

A Theory of Efficient Negotiations

Matthew O. Jackson, Hugo F. Sonnenschein, and Yiqing Xing*

Draft: March 2016

Abstract

Negotiation involves determining not only an agreement's price, but also its content, which typically has many aspects. We model such negotiations and provide conditions under which negotiation leads to efficient outcomes, even in the face of substantial asymmetric information regarding the value of each aspect. With sufficient information about the overall potential surplus, if the set of offers that agents can make when negotiating is sufficiently rich, then negotiation leads the agents to efficient agreements in all equilibria. Furthermore, no "planner" or "mechanism designer" who knows the statistical structure of information is required: the same negotiation game works regardless of the setting. The theory and examples explore the anatomy of negotiation and may shed light on why many situations with significant asymmetric information exhibit little inefficiency.

Keywords: Negotiation, Bargaining, Exchange, Trade, Multiple Items, Linking, Contract Theory, Mechanism Design, Implementation, Efficiency, Sequential Equilibrium, Trembling.

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

*Department of Economics, Stanford University, Stanford, California 94305-6072 USA, and Department of Economics, University of Chicago, Chicago, Illinois, 60657. Jackson is also an external faculty member at the Santa Fe Institute and a fellow of CIFAR. Emails: jacksonm@stanford.edu, hfsonnen@uchicago.edu and xingyq@stanford.edu. Financial support from the NSF under grant SES-1155302 as well as ARO MURI award W911NF-12-1-0509. We thank seminar participants at Chicago, Keio, Monash, Peking, Penn State, Princeton, Stanford, U of Tokyo, Yale as well as the Decentralization conference, for helpful comments and suggestions.

1 Introduction

Negotiations are typically much more complex than bargaining: in addition to a price paid upon an agreement, the content of the agreement must also be determined. Moreover, the content being negotiated generally involves many aspects and substantial asymmetric information regarding agents' valuations for each dimension. Under such circumstances should we expect agents to reach efficient agreements, or even any agreement?

The mechanism design literature has shown that, in a nontrivial and important set of such multi-dimensional circumstances, (nearly) efficient agreements are possible. However, those optimistic results are derived assuming that the 'mechanism' through which agents communicate can be carefully tailored to the setting: often involving extremely specific restrictions on how agents can communicate. The mechanisms are carefully designed to balance incentives and rely on exact knowledge of the utility functions of the agents and the uncertainty that they face. Of course, in most settings an omniscient "planner" or "mechanism designer with such knowledge does not exist. Indeed, most negotiations are fairly free form, bearing almost no resemblance to the mechanisms used to prove results in the literature. Typically, one side offers a contract from a very general set of potential contracts (essentially, subject only to what is considered legal in the society), and then the other side either accepts or has a chance to amend the contract, and so forth. Under such a universal form of negotiation should we still expect agents to be able to reach efficient agreements?

As an example, there may be substantial uncertainty regarding how much value "Labor" attaches to one health package versus another, versus time off for child care, pension provisions, sick days, flexible schedule, safety rules, disability insurance, grievance provisions, wages, and so forth; and similarly how "Management" views the relative costs of each of these items. If each item was bargained over in isolation, following Myerson and Satterthwaite (1983) there would be substantial losses in efficiency, regardless of the way in which negotiation is conducted.¹ However, in aggregate, both Labor and Management may have a better idea of the possible gains from trade and take this into account when making their surplus demands and revealing their preferences. For instance, if there are many dimensions and enough independence in the valuations across dimensions then the agents may have a very good idea of the overall surplus from trade, even though they may have little idea of exactly how to realize that surplus. In such settings, if one can carefully design the mechanism through which agents communicate, then as shown by Jackson and Sonnenschein (2007) there exists a "linking mechanism" that results in (near) efficiency. That mechanism requires that each agent announce valuations across the various dimensions that match the expected

¹This presumes that valuations can overlap, and that agents are not forced to participate in an agreement.

frequency distribution. This gives one hope that efficiency may be reached, but we may not be convinced since this resolution assumes that agents are forced to only communicate via a “mechanism” that is somehow imposed by an unmodeled designer who knows the frequency distribution. What happens in a similar circumstance if the agents simply negotiate in a free-form alternating offer style that more closely matches real-world negotiations? As we prove, efficient outcomes still ensue (in all equilibria). We also show that this depends on the ‘free-form’ negotiations being sufficiently free-form: agents must be able to make certain kinds of offers in order to reach efficiency. If they are limited in what offers they can make, then they will fail to reach efficiency.²

This may help provide insight into an empirical puzzle pointed out by Kennan (2005). It is not curious that strikes exist, but it is curious that they are so rare. According to Kennan (2005), in the United States between 1948 and 2005, “idleness due to strikes never exceeded one half of one percent of total working days in any year.”³ In fact, since 1990 average lost time has been about 20 minutes per year per worker in the U.S.; and, even in a more strike-prone country like Spain, the number is less than 1/3 of a day per worker per year (again, according to Kennan (2005)).

Although the alternating-offer negotiation games that we consider can be viewed as “mechanisms”,⁴ the negotiation games introduced above do not use the distributional information about agents’ types. Thus, our work can be thought of in the broader spirit of Wilson’s (1987) criticisms of mechanisms that depend on agents’ (higher-order) beliefs. Satterthwaite, Williams, and Zachariadis (2014) also view such mechanisms as “impractical” as “[the agents’] beliefs are not a datum that is practically available for defining economic institutions” (p.249). In contrast, “robust” and “detail-free” have mainly been used in literature to refer to resolving 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). The exception is Matsushima (2008) who used “detail-free” with a meaning more similar to ours in an auction environment. To avoid confusions in terminologies we use “universality” to capture the feature that a protocol/mechanism is not defined based on any the knowledge of the prior distribution or utility functions of the agents.

²This is in line with practical advice given to practical negotiators. For example, Fisher and Ury 2001 emphasize putting multiple goods on the table and offering the other party many options.

³Kennan notes similar numbers for Canada, where he mentions that lost work time was about a third of a day per year per worker. Kennan also states that “Although the data are not readily available for a broad sample of developed countries, the pattern described above seems quite general: days lost due to strikes amount to only a fraction of a day per worker per annum, on average, exceeding one day only in a few exceptional years.”

⁴There is an extensive mechanism design literature characterizing the feasibility of efficient allocations. See, e.g., Jackson (2001,2003), Segal and Whinston (2012), and the references therein for the static case; and Skrzypacz and Toikka (2014) for the dynamic case.

As a preview, we proceed as follows. We start with several illustrating examples (Section 2). Next, (in Section 3) we present our formal model of multi-item negotiations and introduce several ‘universal’ negotiation mechanisms. After having introduced the universal mechanisms, (in Section 3) 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 negotiation mechanism is rich enough then all equilibria result in fully efficient outcomes, despite the universality of the mechanism and the fact that the agents only know the total available surplus and not how to realize it. We show that this results hinges on the universal negotiating mechanism being rich enough. Agents must be able to propose overall agreements and demand shares of the overall surplus. The intuition behind this result is roughly as follows. 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 method of negotiation thus aligns incentives, helping the agents find the right agreement quickly and efficiently. After that, (in Section 5) we extend the discussions to the case in which the surplus is only approximately known. This introduces some substantial technical hurdles as the set of sequential equilibria (and the set perfect Bayesian equilibria) are not upper-hemicontinuous. Slight amounts of uncertainty lead to many equilibria that rely on extreme updating of beliefs (that survive the usual refinements). We show that introducing slight trembles eliminates those problems and restores continuity at the limit.⁵ We conclude with some examples that address the case in which there is substantial uncertainty about the total available gains from trade.

2 Examples

We begin with an example of negotiations over multiple dimensions. In this example, if only one dimension is considered at a time, then inefficiency necessarily results in any mechanism that is (interim) incentive compatible, individually rational, and budget balanced. However, if negotiations are over all dimensions together then efficiency ensues.

2.1 An exchange of private goods

Alice has n rugs and she can sell some (or all) of them to Bob. θ_{ik} is $i \in \{a, b\}$ ’s value for the k -th rug. Suppose $\theta_{ak} \in \{0, 8\}$ and $\theta_{bk} \in \{2, 10\}$, both equally likely; i.i.d. across rugs and agents.

A decision is a list of which rugs trade at which prices at which times: $\{(N_t, p_t)\}_t$ indicates

⁵This is different from trembling hand perfection as we hold trembles constant and let the uncertainty vanish, rather than the alternative.

that the rugs with indices in the set N_t trade at a total price $p_t \in \mathbb{R}$ (the amount transferred from Bob to Alice).

Agents' (time-0) utilities from such a decision are

$$u_a(\{(N_t, p_t)\}_t) = \sum_t \delta_a^t \left(p_t - \sum_{k \in N_t} \theta_{ak} \right)$$

and

$$u_b(\{(N_t, p_t)\}_t) = \sum_t \delta_b^t \left(\sum_{k \in N_t} \theta_{bk} - p_t \right),$$

where $\delta_a, \delta_b \in [0, 1)$ are the factors at which agents discount the future.

Beyond exchanging rugs, other applications may be two firms negotiating on a contract which includes specifications of a service to be provided, or a good to be produced, and this could include many aspects which need to be specified, such as the time to production, specs of the item (e.g., its weight, performance, durability), the quality of the object, penalties for failure to deliver, and so forth. Or for example it could be a contract between a faculty member and a university, specifying a teaching load, a sabbatical policy, research funding, summer support, a salary, and so forth.

The efficient ('first-best') outcome is to have a rug traded if and only if Bob's value exceeds Alice's; i.e., $(\theta_a, \theta_b) \in \{(0, 2), (0, 10), (8, 10)\}$; and all such rugs are traded at $t = 0$. The 'surplus' is defined as the gain from trade under the efficient outcome, with a per-item expectation of $0.25 \times ((10 - 0) + (10 - 8) + (2 - 0)) = 3.5$.

With a single item: efficiency is not achievable. When there is only one rug ($n = 1$), efficiency cannot be achieved under the requirements of (interim) incentive compatibility and individual rationality (the price lies between the two valuations), (noting that budget balance is satisfied by construction) - which is a variation on the results of Myerson and Satterthwaite (1983).

Indeed, one can verify that the following mechanism maximizes the joint (second-best) surplus subject to incentive compatibility and individual rationality. In the table, q is the probability of trade and p is the price.

	$\theta_b = 10$	$\theta_b = 2$
$\theta_a = 0$	$q = 1, p = 5$	$q = \frac{5}{6}, p = 2$
$\theta_a = 8$	$q = \frac{5}{6}, p = 8$	$q = 0, \text{ no trade}$

Note that this is a direct mechanism: agents communicate their valuations and then the outcome is enforced by the mechanism. The mechanism is designed based on the agents'

utility functions and the probabilities of the various valuations. If those valuations or probabilities differ slightly from what is anticipated by the mechanism designer, the mechanism might lead to very different outcomes.

Negotiations with Four Items Now, let us consider ‘universal’ mechanisms that are not restricted based on the setting. Agents may negotiate in various manners. Let us compare two benchmark ways in which agents may carry out their negotiation.

