Reputational Bargaining with Minimal Knowledge of Rationality*

Alexander Wolitzky

Stanford University and Microsoft Research

December 18, 2011

Abstract

Two players announce bargaining postures to which they may become committed and then bargain over the division of a surplus. The share of the surplus that a player can guarantee herself under first-order knowledge of rationality is determined (as a function of her probability of becoming committed), as is the bargaining posture that she must announce in order to guarantee herself this much. This "maxmin" share of the surplus is large relative to the probability of becoming committed (e.g., it equals 30% if the commitment probability is 1 in 10, and equals 13% if the commitment probability is 1 in 1000), and the corresponding bargaining posture simply demands this share plus compensation for any delay in reaching agreement.

Keywords: *bargaining, knowledge of rationality, posturing, reputation*

*I thank the editor, Philippe Jehiel, and three anonymous referees for helpful comments. I thank my advisors, Daron Acemoglu, Glenn Ellison, and Muhamet Yildiz, for extensive discussions and comments as well as continual advice and support. For additional helpful comments, I thank Sandeep Baliga, Abhijit Banerjee, Alessandro Bonatti, Matthew Elliott, Jeff Ely, Drew Fudenberg, Xavier Gabaix, Bob Gibbons, Michael Grubb, Peter Klibanoff, Anton Kolotilin, Mihai Manea, Parag Pathak, Jeroen Swinkels, Juuso Toikka, Iván Werning, and (especially) Gabriel Carroll, as well as many seminar audiences. I thank the NSF for financial support.

1 Introduction

Economists have long been interested in how individuals split gains from trade. Recently, "reputational" models of bargaining have been developed that make sharp prediction about the division of surplus independently of many details of the bargaining procedure (Myerson, 1991; Abreu and Gul, 2000; Kambe, 1999; Compte and Jehiel, 2002; Abreu and Pearce, 2007). In these models, players may be committed to a range of possible bargaining strategies, or "postures," before the start of bargaining, and bargaining consists of each player attempting to convince her opponent that she is committed to a strong posture. These models assume that the probabilities with which the players are committed to various bargaining postures (either ex ante or after a stage where players strategically announce bargaining postures) are common knowledge, and that play constitutes a (sequential) equilibrium. In this paper, I study reputational bargaining while assuming only that the players know that each other is rational (so that, in particular, players do not know each other's beliefs or strategies). I show that each player can guarantee herself a share of the surplus that is large relative to her probability of being committed by announcing the posture that simply demands this share plus compensation for any delay in reaching agreement. Furthermore, announcing any other posture does not guarantee her as much.

More precisely, I assume that there is positive number ε such that, if a player announces any bargaining posture (i.e., any infinite path of demands) at the beginning of the game, she then becomes committed to that posture with probability at least ε (or, equivalently, she convinces her opponent that she is committed to that posture with probability at least ε). Player 1's "maxmin" payoff is then the highest payoff u_1 with the property that there exists a corresponding posture (the "maxmin posture") and bargaining strategy such that player 1 receives at least u_1 whenever she announces this posture and follows this strategy and player 2 plays *any* best-response to *any* belief about player 1's strategy that assigns probability at least ε to player 1 following her announced posture.

The main result of this paper characterizes the maxmin payoff and posture when only one player may become committed to her announced posture; as discussed below, a very similar characterization applies when both players may become committed. The maxmin

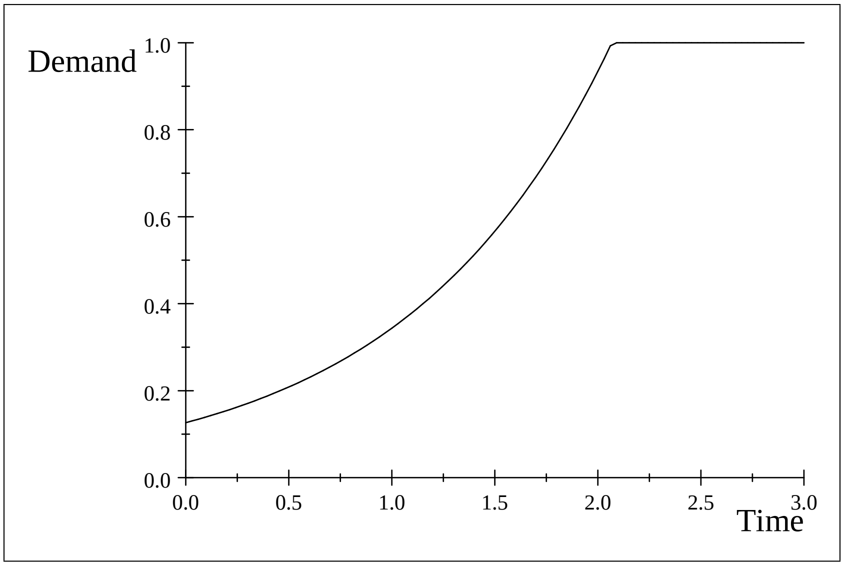

Figure 1: The Unique Maxmin Bargaining Posture for $\varepsilon = 1/1000$ and $r = 1$

payoff equals $1/(1 - \log \varepsilon)$. This equals 1 when $\varepsilon = 1$ (i.e., when the player makes a take-it-or-leave-it offer), and goes to 0 very slowly as ε goes to 0. For example, a bargainer can guarantee herself approximately 30% of the surplus if her commitment probability is 1 in ten; 13% if it is 1 in 1 thousand; and 7% if it is 1 in 1 million. In addition, the unique bargaining posture that guarantees this share of the surplus simply demands this share in addition to compensation for any delay; that is, the demand increases at rate equal to the common discount rate, r.[1] This compensation amounts to the entire surplus after a long enough delay, so the unique maxmin posture demands

$$\min\left\{e^{rt}/(1 - \log \varepsilon), 1\right\}$$

at every time t. This posture is depicted in Figure 1, for commitment probability $\varepsilon = 1/1000$ and discount rate $r = 1$.

The intuition for why the maxmin payoff is large relative to ε is that, when player 1's demand is small, player 2 must accept unless he believes that he will be quickly rewarded for

[1] To my knowledge, this is the first bargaining model that predicts that such a posture will be adopted, though it seems like a reasonable bargaining position to stake out. For example, in most U.S. states defendants must pay "prejudgment interest" on damages in torts cases, which amounts to plaintiffs demanding the initial damages in addition to compensation for any delay (e.g., Knoll, 1996); similarly, unions sometimes include payment for strike days among their demands.

rejecting. In the latter case, if player 1 does not reward player 2 for rejecting, then player 2 quickly updates his belief toward player 1's being committed to her announced posture (i.e., player 1 builds reputation at a high rate), and player 2 accepts when he becomes convinced that she is committed.[2] Hence, player 1 is able to compensate for having a small commitment probability by reducing her demand and thereby increasing the rate at which she builds reputation. This exponentially reduces the cost of the delay before her demand is accepted and thus guarantees her a relatively large payoff.

