must be delivered to the other agent with types in  $\tilde{\Theta}_{i(t+1)}(d, 2T - t, \iota_{2T-t})$ ; and

$L_{i(t)}^t > (1 - \beta^{2T-t}\iota_{2T-t})(\bar{s} - d - \delta_{i(t+1)}M_{i(t+1)}^{t+1} - \Delta)$ , since any offer that delivers at least  $\delta M_a^{2T}$  is accepted for sure by the other agent with types in  $\tilde{\Theta}_{i(t+1)}(d, 2T - t, \iota_{2T-t})$ .

...

The above process bounds time-0 payoffs for the types in sets  $\tilde{\Theta}_i(d, 2T, \iota_{2T})$ . It is easy to verify that (for any fixed  $T$ ) when  $d, \iota_0, \iota_1, \dots, \iota_{2T} > 0$  and  $\Delta > 0$  go to 0, the time-0 bounds  $M_a^0$  and  $L_a^0$  become arbitrarily close to each other, hence the expected payoff for any type in  $\tilde{\Theta}(d, 2T, \iota_{2T})$  is approximately determined. In addition, that payoff can be arbitrarily close to the corresponding Rubinstein share with a surplus  $\bar{S}$ , for large enough  $T$ .

The final step is to show that the above bounds have bite for most types (according to prior probabilities). Formally, notice that (fixing any  $T$ ) for large enough  $m$ , i.e. as the prior knowledge over surplus becomes precise enough, we can find small enough  $d, \iota_0, \iota_1, \dots, \iota_{2T} > 0$  while keeping  $\Pr(\tilde{\Theta}_i(d, 2T, \iota_{2T}))$  close enough to 1.

To do so, for any  $d > 0$ , let  $p_d = \Pr(|S - \bar{S}| < d)$  which converges to 0 according to the definition of converging surplus. We construct  $\iota$ 's from  $\alpha_0$ :

- $\iota_0 = \sqrt{\alpha}$ , easy to verify that  $f_i(\tilde{\Theta}_i(d, 0, \iota_0)) \geq \frac{\alpha}{\sqrt{\alpha}} = \iota_0$ ,

hence  $\Pr(|S - \bar{S}| < d \text{ and } \theta_j \in \tilde{\Theta}_j(d, 0, \iota_0)) \geq 1 - (\alpha + \iota_0)$ ;

- $\iota_1 = \sqrt{\alpha + \iota_0}$ , easy to verify that  $f_i(\tilde{\Theta}_i(d, 0, \iota_0)) \geq \frac{\alpha + \iota_0}{\sqrt{\alpha + \iota_0}} = \iota_1$ ; and

...

- $\iota_t = \sqrt{\alpha + \iota_{t-1}}$ ;

...

Continue this process until we get  $\iota_{2T}$ , which converges to 0 as  $\alpha$  goes to 0 (i.e. with large enough  $m$ ).

In summary, we can approximately determine expected payoffs for all types in  $\tilde{\Theta}_i(d, 2T, \iota_{2T})$ , whose (prior) probability is at least  $1 - \iota_{2T}$ , for arbitrarily small  $\iota_{2T}$  as  $m$  becomes large. The rest of the proof parallels the corresponding parts of the proof of Theorem 3. ■

### B.1.2 Trembles with Fixed Number of Items.

When there is no ambiguity we write  $M_i(\alpha_a, \alpha_b)$  and  $L_i(\alpha_a, \alpha_b)$ , though the payoff bounds do depend on  $(d, n)$ . We now work with a sequence of economies, in which the number(s) of items are bounded above (or fixed), and in which uncertainty over total surplus vanishes but substantial uncertainty about each item remains. This captures agents have accurate information about the surplus rather than relying on laws of large numbers to give them accurate information about the surplus.

The bound on the number of items implies boundedness of sizes of action spaces (given any increment  $\Delta > 0$  in the grids of surplus), and thus allows for an approximate efficiency results derived with trembles, instead of putting an artificial restriction on beliefs.

A bounded number of items and the vanishing uncertainty over overall surplus need not contradict with each other: the assumption captures an environment in which agent's knowledge of each other's total valuation is strong, which is natural in many settings.