Item-by-Item Negotiations The agents negotiate over the n items, but the negotiation is independent across items - so each item is negotiated upon via its own sequence of offers and counter-offers (the naming of any price for an item). It is possible that agents may tie the negotiation together across items, but only via their equilibrium actions, as it is simply n separate bargaining processes that are conducted in parallel.

Combinatorial Negotiations Here we move to the opposite extreme: instead of negotiating item-by-item, agents discuss and “price” all possible subsets of items. An offer is a list of prices - one for each possible subset. If the responder accepts the offer, s/he picks the subset of items he/she desires to trade at the price offered in the list. If the responder rejects an offer, s/he makes a counter-offer by naming a list of prices for each subset of items (after one period of discounting).

These two different forms of negotiations result in very different levels of efficiency, as illustrated in the following example.

EXAMPLE 1 *Consider a case with $n = 4$ rugs, and to keep the setting simple, suppose that each of the four possible matchups $(0, 2)$, $(0, 10)$, $(8, 2)$, $(8, 10)$ appears exactly once - and the 24 different possible orderings of these four matchups are all equally likely. Thus, agents know the total gains from trade, but they do not know which items to trade. The seller has value 8 for two of the items, but does not know which one will trade, and similarly, the buyer has value 2 for two of the items and does not know which one will trade.*

1. *there exists a direct mechanism under which truth is an equilibrium (in fact weakly dominant).*
2. *under an item-by-item negotiation, none of the (weak) perfect Bayesian equilibria is efficient.*
3. *under a combinatorial negotiation, all (weak) perfect Bayesian equilibria are efficient.*

Point 1 is seen as follows.⁶ Here, have trade occur at a fixed price of 15, and have agents simultaneously announce their valuation vectors and execute the efficient trades. The key is that agents are restricted to announce exactly two 0's and two 8's for Alice, and two 2's and two 10's for Bob. The only choice is which items on which to list which valuations. This is incentive compatible.

Point 2 can be seen roughly as follows. Let's consider Alice and the two items that of value 0. In order to trade efficiently, both prices she quotes in the first period would have to be no more than 2, as otherwise Bob will certainly reject one of them. Instead, she has an incentive to charge higher prices to get Bob to trade the higher value item first, and then to later trade the other item. The proof that this is true of all equilibria can be seen in the appendix (see the proof of Example 2)

Point 3 follows from our results below. The intuition is as follows. By being able to negotiate over combinations of goods, the problem of trying to screen items is circumvented. Alice can quote an overall price for various groups of three items (the two she values at 0 plus either of the ones she values at 8) - knowing that the total cost of the goods that should efficiently trade is 8 and the total value to Bob is 22. The equal split price would be 15, but under alternating offer bargaining she will ask for a higher price depending on the discount factor. Effectively she avoids having to detect which of the goods Bob has the 2 on and the 10 on, since she can quote a price for the overall group. Showing that all equilibria satisfy this involves an argument that any alternative must have some inefficiency that could be improved by a mutually beneficial deviation.

The above example shows the sharp contrast between the performances of the two benchmark forms of universal negotiations. Under an item-by-item negotiation, the result is very similar to the case with one single item, and Myerson-Satterthwaite result of inefficiency applies. In contrast, the combinatorial negotiation provides agents the chance to negotiate on "overall" terms. This comparison is formalized in Section 3, and a general condition is provided to characterize efficient negotiations.

This result is reminiscent of the bundling literature, where the ability to sell packages of goods can be welfare improving under some circumstances. Here we show that full efficiency holds in a negotiating setting with certain knowledge surplus, in all equilibria, and we identify which sorts of negotiations admit such full efficiency.

The above example is extreme in that the surplus is exactly known. Such an example can be viewed as a limit case of having many independent items, where by Law of large numbers each of the four possible matchups appears on roughly a quarter of the items. For our formal analysis, we first focus on the case where surplus is known (Section 4), and then

⁶See also the discussion of a similar example in the working paper version of Jackson and Sonneschein (2007).

extend the discussions to the case where surplus is approximately known (Section 5).

2.2 Other applications

Our framework of negotiations covers applications other than the exchange of private goods.

2.2.1 Public Good Provision / Household Investments

Alice and Bob decide whether to invest in n different projects that yield benefits to both agents and which cost $c = 6$ in total per project. Agent $i \in \{a, b\}$'s benefit from project k is $\theta_{ik} \in \{0, 8\}$, which is random and equally likely; we also assume independence across agents' types. Agent i 's values are privately known to agent i at the time of collective decision making.

Possible choices for each item has $x_k \in X = \{1, 0\}$ and with $x_k = 0$ representing not investing in the k -th project and $x_k = 1$ representing investing. Each agent pays a cost of 3 if they decide to invest in the project.

Agents' preferences are quasi-linear in money, and additively separable across items. So their net utilities (as functions of a joint decision x and a transfer p from Bob to Alice) is $u_a(\theta_a, x, p) = \sum_k x_k(\theta_{ak} - 3) + p$ and $u_b(\theta_b, x, p) = \sum_k x_k(\theta_{bk} - 3) - p$. Agents discount the future with factors $\delta_a, \delta_b < 1$.

2.2.2 Task Assignment / Household Chores

Two workers, Alice (a) and Bob (b) have n tasks or chores to complete.

Agents each get a benefit of 6 from a completed task. An agent is of one of two types for each task, $\theta_{ik} \in \{0, 1\}$, where 1 represents that the agent is competent to work on the task, and 0 representing incompetence. If both agents are competent then they can both work on the task.

Possible choices for each task has $x_k \in X = \{a, b, ab, 0\}$ and with $x_k = a$ representing that a works on it alone, $x_k = b$ representing that b works on it alone, $x_k = ab$ representing that they work on it together, and $x_k = 0$ representing not completing the task.

A task costs 10 if completed by one competent agent, and costs a total of 2 (1 for each) if completed jointly by two competent agents, and the task is not completed if an incompetent agent works on it.

Agents' preferences are quasi-linear in money, and additively separable across items; and agents discount the future with factors $\delta_a, \delta_b < 1$.

2.2.3 Three Examples in One

Although the three applications introduced in this section seem quite different, they share a similar mathematical structure. In particular, each example has two agents, each of whom has two types. The two types can be relabelled as h and l , where h representing the type that generates more surplus (under the optimal decision).⁷

The surplus S is generally defined as the sum of the agents' net utilities under the optimal decision (depending on agents' types). The three examples then have the same surplus as a function of the relabelled types:

S_k	$\hat{\theta}_b = h$	$\hat{\theta}_b = l$
$\hat{\theta}_a = h$	10	2
$\hat{\theta}_a = l$	2	0

If both agents have type l , then efficiency requires the null choice to be taken (no trade, no investment, no task).

Due to the above similarity, the results discussed in the context of the exchange application (Section 2.1) also hold for the other two applications. In particular, the efficiency is not achievable when there is a single item (with equally likely independent types). In contrast, with multiple items, as in Example 1, there exist efficient mechanisms. Whether efficiency can be achieved in free-form negotiation depends crucially on how agents negotiate - it fails if agents negotiate one decision at a time, but holds in *all* equilibria if they can negotiate agreements over all dimensions with one overall transfer payment.

2.3 Features in Multi-Item Negotiations

In negotiations with multiple items, with sufficient independence as the number of items becomes large, uncertainty about the total surplus may vanish. However, there are still two forces that push in different directions regarding whether we get efficiency as we have many items for negotiation:

- There is less uncertainty about overall surplus.
- There remains substantial uncertainty about which decision to make on each dimension.

Thus, it is still unclear whether the results from the above examples will generalize. Should we expect some universal negotiation protocols to lead to nearly efficient outcomes in a

⁷In particular, for the application 2.1 let $\hat{\theta}_{ik} = h$ if $\theta_{ik} = 8$, i.e. the type with the larger benefit; for the example 2.2.2 let $\hat{\theta}_{ik} = h$ if $\theta_{ik} = 1$, i.e. the competent type; whereas for the example 2.1 $\hat{\theta}_{ak} = h$ if $\theta_{ak} = 0$, and $\hat{\theta}_{bk} = h$ if $\theta_{bk} = 10$, i.e. the lower-cost type for the seller and the higher-value type for the buyer. Let $\hat{\theta}_{ik} = l$ represent the remaining type in each example.

general set of such settings - and if so what features does the negotiation protocol have to have?

We introduce the general model before answering these questions.

3 Multi-Item/Aspect Negotiations: The Model

3.1 Multiple Aspects and Decisions

A multi-aspect negotiation problem consists of:

- two agents, Alice a and Bob b ,
- a joint decision to be made that concerns a finite number, n , of *items* (or *decisions*, *objects*, *aspects*, *tasks*, *etc.*),
- X a choice space for each item, and X^n the space of joint choices, with representative element $x = (x_1, \dots, x_k, \dots, x_n)$, and
- monetary transfers may be made from Bob to Alice and a transfer in period t is denoted $p_t \in \mathbb{R}$.

Our framework also covers settings where transfers are precluded. Below we show that approximate efficiency can be obtained even in such no-transfer environments, provided that there are enough items to be considered. The case with transfers allows for a less cluttered analysis, and so we presume that transfers are allowed, unless otherwise noted, returning to the case without transfers at the end.

3.2 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 Θ , with f_i denoting the marginal of f on Θ_i^n , and
- discount factors $\delta_i \in (0, 1)$, $i \in \{a, b\}$ that are 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 θ_{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). Our main results include conditions on f under which efficient negotiation holds.

For now, we assume that the agents' payoffs across items are additively separable, but this is not essential to the analysis (see Section 4.4).

We allow decisions over different items to be made 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. In this case, let $N_t \subset \{1, \dots, n\}$ be the subset of problems whose decisions are made at period t s.t. $N_t \cap N_s = \emptyset, \forall s \neq t$, and recall that p_t is the transfer made (from Bob to Alice) at t . Then the agents' time-0 utilities are

- for Alice: $U_a = \sum_t \delta_a^t (\sum_{k \in N_t} u_a(x_k, \theta_{ak}) + p_t)$;
- for Bob: $U_b = \sum_t \delta_b^t (\sum_{k \in N_t} u_b(x_k, \theta_{bk}) - p_t)$.

Discounting can be interpreted in at least two (standard) ways. 1) Utilities from decisions made in period t are realized in period t . This interpretation applies to the example of strikes between a firm and a union, where the costs and benefits are held until the employment relationship is restored (i.e., the agreement is reached, at t), to the task-assignment example where the task is done (and completed immediately) at t , and/or to an exchange setting where the seller produces the goods and the buyer consumes them at t . 2) Particular to an exchange setting: 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. θ_{sk} is the time- t value of flow payoffs the seller could get from item k .