The intuition for why the unique maxmin posture demands compensation for delay involves two key ideas. First, as I have argued, player 1's demand is accepted sooner when it is lower (when player 2's beliefs are those that lead him to reject for as long as possible). Second, the maxmin posture can never make demands that would give player 1 less than her maxmin payoff if they were accepted, because otherwise player 2 could simply accept some such demand and give player 1 a payoff below her maxmin payoff, which was supposed to be guaranteed to player 1 (though it must be verified that such behavior by player 2 is rational). Combining these ideas implies that player 1 must always demand at least her maxmin level of *utility* (hence, compensation for delay), but no more.

I also characterize the maxmin payoffs and postures when both players may become committed to their announced postures. Each player's maxmin posture is exactly the same as in the one-sided commitment model, and each player's maxmin payoff is close to her maxmin payoff in the one-sided commitment model as long as her opponent's commitment probability is small. Thus, the one-sided commitment analysis applies to each player separately.

The paper most closely related to mine is Kambe (1999), which endogenizes the "behavioral types" of Abreu and Gul (2000) by having players strategically announce postures to which they may become committed (as in my model).[3,4] There are two differences be-

[2]This is related to the argument in the existing reputational bargaining literature that player 1 builds reputation more quickly *in equilibrium* when her demand is small, though my analysis is not based on equilibrium.

[3]Other important antecedents include Kreps and Wilson (1982) and Milgrom and Roberts (1982), who pioneered the incomplete information approach to reputation-formation, and Chatterjee and Samuelson (1987, 1988), who study somewhat simpler reputational bargaining models.

[4]It is not essential for the main points of my paper that commitment comes from strategic announcements

tween Kambe's model and mine. First, Kambe requires that players announce postures that demand a constant share of the surplus (as do Abreu and Gul), while I allow players to announce non-constant postures (and players do benefit from announcing non-constant postures in my model). Second, and more fundamentally, Kambe studies sequential equilibria (as does the rest of the existing reputational bargaining literature), while I study maxmin payoffs and postures. My approach entails weaker assumptions on knowledge of commitment probabilities (i.e., second-order knowledge that commitment probabilities are at least ε, rather than common knowledge of exact commitment probabilities) and on behavior (i.e., first-order knowledge of rationality, rather than sequential equilibrium), but does not yield unique predictions about the division of surplus or about the details of how bargaining will proceed. One motivation for this complementary approach is that behavioral types are sometimes viewed as "perturbations" reflecting the fact that a player (or an outside observer) cannot be sure that the model captures all of the other player's strategic considerations, and assuming that the distribution over perturbations is common knowledge goes against the spirit of introducing perturbations.

This paper is related more broadly to the literature on commitment tactics in bargaining dating back to Schelling (1956), who discusses observable factors that make announced postures more credible, corresponding to a higher value of ε in my model.[5] It is also related to the literatures on bargaining with incomplete information either without common priors (Yildiz 2003, 2004; Feinberg and Skrzypacz, 2005) or with rationalizability rather than equilibrium (Cho, 1994; Watson, 1998), in that players may disagree about the distribution over outcomes of bargaining. Finally, this paper weakens the solution concept from equilibrium to knowledge of rationality in reputational bargaining models in the same way that Watson (1993) and Battigalli and Watson (1996) weaken the solution concept from equilibrium to knowledge of rationality in Fudenberg and Levine's (1989) model of reputation in repeated games. They find that Fudenberg and Levine's equilibrium predictions also apply under knowledge of rationality, whereas my predictions differ dramatically from the existing rep-

rather than exogenous behavioral types. What matters is the posterior probability with which a player's opponent thinks she is committed after she stakes out a posture.

[5]Subsequent contributions include Crawford (1982), Fershtman and Seidmann (1993), Muthoo (1996), Compte and Jehiel (2004), and Ellingsen and Miettinen (2008).

utational bargaining literature. The reason for this difference is that in repeated games with one long-run player and one short-run player, or in bargaining where one player is infinitely more patient than her opponent, the long-run or patient player receives close to her Stackelberg payoff (which in bargaining equals the entire surplus) under either knowledge of rationality or equilibrium. Thus, the predictions of Watson and Battigalli and Watson coincide with those of Fudenberg and Levine, just as my predictions coincide with those of Abreu and Gul (and others) *in the special case where one player is infinitely more patient than the other* (Section 4.2). However, the main focus of this paper is on the case of equally (or at least comparably) patient players, where predictions under equilibrium and knowledge of rationality differ for both repeated games and bargaining.

The paper proceeds as follows: Section 2 presents the model and defines maxmin payoffs and postures. Section 3 analyzes the baseline case with one-sided commitment and presents the main characterization of maxmin payoffs and postures. Section 4 presents three brief extensions. Section 5 considers two-sided commitment. Section 6 concludes. Omitted proofs are in the appendix. A supplementary appendix shows that the main characterization is robust to details of the bargaining procedure such as the order and relative frequency of offers (so long as offers are frequent), as well as to strengthening the solution concept from knowledge of normal-form rationality to iterated conditional dominance.[6]

2 Model and Definition of Maxmin Payoff and Posture

2.1 Model

Two players ("she," "he") bargain over one unit of surplus in two phases: a "commitment phase" followed by a "bargaining phase." I describe the bargaining phase first. It is intended to capture a continuous bargaining process where players can change their demands and accept their opponents' demands at any time, but in order to avoid well-known technical issues that emerge when players can condition their play on "instantaneous" actions of

[6] Rubinstein (1982) implies that the main result is *not* robust to *simultaneously* specifying a discrete-time bargaining procedure and strengthening the solution concept to iterated conditional dominance. See the supplementary appendix for details.

their opponents (Simon and Stinchcombe, 1989; Bergin and MacLeod, 1993) I assume that players can revise their paths of demands only at integer times (while letting them accept their opponents' demands at any time).[7]