In particular, we consider a sequence of negotiation problems  $\{(n, \Theta, f^m)\}_{m=1,2,\dots}$  (note the additional restriction that  $n_m = n, \forall m$ ), and again we work on the finite grids surplus  $V^{n,\Delta}$ . Trembles are introduced similar to those in subsection 4.1.4: consider trembles in any period by all types of any agent with probability  $\gamma$ , uniformly to each of the feasible actions.

**THEOREM 5** *Consider a sequence of negotiation problems  $\{(n, \Theta, f^m)\}_{m=1,2,\dots}$  such that the distributions  $\{f^m\}$  have a converging per-item surplus  $\bar{s} > 0$ , and the protocol includes fraction-demanding offers. For any  $\varepsilon > 0$ , there exist a small enough tremble probability  $\gamma(\varepsilon) > 0$  and increment of grids  $\Delta(\varepsilon) > 0$  such that for any  $\gamma \in (0, \gamma(\varepsilon))$  and  $\Delta \in (0, \Delta(\varepsilon))$  there exists  $m_{\gamma,\Delta}$  such that if  $m > m_{\gamma,\Delta}$  then:*

1. *There exist (perfect) Bayesian equilibria with trembles;*
2. *In any such equilibrium, with probability at least  $1 - \varepsilon$ :*
  - *agreement is reached in the initial period;*
  - *the realized surplus is at least  $(1 - \varepsilon)\bar{s}$ ; and*
  - *Expected payoff / 'Full-Information Rubinstein share for  $\bar{s}$ ' for each agent lies in  $(1 - \varepsilon, 1 + \varepsilon)$ .*

We omit a formal proof of Theorem 5. The existence part is straightforward. The rest of the theorem follows by the same logic as Theorem 4, since with fixed number of items (hence fixed action space) and trembles, beliefs are updated at bounded rates (per-period) that are invariant to  $m$ . The only exception involves the errors directly due to the trembles, which are arbitrarily small as the total size of trembles goes to 0.

## B.2 Unknown Surplus but Known Frequencies with Two Costs/Values

This part generalizes Example 3 in Section 4.2. We first provide an example with  $n = 4$  items, which corresponds to our experimental treatment “4 goods - unknown surplus, known frequencies”. Then we generalize the observation to a situation with arbitrary number of items, in which the cost and value for each item take on just two values.

**EXAMPLE 4 (4 goods, known frequencies, unknown surplus)** *There are four items and Alice’s costs are  $(0, 0, 40, 40)$  or its permutations, and Bob’s values are  $(10, 10, 50, 50)$  or its permutations. The possible surpluses are 40, 70, or 100.*

*Under the combinatorial negotiation, there exists an “appealing” efficient equilibrium: at  $t = 0$ , Alice with  $(0, 0, 40, 40)$  offers the following*

- *trade the first two items at  $p = \frac{100}{1+\delta}$ ; or*
- *trade the first two, plus either the 3rd or the 4th (three items in total), at  $p = 40 + \frac{70}{1+\delta}$ ; or*
- *trade all the four items, at  $p = 80 + \frac{40}{1+\delta}$ .*

*On path, Bob accepts, and picks the option that achieves the maximal realized surplus.*

Now we generalize Examples 3 and 4 to allow for more general type spaces. We focus on the case in which the cost and value for each item take on two values. We also suppose that the frequencies of high’s and low’s are commonly known, while the orderings across items can be arbitrary. Formally:

**ASSUMPTION 1 (KNOWN FREQUENCY WITH TWO COSTS/VALUES)** *For each item:  $k \in N$ ,  $\theta_{ak} = c_L, c_H$  and  $\theta_{bk} = v_L, v_H$ , with  $c_L < v_L < c_H < v_H$ . In addition, types have known frequencies:  $\#_k(\theta_{ak} = c_L) = l_a$ ,  $\#_k(\theta_{bk} = v_L) = l_b$ ,  $\forall \theta_a \in \Theta_a^n, \theta_b \in \Theta_b^n$ .*

We note that the above assumption only limits the type space. It does *not* require any specific distribution over the space, nor independence between the seller’s and the buyer’s types. In fact, we put no restrictions on the common prior  $f$ .

In addition, there is substantial uncertainty about the surplus, as in Examples 3 and 4.

Now we are ready to present a positive result that generalizes the two examples.