The welfare from a decision x_k for an item k is $u_a(x_k, \theta_{ak}) + u_b(x_k, \theta_{bk})$. The social surplus from an efficient decision is

$$S(\theta_{ak}, \theta_{bk}) \equiv \max_{x_k} u_a(x_k, \theta_{ak}) + u_b(x_k, \theta_{bk}).$$

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

$$S(\theta_a, \theta_b) \equiv \max_{x \in X^n} \sum_k [u_a(x_k, \theta_{ak}) + u_b(x_k, \theta_{bk})].$$

We assume throughout that for any $(\theta_{ak}, \theta_{bk}) \in \Theta_s \times \Theta_b$, $S(\theta_{ak}, \theta_{bk}) \geq 0$, so that for each possible combination of types (for each item) there is some decision that is not painful to the agents (which could include, not trading, or not doing any tasks, etc., depending on the circumstances), and that $S(\theta_{ak}, \theta_{bk}) > 0$ for some $(\theta_{ak}, \theta_{bk}) \in \Theta_s \times \Theta_b$, so there is, at least potentially, non-trivial welfare gain. With a finite set of types, this is essentially just a normalization.

To achieve efficiency, the optimal choices (those maximizing the joint surplus) should be made for all items, and at $t = 0$.

3.3 Benchmark Negotiation Protocols

This section presents several examples of ways in which agents may negotiate.

These bargaining protocols are all ‘universal’ in the sense that the game forms are independent of agent’s utility functions

3.3.1 Item-by-Item Negotiations

This is a formal description of the item-by-item negotiations introduced in Subsection 2.1, where each item is independently negotiated via Rubinstein-Stahl alternating offer bargaining. In particular:

- One of the agents, say Alice, announces a string of decisions $(x_1, p_1; x_2, p_2; \dots; x_n, p_n)$, with p_k being the transfer associated with item k .
- The other agent, say Bob, can accept any subset of the decisions, which are then implemented with the agreed transfers.
- If for any items a decision has not been made, we start again with the roles of the agents reversed (and one period of discounting ensues). Bob then announces a decision (x_k, p_k) for each item that still remains undetermined.
- Alice can accept any subset of those remaining items, which are then dealt at Bob suggested decisions.
- If any items are undecided, we start again with the roles reversed. We continue in this manner indefinitely or until all goods are traded.⁸

As examples, 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 previous protocol. The current protocol also applies when a wife and a husband are discussing household consumption, where different goods can be invested in and consumed at different times.

In contrast, we note that in some cases, with negotiating over a contract with many aspects, this sort of negotiation may raise issues of feasibility and interpretation if agents cannot consume any until all of the aspects are agreed upon. For instance, an employment contract would have to specify wages, pension plan, hours, holidays, etc., and employment might not be feasible until all of the aspects were 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 are possible, as described below.

⁸Items that remain untraded indefinitely simply result in 0 gains from a decision.

3.3.2 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 combinations of decisions.

- One of the agents, say Alice, names a transfer for each possible (joint) choice $x \in X^n$. Formally, an offer is a mapping $p : X^n \rightarrow \mathbb{R}$.
- The other agent, say Bob, accepts or rejects.
- If accepted, Bob selects a choice x , and transfers the amount $p(x)$ to Alice.
- If rejected we start again with the roles of the agents reversed (and one period of discounting ensues).

This protocol provides an important theoretical benchmark: in terms of choices, this protocol allows for the richest offer space.

The richness in message space of the combinatorial protocol can be a disadvantage in practice: the size of an offer is $|X|^n$, exploding at exponential rate as n grows. For instance, in an exchange example with $n = 20$ rugs, each offer needs to specify $2^{20} \approx 1048k$ prices, which is not very realistic in practice. Instead, agents tend to use “reduced forms” like the ones introduced in 3.3.3 and 3.3.4 where a much smaller message spaces suffices to convey the essential information.

3.3.3 “Overall negotiation”: Rubinstein-Stahl negotiation with many items

We now present an intermediate universal form of negotiation, in which agents negotiate in terms of 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 offer space from $2^n + 1$ to $n + 1$ - agents announce how much they value each rug and then a net gain in utility from the transfer price that is needed. They then allow the other agent to choose the trades.

- One of the agents, say Alice, announces (not necessarily truthfully) her types $\hat{\theta}_a = (\hat{\theta}_{a1}, \dots, \hat{\theta}_{an})$ and demands a payoff of $v_a \in V$, from some compact set V of possible total values.
- The other agent, say Bob, accepts or rejects.
- If accepted, Bob picks a joint decision, $x = (x_1, \dots, x_n) \in X^n$, and delivers a net payoff equal to v_a based on Alice’s announced values: so, the total transfer given to Alice by Bob is equal to $p = v_a - \sum_k u_a(x_k, \hat{\theta}_{ak})$. The game ends.

(In the case where Bob announced $\hat{\theta}_b = (\hat{\theta}_{b1}, \dots, \hat{\theta}_{bn})$ and demanded a payoff $v_b \in V$, Alice picks x and the transfer made by Bob is $p = \sum_k u_b(x_k, \hat{\theta}_{bk}) - v_b$.)

- If rejected we start again with the roles of the agents reversed (and one period of discounting ensues).

The above procedure has clear correspondences in practice. For instance, in the exchange of private goods example, a seller (Alice) claims her costs for the goods and demands a additional gain in payoff; and a buyer (Bob) if accepting an offer chooses which goods to buy at a price that exceeds the sum of the seller’s reservation values (for those selected goods) by the demanded margin. In a task assignment problem, an agent (say Alice) claims her costs for the tasks and targets a total cost she wants to deliver, and the other agent can decide who should work on which task but must subsidize Alice if the sum of costs for the tasks assigned to her exceeds Alice’s target.

3.3.4 Negotiation over frequencies

As another example of reduced-form negotiations, agents may negotiate over frequencies that are then used in a second stage (voluntary) game. This also fits with many settings in which people bargain over “basic terms”, and then after they have reached a tentative agreement then they fill in details. In particular, a “negotiation over frequencies” consists of two phases:

Phase 1 (alternating offers of games characterized by frequencies):

- The offerer, $i \in \{a, b\}$, quotes a frequency distribution $\widehat{\phi}_i^n \in \Phi_i^n$,⁹ and a target payoff $v_i \in V$.
- The recipient accepts or rejects.
- If accepted we move to Phase 2.
- If rejected we begin Phase 1 again with the roles of the agents reversed (and one period of discounting ensues).

Phase 2 (the game is played):

- The offerer i announces $\widehat{\theta}_i^n \in \Theta_i^n$ that has a frequency distribution $\widehat{\phi}_i^n$.
- Either the recipient picks a decision $x \in X^n$, and makes a payment of at least $v_i - \sum_k u_i(x, \widehat{\theta}_{ik}^n)$ to the offerer or says “No” and no decision is made.
- The game ends.

For instance, in a task-assignment example, to quote a frequency is to specify the number of tasks an agent is competent on; and in an exchange example a buyer can quote for how many of the rugs she has high values.

⁹ Φ_i^n is the set of possible frequencies (with n items). For instance, one frequency could be “ $\frac{1}{3}$ 0’s and $\frac{2}{3}$ 8’s”. Note that the quoted frequency $\widehat{\phi}_i^n$ may differ from i ’s true frequency.

The message space in negotiations over frequencies is the smallest among all the four ways of negotiations introduced so far: the space of frequencies has a size of less than $n^{|\Theta_i|-1}$, which can be much smaller than the space of types, let alone the space of all possible choices.

3.4 A General Definition of Alternating-Offer Negotiations

With these various examples of alternating-offer negotiation protocols in hand, we provide a definition of what we mean in general by an *alternating-offer negotiation*, Γ , with n items.

- 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 recipient), responds from a set $A_j^0(a_i^0)$, with a generic response denoted by a_j^0 .
- As a function of a_i^0, a_j^0 , agents agree on some items $N_0(a_i^0, a_j^0) \subset \{1, \dots, n\}$ (possibly the empty set), on which choices $x_{N_0}(a_i^0, a_j^0) \in X^{|N_0|}$ are made, and some transfer $p_0(a_i^0, a_j^0)$ is made from Bob to Alice.
- One period of discounting ensues.
- ...
- Inductively, in period t , agent $i(t)$ makes offers from a finite set $A_{i(t)}(h^{t-1})$ which could depend on $h^{t-1} \equiv (a_{i(0)}^0, a_{j(0)}^0, \dots, a_{i(t-1)}^{t-1}, a_{j(t-1)}^{t-1})$, the full history of negotiations through the last period.¹⁰
- The other agent $j(t)$ reacts from a set $A_{j(t)}(h^{t-1}, a_{i(t)}^t)$.
- As a function of $h^t = (h^{t-1}, a_{i(t)}^t, a_{j(t)}^t)$, agents agree on some of the items that remain, $N_t \subset N \setminus (\bigcup_{s < t} N_s)$, choices x_{N_t} over those are made, and a transfer p_t is made from Bob to Alice.
- This continues as long as there are goods remaining to be traded.

The negotiations introduced in 3.3.1 - 3.3.3 are all examples of the above general definition. In addition, negotiations over frequencies (3.3.4) can be viewed as a special case of a slightly generalized version of the above definition, where the decisions ($N_t(\cdot), x_{N_t}(\cdot), p_t(\cdot)$) are determined via a second phase of announcements after the acceptance.

We focus on the “alternating-offer” protocols, in which the seller offers at $t = 0, 2, 4, \dots$ and the buyer offers at $t = 1, 3, 5, \dots$. The results below extend directly, except with changes to the expressions for the split of the surplus, for other alternation patterns.¹¹

¹⁰For instance, only some items might still need to be negotiated, and a protocol might condition upon that set. This specification also allows for various rules for choosing the proposer - an easy extension is to allow for random recognition of $i(t)$.

¹¹For instance, if the seller makes all of the offers then the seller will get all of the surplus. The pattern

3.5 Universality

The protocols introduced in Sections 3.3.1 - 3.3.4 are “universal”, in the sense that the same game forms will result in efficient equilibria across many environments (specifications of preferences and 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 nobody who would actually know all the relevant statistical details of the setting and also be able to impose a mechanism knowledge before the agents play. Here we find that even when agents negotiate in much more general ways, they can still reach efficiency, something that has not been investigated previously in the broad mechanism design literature.

3.6 Equilibrium

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

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

A *belief system* for agent i is a function $\tilde{f}_i : H \times \Theta_i \rightarrow \Delta(\Theta_{-i})$ 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 collections of the opponent’s types) E_{-i} , conditional on a history h and the agent’s own type θ_i . To capture the idea that these beliefs apply to nodes in the game, we require that a belief system only place positive probability on θ_{-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.

Agent i ’s *strategy*, σ_i , specifies a distribution over the current action space, $\sigma_i(h, \theta_i) \in \Delta(A_i(h))$, at each node $(h, \theta_i) \in H_i \times \Theta_i^n$.

Beliefs are *consistent* if for each i and θ_i they correspond to a conditional distribution (relative to the common prior f) at almost every h in the support of $\sigma_{-i}, \sigma_i(\theta_i)$.¹²

of alternation must be either known in advance or random, but not depend on the history of the game.