Time runs continuously from $t = 0$ to ∞. At every integer time $t \in \mathbb{N}$ (where $\mathbb{N}$ is the natural numbers), each player $i \in \{1, 2\}$ chooses a path of demands for the next length-1 period of time, $u_i^t : [t, t+1) \to [0, 1]$, which is required to be the restriction to $[t, t+1)$ of a continuous function on $[t, t+1]$. Let $\mathcal{U}^t$ be the set of all such functions, and let $\Delta(\mathcal{U}^t)$ be the space of probability measures on the Borel σ-algebra of $\mathcal{U}^t$ endowed with the product topology. The interpretation is that $u_i^t(\tau)$ is the demand that player i makes at time τ (this is simply denoted by $u_i(\tau)$ when t is understood; note that $u_i : \mathbb{R}_+ \to \mathbb{R}_+$ can be discontinuous at integer times but is everywhere right-continuous with left limits). Even though player i's path of demands for $[t, t+1)$ is decided at t, player j only observes demands as they are made. Intuitively, each player i may *accept* her opponent's demand $u_j(t)$ at any time t, which ends the game with payoffs $(e^{-rt}(1 - u_j(t)), e^{-rt}u_j(t))$, where $r > 0$ is the common discount rate (throughout, $j = -i$). Formally, every instant of time t is divided into three *dates*, $(t, -1)$, $(t, 0)$, and $(t, 1)$ (except for time 0, which is divided only into dates $(0, 0)$ and $(0, 1)$), with the following timing: First, at date $(t, -1)$, each player i announces *accept* or *reject*. If both players reject, the game continues; if only player i accepts, the game ends with payoffs $(e^{-rt}(1 - \lim_{\tau \uparrow t} u_j(\tau)), e^{-rt} \lim_{\tau \uparrow t} u_j(\tau))$; and if both players accept, the games ends with payoffs determined by the average of the two demands, $\lim_{\tau \uparrow t} u_1(\tau)$ and $\lim_{\tau \uparrow t} u_2(\tau)$. Next, at date $(t, 0)$, both players simultaneously announce their time-t demands $(u_1(t), u_2(t))$ (which were determined at the most recent integer time); if t is an integer, this is also the date where each player i chooses a path of demands for the next length-1 period, u_i^t. Finally, at date $(t, 1)$, each player i again announces *accept* or *reject*. If both players reject, the game continues; if only player i accepts, the game ends with payoffs $(e^{-rt}(1 - u_j(t)), e^{-rt}u_j(t))$; and if both players accept, the game ends and the

[7] In this paper, working in continuous time not only allows the players more flexibility than does discrete time but also yields simpler results. In particular, the continuous-time maxmin posture demands exact compensation for delay and thus changes over time in a simple way, while I conjecture that the discrete-time maxmin posture demands approximate compensation for delay but that the details of the approximation are complicated and depend on the exact timing of offers.

demands $u_1(t)$ and $u_2(t)$ are averaged. This timing ensures that there is a first and last date at which each player can accept each of her opponent's demands. In particular, at integer time t, player i may accept either her opponent's "left" demand, $\lim_{\tau \uparrow t} u_j(\tau)$, or her time-$t$ demand, $u_j(t)$. I say that *agreement is reached at time t* if the game ends at time t (i.e., at date $(t,-1)$ or $(t,1)$). Both players receive payoff 0 if agreement is never reached.

The public history up to time t excluding the time-t demands is denoted by $h^{t-} = (u_1(\tau), u_2(\tau))_{\tau<t}$, and the public history up to time t including the time-t demands is denoted by $h^{t+} = (u_1(\tau), u_2(\tau))_{\tau \leq t}$ (with the convention that this corresponds to all offers having been rejected, as otherwise the game would have ended). A generic time-t history is denoted by h^t. Since $\lim_{\tau \uparrow t} u_i(t) = u_i(t)$ for non-integer t, I generally distinguish between h^{t-} and h^{t+} only for integer t. Formally, *a bargaining phase (behavior) strategy* for player i is a pair $\sigma_i = (F_i, G_i)$, where F_i is a map from histories h^t into $[0,1]$ with the properties that $F_i(h^t) \leq F_i\left(h^{t'}\right)$ whenever $h^{t'}$ is a successor of h^t and $F_i(h^{t+})$ is a right-continuous function of t; and G_i is a map from histories h^{t-} with $t \in \mathbb{N}$ into $\Delta(\mathcal{U}^t)$. Let Σ_i be the set of player i's bargaining phase strategies. The interpretation is that $F_i(h^{t-})$ is the probability that player i accepts player j's demand at or before date $(t,-1)$, $F_i(h^{t+})$ is the probability that player i accepts player j's demand at or before date $(t,1)$, and $G_i(h^{t-})$ is the probability distribution over paths of demands $u_i^t : [t, t+1) \to [0,1]$ chosen by player i at date $(t,0)$. This formalism implies that player i's *hazard rate of acceptance at history h^t*, $f_i(h^t) / (1 - F_i(h^t))$, is well-defined at any time t at which the realized distribution function F_i admits a density f_i; and in addition player i's *probability of acceptance at history h^{t+}* (resp., h^{t-}), $F_i(h^{t+}) - F_i(h^{t-})$ (resp., $F_i(h^{t-}) - \lim_{\tau \uparrow t} F_i(h^\tau)$), is well-defined for all times t. However, so long as one bears in mind these formal definitions, it suffices for the remainder of the paper to omit the notation (F_i, G_i) and instead simply view a (bargaining phase) strategy $\sigma_i \in \Sigma_i$ as a function that maps every history h^t to a hazard rate of acceptance, a discrete probability of acceptance, and (if $h^t = h^{t-}$ for $t \in \mathbb{N}$) a probability distribution over paths of demands u_i^t. A *pure bargaining phase strategy* is a strategy σ_i such that $F_i(h^t) \in \{0,1\}$ for all h^t and $G_i(h^{t-})$ is a degenerate distribution for all h^{t-}.

At the beginning of the bargaining phase, player i has an initial belief π_i about the behavior of her opponent. Formally, $\pi_i \in \Delta(\Sigma_j)$, the set of finite-dimensional distributions over Σ_j, so π_i is a finite-dimensional distribution over behavior strategies σ_j; note that π_i can alternatively be viewed as an element of Σ_j by reducing lotteries over behavior strategies. Let $\mathrm{supp}(\pi_i) \subseteq \Sigma_j$ be the support of π_i, let $u_i(\sigma_i, \sigma_j)$ be player i's expected utility given strategy profile (σ_i, σ_j), let $u_i(\sigma_i, \pi_i)$ be player i's expected utility given strategy σ_i and belief π_i, and let $\Sigma_i^*(\pi_i) \equiv \mathrm{argmax}_{\sigma_i} u_i(\sigma_i, \pi_i)$ be the set of player i's (normal-form) best-responses to belief π_i (which may be empty; see footnote 8). An action (accepting, rejecting, or choosing a demand path u_i^t) is *optimal* at history h^t under belief π_i if there exists a pure strategy σ_i that prescribes that action at h^t such that $\sigma_i \in \Sigma_i^*(\pi_i)$.