**PROPOSITION 2** *Suppose Assumption 1 holds and agents negotiate under rich negotiations. There exists an equilibrium that is efficient and leads to the Rubinstein shares. On path, in the initial period the seller proposes a fraction- $\frac{1}{1-\delta}$  demanding offer for her type  $\theta_a$ . The buyer accepts, trades the efficient set of items, and the game ends.*

**Proof of Proposition 2.** The proof of this Proposition generalizes the proof of Example 3. We construct an equilibrium that is efficient and leads to the Rubinstein shares.

On-path, the seller, as the initial offerer, proposes a fraction- $\frac{1}{1-\delta}$  demanding offer. It follows from Assumption 1 that such offers are different for different types of the offerer, and therefore the seller's strategy is fully separating.

The initial responder, the buyer, accepts and trades optimally. The game ends. It follows from the definition of fraction demanding offers that the above on-path outcome leads to the efficient set of items being traded and the Rubinstein shares, for any  $(\theta_a, \theta_b)$ .

Should there be any further period(s), the offerer  $i(t)$  always proposes a fraction- $\frac{1}{1-\delta}$  demanding offer given own type  $\theta_{i(t)}$ .

Now we characterize the buyer's belief system. On-path, upon receiving the initial offer, his posterior becomes degenerate, since the seller's on-path offer reveals her type. Off-path, the buyer believes that the seller has the type(s) that lead to the maximum possible surplus,  $\bar{S}(\theta_b) \equiv \max_{\theta_a \in \Theta_a^n} S(\theta_a, \theta_b)$ , when matched with his own type  $\theta_b$ .

The seller's belief system is similar, and only depends on the buyer's offer at  $t = 1$  (if he rejected the seller's initial offer). If the buyer's offer is fraction- $\frac{1}{1-\delta}$  demanding for some type  $\theta_b \in \Theta_b^n$ , then the seller believes that the buyer has that type for sure. Otherwise, the seller believes that the buyer has the type(s) that lead to the maximum possible surplus,  $\bar{S}(\theta_a) \equiv \max_{\theta_b \in \Theta_b^n} S(\theta_a, \theta_b)$ , when matched with her own type  $\theta_a$ .

In addition, off the equilibrium path, the responder  $i(t)$  accepts the current offer if and only if that offer gives him a payoff of at least  $\bar{S}(\theta_{i(t)})$ .

Finally, we follow a similar argument as in the proof of Example 3 to verify the optimality of the above described strategies. Details are omitted here. ■

Proposition 2 holds for any rich negotiations for the current negotiation problem. One example is the combinatorial negotiations. If agent negotiate under the combinatorial negotiations, then the on-path offer can be characterized as follows (let  $\bar{S} \equiv \max_{(\theta_a, \theta_b) \in \Theta_a^n \times \Theta_b^n} S(\theta_a, \theta_b)$  be the maximum possible surplus):

- price a base package of *all* the low-cost items at  $p^{\text{base}} \equiv \frac{\bar{S}}{1+\delta}$ ;
- price *each* high-cost item at an additional  $p^{\text{add}} \equiv c_H - \frac{c_H - v_L}{1+\delta} \in (v_L, c_H)$ .

The buyer accepts, buying the base package (of all the low-cost items) and picking every additional item for which he has the high value.

In the rest of this part we further illustrate the role of Assumption 1. We show that it guarantees the richness of combinatorial negotiations (Lemma 3).

Recall that the surplus is

$$S(\theta_a, \theta_b) \equiv \sum_k \max\{\theta_{bk} - \theta_{ak}, 0\},$$

and let  $N^*(\theta_a, \theta_b) \subset N$  be the *efficient set of items to trade*

$$N^*(\theta_a, \theta_b) \equiv \{k \mid \theta_{bk} > \theta_{ak}\}.$$

We observe that the type spaces we focus on here satisfy the following condition:

**LEMMA 2** *Under Assumption 1, for any agent  $i$ 's type  $\theta_i \in \Theta_i^n$ , if two types of the other agent lead to the the same efficient set of items to trade, they must also lead to the same surplus. That is,*

$$N^*(\theta_i, \theta_{-i}) = N^*(\theta_i, \theta'_{-i}) \implies S(\theta_i, \theta_{-i}) = S(\theta_i, \theta'_{-i}); \quad \forall i, \theta_i, \theta_{-i}, \theta'_{-i}. \quad (14)$$

The following Lemma illustrates the importance of condition (14) in terms of allowing for rich negotiations (as defined in 3.3).