¹²The usual definitions of consistency apply to finite action spaces, whereas we allow for games with continuum of actions and thus conditional probability measures may need to be defined via Radon-Nikodym derivatives and so are only tied down up to sets of measure 0.

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

A strategy profile σ satisfied sequential rationality (relative to a belief system \tilde{f}) if σ_i maximizes $U_i(\sigma_i, \sigma_{-i}, \tilde{f}_i, h, \theta_i)$ for each i , θ_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.

Given that sequential equilibria are difficult to define for games with continua of actions, and that we are proving results that hold for all equilibria (and so a weaker concept leads to stronger results), working with the concept of weak perfect Bayesian equilibria - adapted to continuum games - makes sense here.

4 Multi-Aspect Negotiations with Commonly Known Surplus

We first focus on a case in which the surplus is commonly known. We discuss cases with unknown surplus in later sections.

4.1 Known Surplus

A negotiation problem (n, u, X, Θ, f) (as defined above) has a known total surplus $\bar{S} > 0$ if there exists $\bar{S} > 0$ such that

$$S(\theta_a, \theta_b) \equiv \max_{x \in X^n} \sum_k [u_a(x_k, \theta_{ak}) + u_b(x_k, \theta_{bk})] = \bar{S}$$

for every $(\theta_a, \theta_b) \in \Theta$.

The limiting case in which the overall surplus is commonly known serves as a proxy. There are many justifications for this case, but let us mention two of them.

One is that 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 a clear intuition, while in the face of uncertainty the growing strategy spaces as the number of items gets large make the arguments more complex - an issue which we handle separately below.

Second, the known surplus case captures the idea that the two agents negotiating know enough about the essentials to negotiate well. For example, consider a partnership with ten tasks to do. The agents might not know the full surplus exactly, but do know a couple of critical things: (a) it is either necessary, or at least 'fair', to have each partner do half of the tasks (e.g., due to time and labor constraints). (b) they each know that they will have

some preference for one task over another, or some ranking over tasks. They might not know the exact monetary equivalents of intensities, but do expect that they can each rank the tasks from 1 to 10. A negotiation would be that Alice states her ranking over tasks, and then allows Bob to choose who does which ones, but subject to giving Alice an average ranking of 4. So, if Bob wants to do Alice's favorite task, then he has to compensate her by giving her those ranked 2, 3, 4, 5, 6 in her ranking order to average to 4. But if Bob picks Alice her 1, 2, 4, 6, 7, and so forth. With ten tasks, they can assign tasks pretty well. If they just randomly assigned tasks the average ranking would be 5.5. If they happened to have completely disjoint five favorites, they would each average a 3 ranking (we each get our 1,2,3,4,5). Asking to average a 4 means that they have to do better than random, but not expecting to match up perfectly - a quite reasonable and high probability possibility with ten tasks. Here, we stylize this by assuming that they know that their types will matchup so that they can each get an average ranking of 4. Thus, although known total surplus is a limiting case, it still captures the important idea that agents have a belief that the probability that they should be able to reach the expected surplus is better than just on one draw.

4.2 An Efficiency Result

Our first result show that under any negotiation game introduced in (3.3.2) or (3.3.3) , despite the substantial uncertainty about the value and efficient choice for any given item, *all* weak perfect Bayesian equilibria lead to to immediate and efficient decisions and a unique division of the total surplus.

THEOREM 1 *If a negotiation problem (n, u, X, Θ, f) has a known surplus $\bar{S} > 0$, then in all weak perfect Bayesian equilibria of the negotiation protocols introduced in Sections 3.3.2 or 3.3.3:*

- *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., $\frac{(1-\delta_b)\bar{S}}{1-\delta_b\delta_a}$ for Alice, and $\frac{\delta_b(1-\delta_s)\bar{S}}{1-\delta_b\delta_a}$ for Bob.*

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

We again emphasize that both of these negotiation games are 'universal' in that the above result above holds for exactly the same protocol for *any* u, Θ, f with known surplus. This distinguishes the result from a mechanism-design approach in which the mechanism is tailored to the f and the mechanism would have to change with u, Θ, f . Thus, this is not only distinguished because we are taking a positive perspective (examining a mechanism which

seems ‘natural’ in terms of how people actually negotiate) as opposed to a normative one (using direct mechanisms to prove that efficiency is possible if the ‘designer’ has sufficient knowledge), but also because the negotiation games are simple and directly adapt with the environment.

The intuition behind Theorem 1 is as follows. If there were any inefficiency on the anticipated equilibrium path, then since the agents know the potential surplus and can make demands for shares 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 extension of that by Shaked and Sutton (1984).

4.3 The Structure of the Negotiation Game Matters

One may conjecture that with known overall surplus, the above result about efficient decisions would extend to any negotiation game. This is not the case. Example 1 (Section 2.1) already serves as a counter-example: When agents follow the item-by-item negotiation, none of the equilibria are efficient even when the surplus is exactly known.

Below we provide another example. This example is such that it should be easiest to reach efficiency since it is commonly known that all items should trade and there is only one-sided uncertainty about whether trade should occur.¹³ Nevertheless, even in such a simple environment, *none* of the sequential equilibria are efficient with item-by-item negotiation.

EXAMPLE 2 *Consider an exchange example (following Subsection 2.1) in which agents negotiate over $n = 2$ items under the item-by-item protocol. Alice, the seller, has type $(\theta_{a1}, \theta_{a2}) = (0, 0)$, and there are two equally likely type profiles for Bob, the buyer: $(\theta_{b1}, \theta_{b2}) = (2, 10)$ or $(10, 2)$. In all equilibria at most one of the two items is traded in the initial period.*

The uncertainty about specific matchups distorts the agents’ incentives in the item-by-item negotiation game: the inefficiency in Example 2 derives from each agent’s incentives to screen the the other’s type to try to obtain a better price on any given item. Such incentives are mitigated when agents can negotiate on items “overall”, as agents have better knowledge about the overall surplus. Instead, in protocols that entail separate negotiation, even though agents can coordinate their actions across items they cannot take advantage of the greater knowledge that they have about the overall value of agreement as they cannot make offers that are contingent on combinations of items and so the ensuing selection problems cannot be overcome.

¹³For reviews of (single-item) bargaining with one-sided uncertainty, see, e.g. Fudenberg, Levine and Tirole (1985) and Ausubel, Cramton, and Deneckere (2002).

The above insights, although presented with the private-good exchange example, clearly apply generally to the applications of our framework.

4.4 A General Feature of a Negotiation Game that Ensures Efficiency

The negotiation protocol 3.3 allows the agents to negotiate in an integrated manner that takes advantage of their knowledge of the overall surplus, while item-by-item negotiation does not. There is a general sense in which any protocol that admits such integration leads to full efficiency, in all equilibria, as we now formalize.

Again, consider a negotiation problem (n, u, X, Θ, f) with known surplus \bar{S} .

DEFINITION 1 (SHARE-DEMANDING OFFERS AND PROTOCOLS) *An alternating offer negotiation Γ includes a share- v demanding offer in some period t , for some $i(t)$, $\theta_{i(t)} \in \Theta_i^n$, $v \in [0, \bar{S}]$, and history h^{t-1} , if there exists $a_i \in A_{i(t)}(h^{t-1})$ such that¹⁴*

- *for every $a_j \in A_{j(t)}(h^{t-1}, a_{i(t)}^t)$ either there is no agreement on any items, or the realized payoff for $i(t)$ in the current period is at least v , and*
- *for any $\theta_{j(t)}$ for which $f(\theta_{i(t)}, \theta_{j(t)}) > 0$: there exists $a_j \in A_{j(t)}(h^{t-1}, a_{i(t)}^t)$ for which the realized payoff in the current period for $i(t)$ is v and the realized payoff in the current period is $\bar{S} - v$ for $\theta_{j(t)}$ for $j(t)$ with a type $\theta_{j(t)}$.*

An alternating offer negotiation includes share-demanding offers for some set $V \subset [0, \bar{S}]$, if at any point of the protocol through which no agreement has yet occurred, the current offerer $i(t)$ has a share- v demanding offer for each $v \in V$ and type $\theta_{i(t)}$ for which $f_{i(t)}(\theta_{i(t)}) > 0$.

When there is no ambiguity, we say that an alternating offer negotiation “includes share-demanding offers” if it includes share-demanding offers for $[0, \bar{S}]$.

THEOREM 2 *If a negotiation problem with n items has a known surplus $\bar{S} > 0$ and the alternating offer negotiation Γ includes share-demanding offers, then in all weak perfect Bayesian equilibria:*

- *agreement is reached immediately,*
- *the full surplus is realized, and*
- *the agents’ expected payoffs equal to their Rubinstein shares; i.e., $\frac{(1-\delta_b)\bar{S}}{1-\delta_b\delta_a}$ for Alice, and $\frac{\delta_b(1-\delta_a)\bar{S}}{1-\delta_b\delta_s}$ for Bob.*

¹⁴Payoffs expressed here are not-discounted; i.e., they are evaluated in the current period.

Theorem 1 is a corollary to Theorem 2, since both negotiation games (Sections 3.3.3 and 3.3.2) have share-demanding offers. In particular, for the negotiation in Section 3.3.3: at any point of the game and for any share v , a share- v demanding offer is (θ_i, v) - the current offerer lists the types truthfully and demands a total share of v . Such an offer, once accepted, gives the offerer exactly payoff of v regardless of the responder's decisions, and gives the responder $\bar{S} - v$ if the responder chooses the efficient choices, given the offerer's (listed) types and own types.

As for the combinatorial negotiation (Section 3.3.2): actually any offer available in the negotiation from Section 3.3.3 has an equivalent offer in the combinatorial negotiation.¹⁵ In this sense, the combinatorial negotiation has a “richer” message space. As a result, the combinatorial negotiation also has share-demanding offers. In general, any expansion of the offers space (in which the responder can only accept one offer), still has share-demanding offers and so results in efficiency.

Moreover, the combinatorial negotiation has the advantage of allowing for a general payoff structure. In particular, the agents' utilities can be non-additively separable across items, but quasi-linear in money; i.e. (assuming all choices made in the same period)

$$\begin{aligned} U_a^0 &= \delta_a^t(u_a(x, \theta_a) + p), \\ U_b^0 &= \delta_b^t(u_b(x, \theta_b) - p), \\ S(\theta_a, \theta_b) &= \max_{x \in X^n} (u_a(x, \theta_a) + u_b(x, \theta_b)), \end{aligned}$$

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

In contrast, the item-by-item negotiation (3.3.1) does not have share-demanding offers: the offerer's payoff depends on which items the recipient accepts - and the offerer cannot request an overall surplus that must be taken as a whole rather than in part. This creates a tension between the offerer's incentives and the realization of full surplus. For instance, in Example 2 the seller is able to demand a value of $v = 6$, say by asking for a price of 3 on each item. However, the buyer would rather just accept the price of 3 on the more valuable item, thus realizing a total surplus of 7, instead of accepting both and only getting a surplus of 6. This would only lead to a surplus of 3 for the seller.