At the beginning of the game (prior to time 0), player 1 (but not player 2) publicly announces a *bargaining posture* $\gamma : [0, \infty) \to [0, 1]$, which must be continuous at non-integer times t and be everywhere right-continuous with left limits. Slightly abusing notation, a posture γ is identified with the strategy of player 1's that demands $\gamma(t)$ for all $t \in \mathbb{R}_+$ and always rejects player 2's demand; with this notation, $\gamma \in \Sigma_1$. In other words, a posture is a pure bargaining phase strategy that does not condition on player 2's play or accept player 2's demand. After announcing posture γ, player 1 becomes committed to γ with some probability $\varepsilon > 0$, meaning that she must play strategy γ in the bargaining phase. With probability $1 - \varepsilon$, she is free to play any strategy in the bargaining phase. Whether or not player 1 becomes committed to γ is observed only by player 1.

2.2 Definition of Maxmin Payoff and Posture

This subsection defines player 1's maxmin payoff and posture. Intuitively, player 1's maxmin payoff is the highest payoff she can guarantee herself when all she knows about player 2 is that he is rational (i.e., maximizes his expected payoff given his belief about her behavior) and that he believes that she follows her announced posture γ with probability at least ε.

Formally, that player 2 is rational and assigns probability at least ε to player 1 following her announced posture γ means that his strategy satisfies the following condition:

Definition 1 *A strategy σ_2 of player 2's is* rational given posture γ *if there exists a belief π_2 of player 2's such that $\pi_2(\gamma) \geq \varepsilon$ and $\sigma_2 \in \Sigma_2^*(\pi_2)$.*

I assume that player 1's belief π_1 is *consistent with knowledge of rationality given pos-*

ture γ, in that every strategy $\sigma_2 \in \text{supp}(\pi_1)$ is rational given posture γ.[8] Let $\Pi_1^\gamma \equiv \Delta\{\sigma_2 : \sigma_2 \text{ is rational given posture } \gamma\}$ be set of beliefs π_1 that are consistent with knowledge of rationality given posture γ. Then the highest payoff that player 1 can guarantee herself after announcing posture γ is the following:

Definition 2 *Player* 1*'s* maxmin payoff given posture γ *is*

$$u_1^*(\gamma) \equiv \sup_{\sigma_1} \inf_{\pi_1 \in \Pi_1^\gamma} u_1(\sigma_1, \pi_1).$$

A strategy $\sigma_1^*(\gamma)$ *of player* 1*'s is a* maxmin strategy given posture γ *if*

$$\sigma_1^*(\gamma) \in \operatorname*{argmax}_{\sigma_1} \inf_{\pi_1 \in \Pi_1^\gamma} u_1(\sigma_1, \pi_1).$$

Equivalently, $u_1^*(\gamma)$ is the highest payoff player 1 can receive when she chooses a strategy σ_1 and then player 2 chooses a rational strategy σ_2 that minimizes $u_1(\sigma_1, \sigma_2)$; that is,

$$u_1^*(\gamma) = \sup_{\sigma_1} \inf_{\sigma_2 : \sigma_2 \text{ is rational given posture } \gamma} u_1(\sigma_1, \sigma_2).$$

In particular, to guarantee herself a high payoff, player 1 must play a strategy that does well against *any* rational strategy of player 2's.[9]

Finally, I define player 1's maxmin payoff, the highest payoff that player 1 can guarantee herself before announcing a posture, as well as the corresponding maxmin posture.

Definition 3 *Player* 1*'s* maxmin payoff *is*

$$u_1^* \equiv \sup_{\gamma} u_1^*(\gamma).$$

[8] A subtlety here is that the set $\Sigma_2^*(\pi_2)$ may be empty for some beliefs π_2. It is inevitable in bargaining models that players do not have best responses to all beliefs; for example, a player has no best response to the belief that her opponent will accept any strictly positive offer but will refuse an offer of 0. Thus, the assumption that player 2 plays a strategy $\sigma_2 \in \Sigma_2^*(\pi_2)$ for some belief π_2 is in fact a *joint* assumption on his belief and strategy. While the implied assumption on beliefs is certainly not without loss of generality, it is extremely natural and indeed fundamental, as it says precisely that player 2's choice set is non-empty. It is also weaker than any equilibrium assumption, as players play best responses in equilibrium.

[9] A potential criticism of the concept of the maxmin payoff given posture γ is that it appears to neglect the fact that, in the event that player 1 does become committed to posture γ, she is guaranteed only $\inf_{\pi_1 \in \Pi_1^\gamma} u_1(\gamma, \pi_1)$ in the bargaining phase, rather than $\sup_{\sigma_1} \inf_{\pi_1 \in \Pi_1^\gamma} u_1(\sigma_1, \pi_1)$. However, I show in Section 3.3 that these two numbers are actually identical in my model.

A posture γ^* *is a* maxmin posture *if there exists a sequence of postures* $\{\gamma_n\}$ *such that* $\gamma_n(t) \to \gamma^*(t)$ *for all* $t \in \mathbb{R}_+$ *and* $u_1^*(\gamma_n) \to u_1^*$.

I sometimes emphasize the dependence of u_1^* and γ^* on ε by writing $u_1^*(\varepsilon)$ and γ_ε^*.[10] Both the set of maxmin strategies given any posture γ and the set of maxmin postures are non-empty, though at this point this is not obvious.

Note that Definitions 2 and 3 are "non-Bayesian" in that they characterize the largest payoff that player 1 can guarantee herself, rather than the maximum payoff that she can obtain given some belief. However, repeating the analysis with the "Bayesian" version of these definitions (with the order of the sup and inf reversed) would yield the same results (see footnote 22).

Another reason for studying player 1's maxmin payoff is that it determines the entire range of payoffs that are consistent with knowledge of rationality, as shown by the following proposition.

Proposition 1 *For any posture* γ *and any payoff* $u_1 \in [u_1^*(\gamma), 1)$*, there exists a belief* $\pi_1 \in \Pi_1^\gamma$ *such that* $\max_{\sigma_1} u_1(\sigma_1, \pi_1) = u_1$.

3 Characterization of Maxmin Payoff and Posture

This section states and proves Theorem 1, the main result of the paper, which solves for player 1's maxmin payoff and posture. Section 3.1 states Theorem 1 and provides intuition, and Sections 3.2 through 3.4 provide the proof.