**LEMMA 3** *Combinatorial negotiations are rich if and only if condition (14) holds.*

**Proof of Lemma 3.** “if”: Suppose condition (14) holds. Consider any  $t$ , and history  $h^{t-1}$  such that all items are still available ( $N(h^{t-1}) = N$ ). Without loss of generality let us assume the offerer  $i(t) = a$  is Alice, the seller, with type  $\theta_a$ . Here we construct a fraction- $\kappa$  demanding offer ( $\forall \kappa \in [0, 1]$ ). This offer prices a set  $N^t$  according to the following:

- if  $N^t = N^*(\theta_a, \hat{\theta}_b)$  for some  $\hat{\theta}_b \in \Theta_b^n$ , then its price  $p(N^t) = \sum_{k \in N^t} \theta_{ak} + \kappa S(\theta_a, \hat{\theta}_b)$ ;
- otherwise,  $p(N^t) = \sum_{k \in N^t} \theta_{ak} + \kappa \bar{S}$ , where  $\bar{S}$  is the maximum possible realized surplus.

It follows from condition (14) that prices in the first case are well defined. In addition, we note that the pricing for the second part is not unique: we simply require that high enough prices are charged for those sets so that none of them would be picked by the responder.

It then follows from the definition of rich negotiations that the combinatorial negotiations are rich.

“only if”: Suppose condition (14) does not hold. Without loss of generality let us assume that  $i = a$  is Alice, the seller. The argument is similar for the opposite case. The violation of condition (14) implies that there exists some  $\theta_a \in \Theta_a^n$ , and  $\theta_b, \theta'_b \in \Theta_b^n$  such that  $N^*(\theta_a, \theta_b) = N^*(\theta_a, \theta'_b)$  and  $S(\theta_a, \theta'_b) > S(\theta_a, \theta_b)$ .

We aim to show that there exists no fraction- $\kappa$  demanding offer for the offerer  $a$ 's type  $\theta_a$ ,  $\forall \kappa \in (0, 1)$ : by definition, any such offer shall price  $N^*(\theta_a, \theta_b)$  at  $p(N^*(\theta_a, \theta_b)) = \sum_{k \in N^*(\theta_a, \theta_b)} \theta_{ak} + \kappa S(\theta_a, \theta_b)$  to meet the target divisions of surplus for the responder's type  $\theta_b$ , and similarly for the responder's type  $\theta'_b$  we have  $p(N^*(\theta_a, \theta'_b)) = \sum_{k \in N^*(\theta_a, \theta_b)} \theta_{ak} + \kappa S(\theta_a, \theta'_b) > p(N^*(\theta_a, \theta_b))$ . This violates the fact that  $N^*(\theta_a, \theta'_b) = N^*(\theta_a, \theta_b)$ . ■

Condition (14) is crucial to the “richness” of the negotiations, and thus contributes to existence “well-behaved” equilibria (Proposition 2). To illustrate this point, Example 5 presents a case that violates Assumption 1: there are three levels of valuations for each agent, though the frequencies are still known. For instance, when  $(0, 20, 40)$  meets  $(10, 50, 30)$  or  $(30, 50, 10)$ , efficiency requires to trade the first two items in both cases, but the surpluses are 40 and 60, respectively.

**EXAMPLE 5 (A “counter-example” with more than two levels of costs/values)**  $n = 3$ , and the buyer’s cost is  $(0, 20, 40)$  or its permutations, and the seller’s value is  $(10, 30, 50)$  or its permutations. Under the combinatorial negotiation, there exists no equilibrium that (1) achieves full efficiency and (2) divides the surplus according to the Rubinstein shares.

With three values, the particular matchups can lead to different combinations of surpluses, but the same trading of items. This allows people to posture on their valuations without changing which items trade, which then changes the incentives to negotiate and leads to inefficiencies.

**Proof of Example 5.** Suppose there exists such an equilibrium. Consider the following two realized matchups:

- (a)  $(0, 20, 40)$  meets  $(10, 50, 30)$ : surplus is 40, with first two items traded;
- (b)  $(0, 20, 40)$  meets  $(30, 50, 10)$ : surplus is 60, with first two items traded.

To achieve the Rubinstein shares, the equilibrium price for the first two goods should be  $p_{12} = 20 + \frac{40}{1+\delta}$  according to matchup (a), but  $p_{12} = 20 + \frac{12}{1+\delta}$  according to (b), which is not possible. ■