With additional assumptions on distributions, the results in Theorem 2 can be extend to the frequency negotiations (3.3.4).¹⁶ The small size of frequency protocol's strategy space is also helpful when we turn to settings with approximately known surplus in Section 5.3.2.

¹⁵Any offer (e.g. from the seller) (θ_s, v_s) in the negotiation protocol from Section 3.3 has an equivalent offer $p(x) = \sum_{k \in x} \theta_{sk} + v_s, \forall x$ in the combinatorial protocol.

¹⁶In particular, if type distributions are exchangeable then a sufficient condition is that the valuations are i.i.d. across items and independent across agents. Then in all exchangeable equilibria (such that an agent adopts the same strategy (in Phase 1) for types that have the same frequency), Phase 2 has a unique

Reopened negotiation for not-yet-agreed Items

Our general definition of negotiation allows the agent to continue the negotiation over the not-yet-agreed-upon items. Furthermore, when a negotiation has share-demanding offers, the corresponding negotiation game allowing for reopening for the not-yet-agreed-upon items also has share-demanding offers. Therefore, as a corollary to Theorem 2, it follows that the efficiency of all equilibria is robust to the opportunity of reopening the negotiation.

The opportunity to reopen exchange usually complicates the analyses of many settings, from Walrasian exchange to auctions to contracting. Here, the robustness comes from the share-demanding offers. Even though there is substantial uncertainty, and the potential to reopen discussions usually distorts incentives for screening, those incentives are completely circumvented with the share-demanding offers.

5 Multi-Item Negotiations with a ‘Nearly-Known’ Surplus

Our analysis so far illustrates that with known surplus, share-demanding offers provide for efficient negotiation and are essential for such results. The ability to bargain over a full bundle, means that the known surplus dominates the screening of particular items. Thus, we end up with a sort of “Rubinstein” result, rather than a “Myerson-Satterthwaite” result, with many items and uncertainty over each item but a known surplus overall.

The exact knowledge of the full surplus is an expository device, as we can imagine that with large numbers of items agents will have a good idea of the total surplus possible, but still have substantial uncertainty about which decisions should be made on each item or aspect. Thus, it is useful to verify that there is not a substantial discontinuity between having the total surplus being ‘nearly-known’ versus exactly-known. And, given that *all* equilibria are efficient in the limit, it is enough to look for upper-hemi continuity.

In this section we explore that continuity. There are several technical difficulties with establishing upper-hemi continuity.

The first is that incomplete information game theory is still not well-understood in the case of continua of types and actions - as measurability issues and other issues of updating beliefs conditional on atomless events is cumbersome (e.g., sequential equilibria are not well defined for such settings, see, Myerson & Reny (2015)). This prompts us to discretize the games - so that measurability issues are avoided. In particular, we require that the transfers between the agents can only be selected from some arbitrarily large but finite grid.

equilibrium: the responder makes the efficient choices given the other’s listed types and own types, provided that the surplus thus obtained exceeds the payoff demanded by the offerer. As a best response, the offerer lists the types truthfully.

The second difficulty is that the updating of beliefs is problematic even in very simple incomplete information games. Notice that in the previous section we did not impose any restriction on belief updating off the equilibrium path. Once there is non-trivial uncertainty of the surplus, restrictions on belief updating off-path are needed to have any hope for upper hemi-continuity, which we show via examples. It is important to note that this is a general problem with incomplete information games and not just our setting. In particular, under standard equilibrium notions including sequential equilibria (even when well-defined) or 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.

The third difficulty is that in order to have the total surplus be ‘nearly-known’ it makes sense to work large numbers of items - so that one can appeal to laws of large numbers. However, that means that the action space in our previously discussed negotiation protocols explodes exponentially. This leads to challenges in characterizing how beliefs evolve in equilibria.

To handle these three issues we work with a fixed number of items with uncertainty that converges to full knowledge, and have agents tremble so that beliefs are tied down; and we analyze protocols in which the strategy space satisfies a size restriction, but still allows for share-demanding offers (such as the frequency protocol) - thus allowing us to bound beliefs and characterize the equilibrium correspondence. In the appendix we also show that similar results hold in more general games if one directly bounds the rate at which beliefs update.

We begin by illustrating some of the technical issues via examples.

5.1 Multi-Item Negotiation with Converging Surplus

We index settings by a sequence m , to allow (but not require) the rate at which uncertainty converges not to depend directly on the number of items. The m -th economy has n_m items or aspects.

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 $\bar{s} > 0$ ¹⁷ if

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

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

$$S(\theta_a, \theta_b) \equiv \max_{x \in X^n} \sum_k [u_a(x_k, \theta_{ak}) + u_b(x_k, \theta_{bk})].$$

¹⁷We use the capital letter S to represent the total surplus, and the lower letter s for the per-item surplus.

This embodies the idea that there can be substantial uncertainty about which decisions should be made. In addition, the above allows for correlations of types across items and between agents. It also allows for heterogeneities across decisions.

Transfer grids Consider some grid of transfers, so that P^Δ is finite with a grid structure $\{-S_{\max}, \dots, 0, \Delta, 2\Delta, \dots, S_{\max}\}$,¹⁸ in which $S_{\max} = \max S(\theta_a, \theta_b)$.

5.2 A Technical Challenge: Failure of Upper-Hemicontinuity of Sequential Equilibria at the Limit of Certainty

We first illustrate the challenge that is substantial but is of a ‘technical’ nature: sequential equilibria often fail a fundamental upper hemi-continuity condition. We view this as a shortcoming of sequential equilibrium and the current tool-box of game theory. Moreover, this is not solved by standard existing refinements.

A game with arbitrarily small uncertainty is very different from its counterpart with certainty, in the sense that some sequences of sequential equilibria of the former game have no limit in the set of sequential equilibria (subgame-perfect equilibria) of the limit game with certainty. This happens because the the notion of sequential equilibria allows for a lot of freedom in off-path beliefs, and as a result lots of outcomes can be supported as part of a sequential equilibrium by extreme off-path beliefs, and leads to a failure of basic conditions like upper hemicontinuity of the equilibrium correspondence¹⁹ This challenge is not specific to our multi-item negotiation games. It applies to many simple games. Here we show that upper hemi-continuity even fails in simple single-item Rubinstein bargaining with the most basic forms of uncertainty.

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). Thus, there is a fundamental sort of discontinuity between equilibrium concepts with slight amounts of incomplete information and the limit of full information, which seems symptomatic of the tools of game theory rather than a real phenomenon. This does not contradict the fact that when both the 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 since we are converging to complete information. Given the importance of

¹⁸The increment Δ can be viewed as a 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} .

¹⁹Upper 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.

the complete information case in the theory and (its approximation) in practice, the failure of upper hemi-continuity is important and disturbing.

Consider a Rubinstein bargaining game with one item and one-sided uncertainty (let $\delta_a = \delta_b = \delta = 0.8$): Bob's value is commonly known as 10, and Alice's cost is either 0 or 8, 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 8 along the sequence. One may conjecture that all sequential equilibria in this game converge to the unique equilibrium in the limiting complete information bargaining game in which Alice's cost is 8 for sure, however this is not the case.

In particular, consider a price grid $P^\Delta = \{0, 1, 2, \dots, 10\}$ in order to have a nice finite game. The unique subgame perfect equilibrium of the limiting game (i.e., a complete information game with $\theta_a = 8$ and $\theta_b = 10$) is immediate trade at a price of 9. Below we show that with arbitrarily small uncertainty, so that $f_a(8) = 1 - \varepsilon$ for any tiny ε , 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 3 *With above parameters, 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 $p = 10$ and are rejected; at $t = 1$, the buyer offers a $p = 9$, which is accepted by both types of the seller.*

To see why the claim is true, consider the following equilibrium specification. The on-path behavior is supported by the buyer's *off-path* belief $\Pr(\theta_a = 8) = 0$ upon seeing any offer $p \neq 10$ at $t = 0$.²⁰ Given this belief, the buyer plays as if in a complete information Rubinstein bargaining with with "0 meets 10", i.e. always offering $p = 5$, and rejects any offer with $p > 5$. 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 problem of substantial inefficiency as presented in Example 3 remains with a discount factor arbitrarily close to 1: one can 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 example ($1 - 0.8 = 20\%$). The same problem also remains with an arbitrarily fine grid of transfers.

With arbitrarily small uncertainty, the set of sequential equilibria allows for substantial multiplicity in outcomes, unlike the uniqueness feature in the limiting game with exactly known surplus. This multiplicity is due to dramatic belief updating off-path that is robust

²⁰It is direct to check 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 10 compared to the 8 types.

to standard refinements²¹: in Example 3, upon seeing one off-path offer, the buyer believes the seller is of a 0 type for sure, completely discarding the prior belief which puts almost all weights on the 8 type. This dramatic change in beliefs results in a continuation of the game that is very different from the one where the agents started with, hence a lack of unique prediction of outcomes. Whereas in the limiting complete information game, there is no room for belief updating. Hence there is a loss of upper hemi-continuity of the set of sequential equilibria at the limit of certainty.

Since we aim to understand the case in which the overall surplus is increasingly known, and in *all* equilibria, additional structure on the game is necessary due to deal with above challenge.

Our main approach is to introduce trembles - which can be small, *but are not forced to 0* - to deal with belief-updating. This places all actions on the equilibrium path with some minimal weight from all types, and precludes the problems due to lack of belief restrictions, and avoids us having to make ad hoc restrictions concerning off-path beliefs. Alternative approaches and results are discussed in Appendix A.

5.3 Approximate Efficiency Results with Trembles

We introduce trembles that naturally regulate the rate at which beliefs can be updated. We first illustrate the idea in the single-item bargaining game.

5.3.1 Single-item bargaining: an illustrative example

First, reconsider Example 3. Consider a variation on the game such that there is 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 (or $\gamma/2$ on each of accept/reject for the responder); and subject to that constraint chooses the remaining probability according to a best response under the agent's beliefs.²² So, it is as if an agent best responds with probability $1 - \gamma$ and then trembles with the remaining probability γ - picking an action uniformly at random. The exact 'uniform at random' aspect of trembles is not needed, as is clear from the proofs; but it is necessary that the trembles not become

²¹Beyond Kreps and Wilson (1982), see Rubinstein (1985), Banks and Sobel (1987), Grossman and Perry (1986), and Cho and Kreps (1987).

²²Kreps 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).

infinitely more likely on some actions than others. With trembles, all nodes are reached and so beliefs are completely tied-down by Bayes' Rule, and so we can work with a trembling version of Bayesian equilibrium in which agents' update beliefs via Bayes' rule at all nodes and they best respond with probability $1 - \gamma$ and tremble with the remaining probability.

EXAMPLE 4 Consider $P^\Delta = \{0, 1, 2, \dots, 10\}$, and $\gamma = 0.11$ so that the probability of trembles to each possible price (in all periods) is $\gamma/|P^\Delta| = .01$. 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 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 $\Pr(a_i | \theta'_i) \geq .9$. ■

When agents tremble at some minimal rate, after one period of belief updating the posterior still puts substantial probability on the most probable types. Therefore, even very small trembles can make the posteriors more congruent with the priors in each period.