3.1 Main Result

The main result is the following:

Theorem 1 *Player* 1*'s maxmin payoff is*

$$u_1^*(\varepsilon) = 1/(1 - \log \varepsilon),$$

[10] The notation $\gamma^*(\cdot)$ is already taken by the time-t demand of posture γ^*. I apologize for abusing notation in writing $u_1^*(\gamma)$ and $u_1^*(\varepsilon)$ for different objects and hope that this will not cause confusion.

and the unique maxmin posture γ_ε^* is given by

$$\gamma_\varepsilon^*(t) = \min\left\{e^{rt}/(1-\log\varepsilon), 1\right\} \text{ for all } t \in \mathbb{R}_+.$$

Theorem 1 shows that player 1's maxmin payoff large relative to her commitment probability ε, and that her unique maxmin posture is simply demanding the maxmin payoff plus compensation for any delay in reaching agreement.[11] I first give intuition for why the maxmin posture demands the maxmin payoff plus compensation for delay, and then give intuition for why the maxmin payoff equals $1/(1-\log\varepsilon)$.

The first step (Sections 3.2 and 3.3) is solving the "min" in maxmin: that is, determining the worst belief that player 2 can have after player 1 announces an arbitrary posture γ. This belief is called the *γ-offsetting belief* and plays an important role in the analysis. A preliminary observation is that player 1 should mimic her announced posture γ forever in order to guarantee herself as much as possible (however ill-chosen γ may be). This is because player 1 is not guaranteed a positive payoff at histories following a deviation from her announced posture, because at such histories player 2's beliefs and strategy are unrestricted. It follows from this observation that the γ-offsetting belief is whatever belief leads player 2 to rejects player 1's demand for as long as possible when player 1 mimics γ.[12]

What belief leads player 2 to reject for as long as possible? It is the belief that player 1 is committed to γ with the smallest possible probability (i.e., ε), and that player 1 concedes the entire surplus to him at the rate that makes him (player 2) indifferent between accepting and rejecting. For if player 1 conceded more slowly, player 2 would accept, and if player 1 conceded more quickly, player 1 would build reputation more quickly when mimicking γ and

[11] The importance of non-constant postures is a difference between this paper and existing reputational bargaining models, where it is usually assumed that players may only be committed to strategies that demand a constant share of the surplus. A notable exception is Abreu and Pearce (2007), where players may be committed to non-constant postures that can also condition their play on their opponents' behavior. However, Abreu and Pearce's main result is that a particular posture that demands a constant share of the surplus is approximately optimal in their model, when commitment probabilities are small.

[12] This is not true if γ ever increases so quickly that delay benefits player 1. For this intuitive discussion, consider instead the "typical" case where delay hurts player 1.

would thus eventually have her demand accepted sooner.[13] Furthermore, this concession rate (call it $\lambda(t)$) is higher when player 1's demand is smaller, because when player 1's demand is smaller player 2 is more tempted to accept. Thus, the γ-offsetting belief is that player 1 is committed to γ with probability ε and concedes the entire surplus at rate $\lambda(t)$ (which depends on γ).

The second step (Section 3.4) is solving the "max" in maxmin: that is, determining the posture γ that maximizes player 1's payoff when player 2 has γ-offsetting beliefs. Since player 1 builds reputation more quickly when her demand is smaller, she benefits from demanding as little as possible, *subject to the constraint that she always demands at least her maxmin payoff plus compensation for delay* (as otherwise player 2 might rationally accept at a time where she demands less than this, leaving her with less than her maxmin payoff). This implies that the maxmin posture demands exactly the maxmin payoff plus compensation for delay, until player 1's reputation reaches 1 (i.e., until player 2 becomes certain that she is committed to γ).[14] It can also be shown that under the maxmin posture player 1's reputation reaches 1 at the same time at which her demand reaches 1, and that her demand can never subsequently drop below 1. So the maxmin posture demands compensation for delay until player 1's demand reaches 1, and subsequently demands 1 forever.

It remains to describe why the maxmin payoff equals $1/(1-\log\varepsilon)$. Consider a posture γ given by $\gamma(t) = \min\{e^{rt}u_1, 1\}$ for all $t \in \mathbb{R}_+$, for arbitrary $u_1 \in \mathbb{R}_+$; that is, γ demands u_1 plus compensation for delay. Observe that if player 1's reputation reaches 1 before her demand reaches 1, then player 2 must accept by the time her reputation reaches 1, as at that time he is certain that player 1's demand will only increase if he rejects further. Since γ demands compensation for delay until player 1's demand reaches 1, it follows that if player 1's reputation reaches 1 before her demand reaches 1 then player 1 is guaranteed a payoff

[13] This point is clearest in the extreme case where player 2 thinks that player 1 will surely concede in one second. Then player 1 needs only mimic γ for two seconds to convince player 2 that she is committed.

[14] A subtlety is that some postures γ that always demand more than the maxmin posture γ^* may have $\gamma'(t) > \gamma^{*\prime}(t)$ for some t and thus $\lambda(t) > \lambda^*(t)$ for some t (see the equation for $\lambda(t)$ below). However, it can be shown that this advantage in the derivative term γ' is always more than offset by the disadvantage in the level term γ when integrating λ over an interval, so that player 1's reputation is always greater with posture γ^* than with γ.

equal to her initial demand u_1. I now argue that player 1's reputation reaches 1 before her demand does whenever $u_1 < 1/(1-\log\varepsilon)$, which proves that player 1 can guarantee herself up to $1/(1-\log\varepsilon)$.

I first compute the concession rate $\lambda(t)$ that makes player 2 indifferent between accepting and rejecting player 1's demand $\gamma(t)$. For player 2, accepting yields flow payoff $r(1-\gamma(t))$, while rejecting yields flow payoff $\lambda(t)\gamma(t)-\gamma'(t)$, and equalizing these flow payoffs gives

$$\lambda(t) = \frac{r(1-\gamma(t))+\gamma'(t)}{\gamma(t)}.$$[15]

Since $\gamma(t) = e^{rt}u_1$ until $\gamma(t)$ reaches 1, it follows that $\lambda(t) = r/(e^{rt}u_1)$ until $\gamma(t)$ reaches 1. Now player 1's reputation reaches 1 at the time T such that the probability that player 1 has not conceded by T equals ε, since if player 1 does not concede by this time then she must be committed to γ (as γ never concedes). This time T is given by

$$\exp\left(-\int_0^T \lambda(t)\,dt\right) = \varepsilon.$$[16]

Substituting $r/(e^{rt}u_1)$ for $\lambda(t)$, this becomes

$$T = -\log(1+u_1\log\varepsilon)/r.$$

On the other hand, player 1's demand reaches 1 at the time T^1 given by $e^{rT^1}u_1 = 1$, or

$$T^1 = -\log(u_1)/r.$$

Hence, player 1's reputation reaches 1 before her demand does if and only if $T < T^1$, that is, if and only if $u_1 < 1/(1-\log\varepsilon)$. So player 1 can guarantee herself up to $1/(1-\log\varepsilon)$.

Finally, player 1 cannot guarantee herself more than $1/(1-\log\varepsilon)$, because it can be shown that any posture that guarantees close to $1/(1-\log\varepsilon)$ must be close to the maxmin posture γ^*.

[15] Two remarks: First, the formal definition of $\lambda(t)$ is provided in Section 3.2 (the current definition assumes that γ is differentiable and that $r(1-\gamma(t))+\gamma'(t)\geq 0$, for example). Second, a slightly more rigorous derivation of $\lambda(t)$ comes from considering the equation for player 2 to be indifferent between accepting at t and $t+dt$, $1-\gamma(t) = \lambda(t)\,dt + (1-\lambda(t)\,dt)(1-rdt)(1-\gamma(t+dt))$, and taking a first-order expansion in dt.

[16] A more general definition of T is provided in Section 3.2.

3.2 Offsetting Beliefs and Strategies

This subsection solves the problem

$$\inf_{(\pi_2,\sigma_2):\pi_2(\gamma)\geq\varepsilon,\sigma_2\in\Sigma_2^*(\pi_2)} u_1(\gamma,\sigma_2). \tag{1}$$

The resulting *(belief-strategy)* pair $(\pi_2^\gamma,\sigma_2^\gamma)$ are the *γ-offsetting belief and strategy.*

The key step in solving (1) is computing the smallest time T by which agreement must be reached under strategy profile (γ,σ_2) for $\sigma_2\in\Sigma_2^*(\pi_2)$. I then show that the value of (1) is simply $\min_{t\leq T} e^{-rt}\underline{\gamma}(t)$, where $\underline{\gamma}(t)\equiv\min\{\lim_{\tau\uparrow t}\gamma(\tau),\gamma(t)\}$.

Toward computing T, let $v(t,-1)$ be the continuation value of player 2 from best-responding to γ starting from date $(t,-1)$ (for integer t), and let $v(t)$ be the corresponding continuation value starting from date $(t,1)$ (for any $t\in\mathbb{R}_+$):

$$\begin{aligned} v(t,-1) &\equiv \max_{\tau\geq t} e^{-r(\tau-t)}\left(1-\underline{\gamma}(\tau)\right), \\ v(t) &\equiv \max\left\{1-\gamma(t), \sup_{\tau>t} e^{-r(\tau-t)}\left(1-\underline{\gamma}(\tau)\right)\right\}, \end{aligned} \tag{2}$$

where $\underline{\gamma}(\tau)\equiv\min\{\lim_{s\uparrow\tau}\gamma(s),\gamma(\tau)\}$. Thus, the difference between $v(t,-1)$ and $v(t)$ is that only $v(t,-1)$ gives player 2 the opportunity to accept the demand $1-\lim_{\tau\uparrow t}\gamma(\tau)$. In particular, $v(t,-1)=v(t)$ if γ (or v) is continuous at t. Note that $\max_{\tau\geq t}e^{-r(\tau-t)}\left(1-\underline{\gamma}(\tau)\right)$ is well-defined because $\underline{\gamma}$ is lower semi-continuous and $\lim_{\tau\to\infty}e^{-r(\tau-t)}\left(1-\underline{\gamma}(\tau)\right)=0$, and that v is continuous at all non-integer times t; let $\{s_1,s_2,\ldots\}\equiv S\subseteq\mathbb{N}$ be the set of discontinuity points of v. Finally, note that v can increase at rate no faster than r. That is, $v(t)\geq e^{-r(t'-t)}v(t')$ for all $t'\geq t$, because if $v(t')=e^{-r(\tau-t')}\left(1-\underline{\gamma}(\tau)\right)$ for some $\tau\geq t'$, then $v(t)\geq e^{-r(\tau-t)}\left(1-\underline{\gamma}(\tau)\right)=e^{-r(t'-t)}v(t')$. This implies that v is continuous but for downward jumps,[17] and that v is differentiable almost everywhere.[18] These are but two of the useful properties of the function v (which are not shared by γ) that reward working with v rather than γ in the subsequent analysis.

Next, I introduce two functions $\lambda:\mathbb{R}_+\to\mathbb{R}_+$ and $p:\mathbb{R}_+\to\mathbb{R}_+$ with the property that if player 1 mixes between mimicking γ and conceding the entire surplus to player 2, then $\lambda(t)$

[17] A function $f:\mathbb{R}\to\mathbb{R}$ is *continuous but for downward jumps* if $\liminf f_{x\uparrow x^*}(x)\geq f(x^*)\geq\limsup_{x\downarrow x^*}f(x)$ for all $x\in\mathbb{R}$.

[18] Proof: Let $f(t)=e^{-rt}v(t)$. Then f is non-increasing, which implies that f is differentiable almost everywhere (e.g., Royden, 1988, p. 100). Hence, v is differentiable almost everywhere.

(resp., $p(t)$) is the smallest non-negative hazard rate (resp., discrete probability) at which player 1 must concede in order for player 2 to be willing to reject player 1's time-t demand, $\underline{\gamma}(t)$. Let

$$\lambda(t) = \frac{rv(t) - v'(t)}{1 - v(t)} \tag{3}$$

if v is differentiable at t and $v(t) < 1$, and let $\lambda(t) = 0$ otherwise; note that $\lambda(t) \geq 0$ for all t, because v cannot increase at rate faster than r.[19] Also, let

$$p(t) = \frac{v(t, -1) - v(t)}{1 - v(t)} \tag{4}$$

if $v(t) < 1$, and let $p(t) = p(0) = 0$ otherwise.

When player 2 expects player 1 to accept his demand at rate (resp., probability) λ (resp., p), he becomes convinced that player 1 is committed to posture γ at the time $\tilde{T}$ defined in the following crucial lemma, which leads him to accept player 1's demand no later than the time T defined in the lemma. In the lemma, and throughout the paper, maximization or minimization over times t should be read as taking place over $t \in \mathbb{R}_+ \cup \{\infty\}$ (i.e., as allowing $t = \infty$, with the convention that $e^{-r\infty}\underline{\gamma}(\infty) \equiv 0$ for all postures γ).

Lemma 1 *Let*

$$\tilde{T} \equiv \sup\left\{t : \exp\left(-\int_0^t \lambda(s)\,ds\right) \prod_{s \in S \cap [0,t)} (1 - p(s)) > \varepsilon\right\},$$

and let

$$T \equiv \max \operatorname*{argmax}_{t \geq \tilde{T}} \begin{cases} e^{-rt}(1 - \gamma(t)) & \text{if } t = \tilde{T} \\ e^{-rt}\left(1 - \underline{\gamma}(t)\right) & \text{if } t > \tilde{T} \end{cases}.$$

Then, for any π_2 such that $\pi_2(\gamma) \geq \varepsilon$ and any $\sigma_2 \in \Sigma_2^*(\pi_2)$, agreement is reached no later than time T under strategy profile (γ, σ_2). In particular,

$$\inf_{(\pi_2,\sigma_2):\pi_2(\gamma)\geq\varepsilon,\sigma_2\in\Sigma_2^*(\pi_2)} u_1(\gamma, \sigma_2) \geq \min_{t \leq T} e^{-rt}\underline{\gamma}(t). \tag{5}$$

[19] If $v'(t) = 0$, then $\lambda(t)$ becomes the concession rate that makes player 2 indifferent between accepting and rejecting the constant offer $v(t)$, which is familiar from the literatures on wars of attrition and reputational bargaining. However, in these literatures $\lambda(t)$ is the rate at which player 1 concedes in equilibrium, while here is the rate at which player 1 concedes according to player 2's γ-offsetting belief, as will become clear.