The following proposition shows that this convergence of beliefs is enough to generally restore continuity of the equilibrium correspondence with respect to small amounts of uncertainty.

PROPOSITION 1 Consider a single-item alternating-offer (Rubinstein) bargaining game. For any $\varepsilon > 0$ and correspondingly fine enough P^Δ , there is a small enough tremble probability $\gamma(\varepsilon) > 0$, such that for any $\gamma \in (\gamma(\varepsilon), 0)$ there exists low enough uncertainty ($\bar{f} < 1$ such that if f places probability at least \bar{f}_γ on a single pair of types $(\bar{\theta}_a, \bar{\theta}_b)$, then in all Bayesian equilibria, with probability at least $1 - \varepsilon$:

- if $\bar{\theta}_b > \bar{\theta}_a$, then the price offered in the initial period is in $((1 - \varepsilon)\bar{p}, (1 + \varepsilon)\bar{p})$ and is accepted, where the price $\bar{p} = \bar{\theta}_a + \frac{(1 - \delta_b)(\bar{\theta}_b - \bar{\theta}_a)}{1 - \delta_b\delta_a}$ is the Rubinstein price associated with the high probability types $\bar{\theta}_a, \bar{\theta}_b$; and
- if $\bar{\theta}_b < \bar{\theta}_a$, then trade does not occur.

Thus, Proposition 1 shows that introducing small trembles can completely eradicate the discontinuities associated with incomplete information. The proof of this proposition is a variation on that of Theorem 3, and so we omit it.

Let us be explicit about the order of the quantifiers, as the order is subtle, but not as restrictive as it might superficially appear. Given any ε it is clear that we need a fine enough grid and small enough trembles to be sure that an ε -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, \bar{f} is chosen as a function of γ , so that belief on ‘very unlikely types’ cannot become too large under trembles.

Effectively, trembles tie down beliefs and avoid the problems of updating off the path that drove the discontinuities in sequential equilibria without sustained trembles. Here, we get a continuity result at the limit (a technique that could also be helpful in other settings, beyond negotiations).

5.3.2 Multi-item negotiation under the frequency protocol.

Next, we illustrate how the near efficiency result with trembles applies to negotiations with multiple items, under the frequency protocol introduced in Subsection 3.3.4. We work with the frequency protocol since it has a “small” strategy space and this makes the handling of beliefs under trembles tractable. We discuss extensions to other negotiation protocols in the appendix.

Recall that there is some grid of transfers, $P^\Delta = \{-S_{\max}, \dots, 0, \Delta, 2\Delta, \dots, S_{\max}\}$, that is finite.

Again, consider trembles in *Phase 1* in any period by all types of any player with probability γ , uniformly to each of the feasible actions.²³ As noted above, uniformity of the trembles is an expository convenience, and all that is needed is 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).

The following distributional assumptions help put the frequency protocol 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 Θ_i , and independent across agents.²⁴ Without loss of generality, let $\bar{\phi}_i(\theta_{ik}) > 0, \forall \theta_{ik} \in \Theta_i$ (i.e., defining Θ_i to be the support).

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

$$\bar{s} = \sum_{\theta_a} \sum_{\theta_b} \bar{\phi}_a(\theta_a) \bar{\phi}_b(\theta_b) \max_{x \in X} [u_a(x_k, \theta_{ak}) + u_b(x_k, \theta_{bk})].$$

With these 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 his or her types that have the same frequency. We can then prove the following approximate efficiency result for the frequency protocol.

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

²⁴These 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 θ : if π is a bijection, then $f^m(\theta^\pi) = f^m(\theta)$ for all θ , where $\theta_k^\pi = \theta_{\pi(k)}$); and there is an exponential rate of precision improvements, sufficient conditions for which are stationarity and with summable covariance (cf. <https://stat.duke.edu/courses/Fall11/sta205/lec/wk-07.pdf> Section 7.2).

THEOREM 3 Consider a sequence of problems indexed by the number of items n with negotiation under the frequency protocol with given δ_a, δ_b , and i.i.d. distributions over types on each item. For any $\varepsilon > 0$, there exists a small enough tremble probability $\gamma(\varepsilon) > 0$, such that for any $\gamma \in (0, \gamma(\varepsilon))$ there exists n_γ such that if $n > n_\gamma$:

1. There exist exchangeable (weak) perfect Bayesian equilibria, subject to the trembles.²⁵
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
 - Expected payoff / ‘Full-Information Rubinstein share’ for each agent lies in $(1 - \varepsilon, 1 + \varepsilon)$.

Although we do not provide rates of convergence, they are direct to deduce for Theorem 3. The inefficiency per-item, when trembles are set to be as small as possible as a function of n (as otherwise they drive the inefficiency), is of order $O(n^{-0.5-\tau}), \forall \tau > 0$. Thus, the inefficiency vanishes at a rate arbitrarily close to the square-root of n . In particular, inefficiency comes from several sources: First, there is a potential inefficiency due to the increment in price grids. Here that naturally disappears with n since negotiation is over total surplus which is normalized by the number of items, so the inefficiency due to the lumpiness in the grid is of the order $O(\frac{1}{n})$. Second, there is inefficiency due to the trembles, which is proportional to γ , which can be picked as disappearing with n , in particular of order $O(\frac{1}{n})$. Third, the realized surplus can be different from the limit surplus, and such a difference induces inefficiency that is of the order of $O(n^{-0.5-\tau}), \forall \tau > 0$, a rate similar to ones provided by standard central limit theorems.²⁶

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

²⁵In this game with trembles all nodes are reached by all types with positive probability, and so perfect Bayesian equilibria, sequential equilibria, and weak perfect Bayesian equilibria coincide.

²⁶Fixing any rate of trembles $\gamma > 0$, and time T such that the time-0 continuation value after period T is negligible, the rate at which time- T posterior on some event may differ from time-0 prior is of the order $O(\gamma^{-T} n^{|\Theta_i|T})$ due to the trembles, where the size of i 's action space is $\sim O(n^{|\Theta_i|})$. Setting $\gamma \sim O(\frac{1}{n})$, this becomes $O(n^{(|\Theta_i|+1)T})$. Let d be an amount allowed between the realized and limit surpluses. By standard concentration inequalities, e.g. Hoeffding (1963), the beliefs that the actual surplus and realized surplus differs by more than d (at some belief condition on a period- T) is at most $\alpha^T \sim O(e^{-2d^2n} \times n^{(|\Theta_i|+1)T})$ which is still of the order $O(e^{-2d^2n})$ since the exponential term dominates.

$\max\{O(d), O(e^{-2d^2n})\}$ is minimized with an optimal selected difference $d \sim O(n^{-0.5-\tau})$, resulting in an overall inefficiency of the order of $O(n^{-0.5-\tau})$. Notice that the square-root rate of convergence cannot be achieved because we need to simultaneously control the difference allowed and the likelihood of the tails.

some types with a given frequency than others. We conjecture that a similar result holds when extending to those equilibria, as the rate at which helpful information is gained is bounded by the trembles, but we have not been able to find a proof or counter-example for the non-exchangeable case.

6 Unknown Surplus

Our attention in the paper has been on the case of (nearly) known overall surplus. Clearly getting approximate efficiency in all equilibria cannot hold generally, as that would violate the Myerson-Satterthwaite Theorem. Nonetheless, the result that full efficiency 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 can provide some intuitive sufficient conditions, that in fact appear to be close to necessary in order to get efficiency in all equilibria.

Consider a negotiation problem, (n, u, X, Θ, f) as defined above. Here, for ease of exposition, we focus on the exchange of items, but the logic extends.

6.1 An Example with Substantial Uncertainty about the Total Surplus

Under a negotiation protocol that is a variation of the combinatorial negotiation, there exists a fully efficient equilibrium. With a refinement on belief updating, all equilibria lead to the fully efficient outcome with a unique division of the gains from trade.

Consider a two-item exchange problem, in which the seller's cost vector $(\theta_{a1}, \theta_{a2})$ is either $(0, 8)$ or $(8, 0)$, and the buyer's value vector $(\theta_{b1}, \theta_{b2})$ is either $(2, 10)$ or $(10, 2)$, all equally likely and independent across agents. The corresponding total surplus from trade is either 4 or 10, equally likely.

The following table depicts the efficient trades.

	(2, 10)	(10, 2)
(0, 8)	trade both	trade 1st
(8, 0)	trade 2nd	trade both

Consider an alternating offer negotiation protocol, in which an offer consists of a triple $(k, p_{\{k\}}, p_{\{1,2\}})$, which specifies an item $k \in \{1, 2\}$ which the offerer is willing to trade individually at a price $p_{\{k\}} \in \mathbb{R}$, as well as a price $p_{\{1,2\}} \in \mathbb{R}$ for both items. The responder, if accepting the offer, can choose to trade either just the k -th item or else both items, at the corresponding price. Note that this negotiation protocol is universal.

We now provide one condition on belief updating. To do so, we say a type is of *category- k* if it is possible that (with that type) the efficient outcome is to trade k -th item individually. In the above example, the seller’s $(0, 8)$ and the buyer’s $(10, 2)$ are of category-1, and $(8, 0)$ and $(2, 10)$ are of category-2.

Here is a condition on belief updating:

Upon receiving an offer $(k, p_{\{k\}}, p_{\{1,2\}})$, the responder updates his or her belief to be sure that the offerer is of type of category- k (provided that belief does not contradict the prior).

For the above example, under the given negotiation protocol, There exists a weak perfect Bayesian equilibrium that is efficient and satisfies the belief updating condition.

7 Concluding Remarks

Despite the fact that ‘real-world’ negotiations frequently involve several aspects of a contract or deal, the formal theory of negotiation, as exemplified by the seminal work of Rubinstein (1982) and the literature that follows, focuses on a situation in which there is a single aspect (often a monetary transfer) to be determined. We extend that theory to encompass negotiations when an agreement includes many aspects. We distinguish two different dimensions of asymmetric information: that about each individual aspect, or the optimal decision regarding each aspect; and that about the overall surplus. Involved parties can 1) have a good knowledge about the overall surplus, even when 2) there is substantial Bayesian uncertainty regarding the valuation of each of these aspects. When 1) holds, we show that all “share-demanding” negotiations (a rich set of negotiation games) lead to efficient exchange in all equilibria, while item-by-item protocols result in inefficiency in all equilibria.

These negotiation games are extensions of Rubinstein (1982) alternating offer bargaining game. Also these protocols are ‘universal’ or ‘detail-free’, in the sense that their game forms are not tailored to the knowledge about the distribution of agents’ types. This is in contrast to many mechanisms, including the linking mechanisms of Jackson and Sonnenschein (2007), which are designed according to the statistical dispersion of information. The investigation of mechanisms that are independent of the (distributional knowledge of the) uncertainty faced by the agents, and operate for many settings, is an important issue for applications beyond negotiations and auctions, as the existence of mechanism designers with precise statistical knowledge is certainly far from ubiquitous.