Thus, Lemma 1 shows that agreement is delayed for as long as possible when player 1's concession rate and probability are given by λ and p. This gives a lower bound for (1).

The remainder of this subsection is devoted to showing that (5) holds with equality, which proves that (1) equals $\min_{t \leq T} e^{-rt} \underline{\gamma}(t)$. The idea is that player 2 may hold a belief that leads him to demand the entire surplus until time $t^* \equiv \min \operatorname{argmin}_{t \leq T} e^{-rt} \underline{\gamma}(t)$ and then accept player 1's offer; this is the γ-offsetting belief.[20] I first define the γ-offsetting belief, and then show that (5) holds with equality.

I begin by introducing a strategy, $\tilde{\gamma}$, which is used in defining the γ-offsetting belief.[21] Let

$$\chi(t) = \max \left\{ \frac{\exp\left(-\int_0^t \lambda(s)\, ds\right) \prod_{s \in S \cap [0,t)} (1 - p(s)) - \varepsilon}{\exp\left(-\int_0^t \lambda(s)\, ds\right) \prod_{s \in S \cap [0,t)} (1 - p(s))}, 0 \right\}; \tag{6}$$

let

$$\hat{\lambda}(t) = \frac{\lambda(t)}{\chi(t)} \tag{7}$$

if $\chi(t) > 0$, and let $\hat{\lambda}(t) = 0$ otherwise; and let

$$\hat{p}(t) = \min \left\{ \frac{p(t)}{\chi(t)}, 1 \right\} \tag{8}$$

if $\chi(t) > 0$, and let $\hat{p}(t) = 0$ otherwise. Thus, $\chi(t)$ is the posterior probability that player 2 assigns to player 1's playing a strategy other than γ at time t when player 1's *unconditional* concession rate and probability are $\lambda(t)$ and $p(t)$, and $\hat{\lambda}(t)$ and $\hat{p}(t)$ are the *conditional* (on not playing γ) concession rate and probability needed for the unconditional concession rate and probability to equal $\lambda(t)$ and $p(t)$.

Definition 4 *$\tilde{\gamma}$ is the strategy that demands $\gamma(t)$ at all t, accepts with hazard rate $\hat{\lambda}(t)$ at all $t < t^*$, accepts with probability $\hat{p}(t)$ at date $(t, 1)$ for all $t < t^*$, and rejects for all $t \geq t^*$, for all histories h^t.*

I now define the γ-offsetting belief. Throughout, a history h^{t-} (resp., h^{t+}) is *consistent with posture* γ if $u_1(\tau) = \gamma(\tau)$ for all $\tau < t$ (resp., $\tau \leq t$).

[20] Note that $\min \operatorname{argmin}_{t \leq T} e^{-rt} \underline{\gamma}(t)$ is well-defined, because $\underline{\gamma}(t)$ is lower semi-continuous (though it may equal ∞, if $T = \infty$). Note also that $\gamma(t) > 0$ for all $t < t^*$. This property of t^* makes it more convenient to define the γ-offsetting belief with reference to t^*, rather than some other element of $\operatorname{argmin}_{t \leq T} e^{-rt} \underline{\gamma}(t)$.

[21] This approach is related to a construction in Wolitzky (2011).

Definition 5 *The* γ-offsetting belief*, denoted* π_2^γ*, is given by* $\pi_2^\gamma(\gamma) = \varepsilon$ *and* $\pi_2^\gamma(\tilde{\gamma}) = 1 - \varepsilon$*. The* γ-offsetting strategy*, denoted* σ_2^γ*, is the strategy that always demands* 1 *and accepts or rejects player* 1*'s demand as follows:*

1. *If* h^t *is consistent with* γ*, then reject if* $t < t^*$*; accept at date* $(t^*, -1)$ *if and only if* $\lim_{\tau \uparrow t^*} \gamma(\tau) \leq \gamma(t^*)$*; accept at date* $(t^*, 1)$ *if and only if* $\lim_{\tau \uparrow t^*} \gamma(\tau) > \gamma(t^*)$*; and reject if* $t > t^*$*.*

2. *If* h^t *is inconsistent with* γ*, then reject.*

Finally, I show that (5) holds with equality, and also that the γ-offsetting *(belief, strategy)* pair $(\pi_2^\gamma, \sigma_2^\gamma)$ is a solution to (1). If $t^* = \infty$, then the following statement that agreement is reached at time t^* means that agreement is never reached:

Lemma 2 *Agreement is reached at time* t^* *under strategy profile* $(\gamma, \sigma_2^\gamma)$*, and* $\sigma_2^\gamma \in \Sigma_2^*(\pi_2^\gamma)$*. In particular, the pair* $(\pi_2^\gamma, \sigma_2^\gamma)$ *is a solution to (1), and* $u_1(\gamma, \sigma_2^\gamma) = \min_{t \leq T} e^{-rt} \underline{\gamma}(t)$*.*

Proof. It is immediate from Definition 5 that agreement is reached at t^* under strategy profile $(\gamma, \sigma_2^\gamma)$, which implies that $u_1(\gamma, \sigma_2^\gamma)$ equals $\min_{t \leq T} e^{-rt} \underline{\gamma}(t)$, the right-hand side of (5). Since $\pi_2^\gamma(\gamma) \geq \varepsilon$, it remains only to show that $\sigma_2^\gamma \in \Sigma_2^*(\pi_2^\gamma)$.

If $t < \min\left\{\tilde{T}, t^*\right\}$ and h^t is consistent with γ, then, by construction of $\tilde{\gamma}$, player 1 accepts player 2's demand of 1 with unconditional hazard rate $\lambda(t)$ and unconditional discrete probability $p(t)$ under π_2^γ. It is established in the proof of Lemma 1 that it is optimal for player 2 to demand $u_2(t) = 1$ and reject at any time $t < \min\left\{\tilde{T}, t^*\right\}$ when player 1 accepts his demand of 1 at rate λ and probability p until time $\tilde{T}$; and that in addition if $t^* < \tilde{T}$ then player 2 is indifferent between accepting and rejecting at time t^* when player 1 accepts with this rate and probability until time $\tilde{T}$. Therefore, it is optimal for player 2 to demand $u_2(t) = 1$ and reject at time t when player 1 accepts with this rate and probability only until time $\min\left\{\tilde{T}, t^*\right\}$.