In showing that our results are robust to small amounts of uncertainty about the overall surplus, we confronted a fundamental game-theoretic difficulty: the failure of upper-hemicontinuity of sequential equilibria at the limit of certainty. Our approach of introducing non-vanishing trembles may be useful as a natural approach to restoring upper hemicontinuity of the equilibrium correspondence in general settings - a topic for further research.

References

- Abreu, D., D. Pearce, and E. Stacchetti (1991): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* 58 (5), 1041-1065.
- 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.
- 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.
- Cohn, Z. (2010): "A Note on Linked Bargaining," *Journal of Mathematical Economics*, 46 (2), 238-247.
- 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.
- (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
- 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.
- 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.
- 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): 97109.
- Rubinstein, A. (1985): “A Bargaining Model with Incomplete Information about Time Preferences,” *Econometrica*, 53 (5), 1151-1172.
- 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)

Appendix

A Proofs

Proof of Theorem 1:

This Theorem is a corollary to Theorem 2 given the fact that the protocol 3.3 includes share-demanding offers. In particular, for any v , a share- v demanding offer is to announce the truth θ_i^n and demand a payoff of v . ■

Proof of Example 2:

Let L_s be the seller's worst continuation payoff in any seller-offer period in any sequential equilibria with both items remaining. This means when the buyer makes an offer, he gets a continuation payoff of at most $\delta(12 - \delta L_s)$ 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-2 item, and accepts p on the value-10 item for sure if $p < \tilde{p}$, s.t.

$$10 - \tilde{p} + \frac{2\delta}{1 + \delta} = (12 - \delta L_s)\delta,$$

where on the left-hand side $10 - p$ is the payoff from the value-10 item and $\frac{2\delta}{1+\delta}$ from the value-2 item (the corresponding Rubinstein share, since it is commonly known that the item left is of value-2).

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

On the other hand, since L_s 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

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

This, combined with the definition of \tilde{p} , gives $(1 - \delta^2)L_s \geq 10(1 - \delta) - \epsilon$, i.e. (since ϵ can be arbitrarily small)

$$L_s \geq \frac{10}{1 + \delta}$$

Finally, for both items to be traded in the initial period the seller's expected payoff is at most 4, which is not possible in any sequential equilibria: 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 2 on each item. The seller gets a payoff of at most 4 from such an offer, and hence an expected payoff of at most 4 from the game since the seller must be indifferent among any strategies used with positive probability. ■

Proof of Theorem 2:

We begin with some notation. Let $\Theta^n(f) \equiv \{(\theta_s, \theta_b) \mid f(\theta_s, \theta_b) > 0\}$ 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 $S(\theta_s, \theta_b) = \bar{S}$ for all $(\theta_s, \theta_b) \in \Theta^n(f)$. In addition, in a sequential equilibrium, 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.

Let M_i^t be the supremum of the expected per-item continuation payoff for agent i , starting at the beginning of period t over all sequential equilibria histories, and all i 's types in $\tilde{\theta}_i \in \Theta_i^n(f)$

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.

We first show that $M_s^0 \leq \frac{1-\delta_b}{1+\delta_s\delta_b}\bar{S}$.

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

$$L_b^{t+1} \equiv \bar{S} - \delta_s M_s^{t+2}.$$

The buyer does so by offering a share- v_{t+1} demanding offer with $v_{t+1} \equiv n(L_b^{t+1} - \eta)$ for $\eta > 0$ arbitrarily small. Such an offer is accepted for sure for a seller with any type $\hat{\theta}_s$ s.t. $(\hat{\theta}_s, \theta_b) \in \Theta^n(f)$: notice that $\hat{\theta}_s \in \Theta_s^n(f)$ by construction, hence $S(\hat{\theta}_s, \theta_b) = \bar{S}$; therefore by accepting the offer (and picking any items on which her cost is lower than the announced buyer values) the seller gets $\delta_s M_s^{t+2} + \eta$, which exceeds $\delta_s M_s^{t+2}$, the present value of the payoff from the continuation of the game if rejecting the offer. Finally, the buyer gets a payoff of $L_b^{t+1} - \eta$ if the above offer is accepted, regardless which items the seller picks to trade.

At t , the seller makes offer. We argue that a seller with any type $\theta_s \in \Theta_s^n(f)$ can get a payoff at most $\bar{S} - \delta_b L_b^{t+1}$: With any type $\theta_b \in \Theta_b^n(f)$, by rejecting an offer at t , the buyer's payoff from the continuation of the game has a present value of at least $\delta_b L_b^{t+1} - \delta_b \epsilon$ for $\forall \eta > 0$. Hence the payoff left to the seller with $\theta_s \in \Theta_s^n(f)$ is at most $\bar{S} - \delta_b L_b^{t+1}$, as the (expected) surplus is \bar{S} by construction.

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

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

$$M_s^0 \leq \frac{1 - \delta_b}{1 + \delta_s \delta_b} \bar{S}.$$

By a similar argument, it follows that

$$L_s^0 \geq \frac{1 - \delta_b}{1 + \delta_s \delta_b} \bar{S}.$$

Therefore, the payoff for the seller with any possible type $\theta_s \in \Theta_s^n(f)$ in any sequential equilibrium is $U_s \leq \frac{1 - \delta_b}{1 + \delta_s \delta_b} \bar{S}$.

From the above, we also know that $M_s^2 = L_s^2 = \frac{1 - \delta_b}{1 + \delta_s \delta_b} \bar{S}$, hence $M_b^1 = L_b^1 = \frac{1 - \delta_s}{1 + \delta_s \delta_b} \bar{S}$ (both in terms of the present value then), i.e. the total surplus realized is at least $\frac{1 - \delta_b}{1 + \delta_s \delta_b} \bar{S} + \delta_b \frac{1 - \delta_s}{1 + \delta_s \delta_b} \bar{S} = \bar{S}$ 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:

To simplify notations we prove this theorem for the private-good exchange application, with θ_{ak} being the seller's cost and θ_{bk} being the buyer's private value for the k -th item.

Notation:

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

Let $\phi[\theta_i] \in \Delta(\Theta_i)$ denote the frequency of a valuation type θ_i . So $\phi : \{\Theta_i^n\}_{i,n} \rightarrow \{\Phi_i^n\}_{i,n}$, and the notation $\phi[\theta_i](\theta_k)$ denotes the fraction of items having a specific value θ_k .²⁷

²⁷Notice 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 $\theta_s = (0, 0, 8, 8, 8)$, s 's true frequency is $\phi[\theta_s] = (\frac{2}{5}, \frac{3}{5})$, where $\phi[\theta_s](0) = \frac{2}{5}$ and $\phi[\theta_s](8) = \frac{3}{5}$.

When there is no confusion, we also use $\phi_i^n \in \Phi_i^n$ for i 's true frequency, and $\widehat{\phi}_i^n$ for a feasible frequency that can be announced.

Recall that $S(\theta_s, \theta_b)$ is the surplus with the corresponding pair of valuation types. With a slight abuse of notation, we extend the definition of this function to capture the expected surplus as a function of a frequencies:

- $S(\theta_s, \phi_b^n) = \sum_k \sum_{\theta_{bk}} \phi_b^n(\theta_{bk}) \cdot (\theta_{bk} - \theta_{sk})_+$
- $S(\phi_s^n, \theta_b) = \sum_k \sum_{\theta_{sk}} \phi_s^n(\theta_{sk}) \cdot (\theta_{bk} - \theta_{sk})_+$
- $S(\phi_s^n, \phi_b^n) = \sum_k \sum_{\theta_{sk}, \theta_{bk}} \phi_s^n(\theta_{sk}) \phi_b^n(\theta_{bk}) \cdot (\theta_{bk} - \theta_{sk})_+$

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

Although the agents' beliefs are defined over each other's valuation types θ_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: $\Phi_i^n(d) \equiv \{\phi_i^n : |\phi_i^n - \bar{f}_i| < d\}$ and $\Phi^n(d) \equiv \{(\phi_s^n, \phi_b^n) : |\phi_i^n - \bar{f}_i| < d, i = s, b\}$, where $|\cdot|$ is sup norm. Note that the sets naturally depend on \bar{f}_i 's, but we omit them in the notation since they are fixed throughout the statement and proof of the theorem.

Let

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

be the time-0 prior on 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 γ 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.

$$\Pr(|\phi_i^n[\theta_i] - \bar{f}_i| \geq d | h_t) \leq \alpha_i^t(n, d, \gamma) \tag{1}$$

When there is no confusion, we simplify notation by using α_i^0 and α_i^t 's, for a given set of parameters (n, d, γ) .

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 $\phi[\theta_i]$, instead of the valuation type θ_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.²⁸

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 θ_i^n copies the type $\phi[\theta_i]$'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 θ_i^n shares the beliefs that the type $\phi[\theta_i]$ 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 an 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 recipient has a unique strict best reply to trade $X = \{k \mid \theta_{bk} > \theta_{sk}\}$ and pay an amount that exactly delivers the payoff demanded by the offerer - provided these lead to positive payoffs, and otherwise to say ‘No’ (and do either if there is indifference),²⁹ If the offerer was truthful in the first Phase on the announced frequency, then given recipient's strategy is to truthfully list valuations. (What happens in other subgames will not be important for the argument below.)

The Phase 2 strategies imply the following expected payoff from offering/accepting an offer in Phase 1: An offerer of type θ_i^n gets a (non-discounted) payoff of \hat{v}_i with a “truthful”

²⁸“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.

²⁹Generally, 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.

offer $(\phi[\theta_i], \widehat{v}_i)$ if the offer is accepted; recall that $\phi[\theta_i]$ is the true frequency of θ_i . A recipient's (non-discounted) expected payoff from accepting an offer $(\widehat{\phi}_i^n, \widehat{v}_i)$ is $\mathbb{E}v_j = S(\theta_j, \widehat{\phi}_i^n) - \widehat{v}_i$. This is true regardless of whether $\widehat{\phi}_i^n$ 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 $i \in \{s, b\}$:

- $M_i(\alpha_s, \alpha_b; d, n)$ [$L_i(\alpha_s, \alpha_b; d, n)$] is the sup [inf] of expected 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 $\Pr(|\phi_s^n[\theta_s] - \bar{f}_s| < d|h_t) \geq 1 - \alpha_s$ and $\Pr(|\phi_b^n[\theta_b] - \bar{f}_b| < d|h_t) \geq 1 - \alpha_b$.

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

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} :

$$|\mathbb{E}S(\phi_s^n, \phi_b^n) - \bar{S}| < 2dS_{max}, \forall (\phi_s^n, \phi_b^n) \in \Phi^n(d). \quad (2)$$

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

$$M_i(\alpha_s^t, \alpha_b^t) \leq \bar{S} + 2dS_{max} - \delta_j(1 - \alpha_j^{t+1})L_j(\alpha_s^{t+1}, \alpha_b^{t+1}) + \frac{1}{n} \quad (3)$$

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

$$L_i(\alpha_s^t, \alpha_b^t) \geq (1 - \gamma)(1 - \alpha_j^t) \left[\bar{S} - 2dS_{max} - \delta_j M_j(\alpha_s^{t+1}, \alpha_b^{t+1}) - \frac{1}{n} \right] \quad (4)$$

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 $\delta_j L_j(\alpha_s^{t+1}, \alpha_b^{t+1})$ 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 ϕ_i^n and any demanded payoff of

nor more than $S - 2dS_{max} - \delta M_j(\alpha_s^{t+1}, \alpha_b^{t+1})$. 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 $M_j(\alpha_s^{t+1}, \alpha_b^{t+1})$, 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 analagous).

First let $\delta_{max} \equiv \max\{\delta_s, \delta_B\}$. Next 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

$$error^{2t} = (\alpha_s^{2t} + \alpha_b^{2t} + \delta_{max}(\alpha_s^{t+1} + \alpha_b^{t+1}))S_{max} + (1 + \delta_{max})2dS_{max} + (1 + \delta_{max})\frac{1}{n} + \delta_{max}\gamma S_{max}.$$

Then it follows that

$$\sum_{t=0}^{T-1} error^{2t} = \sum_{t=0}^{2T-1} [\delta_{max}^t(\alpha_s^t + \alpha_b^t)] S_{max} + \frac{1 - \delta_{max}^{2T-2}}{1 - \delta_{max}} (2dS_{max} + \frac{1}{n} + \delta_{max}\gamma S_{max}).$$

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

$$\begin{aligned} M_s(\alpha_s^0, \alpha_b^0) &\leq \bar{S}(1 - \delta_b) + \delta_s \delta_b M_s(\alpha_s^2, \alpha_b^2) + error^0 \\ &= \bar{S}(1 - \delta_b)(1 + \delta_s \delta_b \dots + \delta_s^T \delta_b^T) + \delta_s^T \delta_b^T M_s(\alpha_s^{2T}, \alpha_b^{2T}) + \sum_{t=0}^{T-1} error^{2t} \\ &\leq \frac{1 - \delta_b}{1 - \delta_s \delta_b} \bar{S} + \left(\frac{\delta_s^T \delta_b^T}{1 - \delta_s \delta_b} S_{max} + \sum_{t=0}^{T-1} error^{2t} \right), \end{aligned} \quad (5)$$

$$\begin{aligned} L_s(\alpha_s^0, \alpha_b^0) &\geq \bar{S}(1 - \delta_b) + \delta_s \delta_b L_s(\alpha_s^2, \alpha_b^2) - error^0 \\ &= \bar{S}(1 - \delta_b)(1 + \delta_s \delta_b \dots + \delta_s^T \delta_b^T) + \delta_s^T \delta_b^T L_s(\alpha_s^{2T}, \alpha_b^{2T}) - \sum_{t=0}^{T-1} error^{2t} \\ &\geq \frac{1 - \delta_b}{1 - \delta_s \delta_b} \bar{S} - \sum_{t=0}^{T-1} error^{2t}. \end{aligned} \quad (6)$$

This implies that

$$\frac{1 - \delta_b}{1 - \delta_s \delta_b} \bar{S} - \sum_{t=0}^{T-1} error^{2t} \leq L_s(\alpha_s^0, \alpha_b^0) \leq M_s(\alpha_s^0, \alpha_b^0) \leq \frac{1 - \delta_b}{1 - \delta_s \delta_b} \bar{S} + (\delta_{max}^{2T} S_{max} + \sum_{t=0}^{T-1} error^{2t}), \quad (7)$$

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

- pick $T \in \mathbb{Z}_+$ such that $\delta_{max}^{2T} S_{max} < \eta/5$
- pick $d > 0$ such that $\frac{2dS_{max}}{1 - \delta_{max}} < \eta/5$
- pick $\gamma(\varepsilon) < \frac{\varepsilon}{5(1 - \delta_{max})S_{max}}$, hence $\frac{\gamma(\eta)S_{max}}{1 - \delta_{max}} < \eta/5$
- for any $\gamma \in (0, \gamma(\varepsilon))$ and d (already picked), pick the threshold number of items, n_γ , so that for $\forall n > n_\gamma$ we have $\sum_{t=0}^{2T-1} \left[\delta_s^{\lfloor \frac{t}{2} \rfloor} \delta_b^{\lfloor \frac{t}{2} \rfloor + 1} (\alpha_s^t(d, n, \gamma) + \alpha_b^t(d, n, \gamma)) \right] S_{max} < \eta/5$. 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);
- ◇ $\alpha_i^0 = \Pr^n \left(|\hat{\phi}_i^n - \bar{f}_i| \geq d \right) \leq 2e^{-2d^2 n}$ (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, if $\frac{1}{n_\gamma(1-\delta_{\max})} \geq \eta/5$, replace n_γ by $\frac{5}{\eta(1-\delta_{\max})}$ so that $\frac{1}{n(1-\delta_{\max})} < \eta/5$ 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

$$\mathbb{E}U_s^0(\phi_s^n) \in \left(\frac{1 - \delta_b}{1 - \delta_s \delta_b} \bar{S} - \eta, \frac{1 - \delta_b}{1 - \delta_s \delta_b} \bar{S} + \eta \right), \forall n > \bar{n}, \forall \phi_s^n \in \Phi_s^n(d) \quad (8)$$

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

$$\mathbb{E}U_b^1(\phi_b^n) \in \left(\frac{1 - \delta_s}{1 - \delta_s \delta_b} \bar{S} - \eta, \frac{1 - \delta_s}{1 - \delta_s \delta_b} \bar{S} + \eta \right), \forall n > \bar{n}, \forall \phi_b^n \in \Phi_b^n(d) \quad (9)$$

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

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

By the construction of η , we have $\mathbb{E}U_s^0(\phi_s^n)$ 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

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

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 $\Pr(\Phi^n(d)) = (1 - \alpha_s^0)(1 - \alpha_b^0) \geq 1 - 2d > \frac{\eta}{5S_{\max}}$ 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:

$$\begin{aligned}
\Pr(\text{delay}) &\leq \Pr(\text{delay} \mid \Phi^n(d)) \Pr(\Phi^n(d)) + 1 - \Pr(\Phi^n(d)) \\
&\leq \frac{4\eta}{\delta(\bar{S} - \eta)} + \frac{\eta}{5S_{max}} \\
&< \varepsilon
\end{aligned} \tag{13}$$

■

Supplementary Appendix: Additional Results with a Nearly-Known Surplus

(to “A Theory of Efficient Negotiations”

by Matthew O. Jackson, Hugo F. Sonnenschein, and Yiqing Xing)

Subsection 5.2 discussed a (technical) challenge associated with the multiplicity of sequential equilibria due to dramatic updating in posteriors. This challenge was handled in subsection 5.3 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 share-demanding offers (see definition 1, including our first protocol from Section 3.3 and the combinatorial protocol from Section 3.3.2, 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.

A.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 3.3 and the combinatorial protocol from Section 3.3.2. In such protocols, at the beginning any period t agents share a common history $h^{t-1} \equiv (a_{i(1)}^1, a_{j(1)}^1, \dots, a_{i(t-1)}^{t-1}, a_{j(t-1)}^{t-1})$, 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 θ_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 Θ_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 $\forall E \subset \Theta^n, (h^{t-1}, a_{i(t)}^t) \in H, a_{j(t)}^t \in A_{j(t)}(h^{t-1}, a_{i(t)}^t)$, and $\theta_i \in \Theta_i^n$:

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

We consider (arbitrarily) large but bounded β s. This means that the assumption only affects events that are very unlikely according to the prior. In particular, when $\tilde{f}(E | h^{t-1}) > \frac{1}{\beta}$, the constraint is not binding for the updating at history h^{t-1} .

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

DEFINITION 2 (EQUILIBRIUM WITH BOUNDED BELIEF UPDATING) *An equilibrium with bounded belief updating at rate β is a profile of the agents' (mixed) strategies and posterior systems $\tilde{f}_j, j = s, b$, 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}_s and \tilde{f}_b satisfy bounded updating at rate β .*

Definition 2 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.³⁰

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 share-demanding offers. For any $\varepsilon > 0$, $\forall \beta \in [1, \infty)$, $\forall \delta_s, \delta_b < 1$, there is $\Delta(\varepsilon) > 0$ such that for any $\Delta \in (0, \Delta(\varepsilon))$ there exists m_Δ such that if $m > m_\Delta$ then:*

1. *There exist equilibria with bounded belief updating at rate β ;*
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 $(1 - \varepsilon, 1 + \varepsilon)$.*

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. The 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 θ_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 $i = s, b$), given some $\iota_0, \iota_1, \dots > 0$

0. $\tilde{\Theta}_i^{n_m}(d, 0, \iota_0) = \{\theta_i \mid f_b(\Theta_j^{n_m}(\theta_i, d), \theta_i) > 1 - \iota_0\}, i = s, b;$
1. $\tilde{\Theta}_i^{n_m}(d, 1, \iota_1) = \{\theta_i \mid f_b(\Theta_j^{n_m}(\theta_i, d) \cap \tilde{\Theta}_j^{n_m}(d, 0, \iota_0), \theta_i) > 1 - \iota_1\}, i = s, b;$
- ...

³⁰An 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 β times the prior, that belief on that event is replaced by β times the prior. Same theorem holds under that alternative notion.

$$t. \tilde{\Theta}_i^{n_m}(d, t, \iota_t) = \left\{ \theta_i \mid f_b \left(\Theta_j^{n_m}(\theta_i, d) \cap \tilde{\Theta}_j^{n_m}(d, t-1, \iota_{t-1}), \theta_i \right) > 1 - \iota_t \right\}, i = s, b;$$

...

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 simply notation we omit the superscript n_m when there is no confusion.

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 $i \in \{s, b\}$, let $M_i^t [L_i^t]$ be the sup [inf] of the expected 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 $\theta_i \in \tilde{\Theta}_i(d, 2T-t, \iota_{2T-t})$.

We derive the bounds recursively:

- Backward from period $2T$:

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

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

...

- 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 $\delta_s M_s^{2T}$ 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_s^0 and L_s^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 $\Pr\left(\tilde{\Theta}_i(d, 2T, \iota_{2T})\right)$ close enough to 1.

To do so, for any $d > 0$, let $\alpha = \Pr(|S - \bar{S}| < d)$ which converges to 0 according to the definition of converging surplus. We construct ι 's from α_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 $\Pr\left(|S - \bar{S}| < d \text{ and } \theta_j \in \tilde{\Theta}_j(d, 0, \iota_0)\right) \geq 1 - (\alpha + \iota_0)$;
- $\iota_1 = \sqrt{\alpha + \iota_0}$, 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 ι_{2T} , which converges to 0 as α 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 ι_{2T} as m becomes large. The rest of the proof parallels the corresponding parts of the proof of Theorem 3. ■

A.2 Trembles with Fixed Number of Items.

When there is no confusion we write $M_i(\alpha_s, \alpha_b)$ and $L_i(\alpha_s, \alpha_b)$, but note that the payoff bounds depend on (d, n) . We now work with sequences of economies, where 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 5.3: consider trembles in any period by all types of any agent with probability γ , 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 share-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.