If $t \in \left[\tilde{T}, t^*\right)$ and h^t is consistent with γ, then under π_2^γ player 2 is certain that player 1 is playing γ at h^t. Since $t^* \leq T$, this implies that it is optimal for player 2 to reject. If a history h^t is not reached under strategy profile $(\pi_2^\gamma, \sigma_2^\gamma)$ (as is the case if h^t is inconsistent with γ or if $t > t^*$), then any continuation strategy of player 2's is optimal. Finally, to see that accepting $\underline{\gamma}(t^*)$ (i.e., accepting at the more favorable of dates $(t^*,-1)$ and $(t^*,1)$) is optimal, note that the fact that $t^* \in \operatorname{argmin}_{t \le T} e^{-rt}\underline{\gamma}(t)$ implies that $\underline{\gamma}(t) \ge \underline{\gamma}(t^*)$ for all $t \in [t^*, T]$. Hence, $t^* \in \operatorname{argmax}_{t\in[t^*,T]} e^{-rt}\left(1-\underline{\gamma}(t)\right)$. Because $\tilde{\gamma}$ coincides with γ after time t^*, it follows that, conditional on having reached time t^*, player 2 receives at most $\sup_{t\in(t^*,T]} e^{-rt}\left(1-\underline{\gamma}(t)\right)$ if he rejects, and receives $e^{-rt^*}\left(1-\underline{\gamma}(t^*)\right)$ if he accepts, which is weakly more. Therefore, $\sigma_2^\gamma \in \Sigma_2^*(\pi_2^\gamma)$. ■

3.3 Maxmin Strategies

This subsection shows that γ itself is a maxmin strategy given posture γ. Henceforth, I write $\tilde{T}(\gamma)$ and $T(\gamma)$ for the times defined in Lemma 1, making the dependence on γ explicit.

Lemma 3 *For any posture γ, $u_1^*(\gamma) = u_1(\gamma, \sigma_2^\gamma) = \min_{t \le T(\gamma)} e^{-rt}\underline{\gamma}(t)$.*

Proof. By Lemma 2, $(\pi_2^\gamma, \sigma_2^\gamma)$ is a solution to (1), so

$$\sigma_2^\gamma \in \underset{\pi_1 \in \Pi_1^\gamma}{\operatorname{argmin}}\, u_1(\gamma, \pi_1). \tag{9}$$

Under strategy σ_2^γ, player 2 always demands 1 and only accepts player 1's demand if she conforms to γ through time t^*. Hence, $\sup_{\sigma_1} u_1(\sigma_1, \sigma_2^\gamma) = e^{-rt^*}\underline{\gamma}(t^*) = u_1(\gamma, \sigma_2^\gamma)$, and therefore

$$\gamma \in \underset{\sigma_1}{\operatorname{argmax}}\, u_1(\sigma_1, \sigma_2^\gamma). \tag{10}$$

Now (9) and (10) imply the following chain of inequalities:

$$\begin{aligned}
\sup_{\sigma_1} \inf_{\pi_1 \in \Pi_1^\gamma} u_1(\sigma_1, \pi_1) &\ge \inf_{\pi_1 \in \Pi_1^\gamma} u_1(\gamma, \pi_1) \\
&= u_1(\gamma, \sigma_2^\gamma) \quad \text{(by (9))} \\
&= \sup_{\sigma_1} u_1(\sigma_1, \sigma_2^\gamma) \quad \text{(by (10))} \\
&\ge \sup_{\sigma_1} \inf_{\pi_1 \in \Pi_1^\gamma} u_1(\sigma_1, \pi_1).
\end{aligned}$$

This is possible only if both inequalities hold with equality (and the supremum and infimum are attained at γ and σ_2^γ, respectively).[22] Therefore, $u_1^*(\gamma) = u_1(\gamma, \sigma_2^\gamma) = \min_{t \le T(\gamma)} e^{-rt}\underline{\gamma}(t)$.
■

[22] The same argument applies with the order of the sup and inf reversed, which is why Theorem 1 continues to hold when the order of the sup and inf are reversed in Definitions 2 and 3.

3.4 Proof of Theorem 1

I now sketch the remainder of the proof of Theorem 1. The details of the proof are deferred to the appendix.

The first part of the proof is constructing a sequence of postures $\{\gamma_n\}$ such that $\lim_{n\to\infty} u_1^*(\gamma_n) = 1/(1-\log\varepsilon)$ and $\{\gamma_n(t)\}$ converges to $\gamma^*(t) \equiv \min\{e^{rt}/(1-\log\varepsilon), 1\}$ for all $t \in \mathbb{R}_+$. Define γ_n by

$$\gamma_n(t) = \min\left\{\left(\frac{n}{n+1}\right)\frac{e^{rt}}{1-\log\varepsilon}, 1\right\} \text{ for all } t \in \mathbb{R}_+.$$

Let T_n^1 be the time where $\gamma_n(t)$ reaches 1. It can be shown that $T_n^1 > \tilde{T}(\gamma_n)$ for all $n \in \mathbb{N}$. Hence, $\gamma_n(t) = \left(\frac{n}{n+1}\right)\frac{e^{rt}}{1-\log\varepsilon}$ for all $t \leq \tilde{T}(\gamma_n)$, and $\gamma_n\left(\tilde{T}(\gamma_n)\right) < 1$. Since $\gamma_n(t)$ is non-decreasing and $\gamma_n\left(\tilde{T}(\gamma_n)\right) < 1$, it follows from the definition of $T(\gamma_n)$ that $T(\gamma_n) = \tilde{T}(\gamma_n)$. Thus, by Lemma 3,

$$u_1^*(\gamma_n) = \min_{t\leq T(\gamma_n)} e^{-rt}\underline{\gamma_n}(t) = \min_{t\leq\tilde{T}(\gamma_n)}\left(\frac{n}{n+1}\right)\frac{1}{1-\log\varepsilon} = \left(\frac{n}{n+1}\right)\frac{1}{1-\log\varepsilon}.$$

Therefore, $\lim_{n\to\infty} u_1^*(\gamma_n) = 1/(1-\log\varepsilon)$.

The second part is showing that no posture γ guarantees more than $1/(1-\log\varepsilon)$. Here, the crucial observation is that any posture γ such that $\gamma(t) \geq e^{rt}/(1-\log\varepsilon)$ for all $t \leq T(\gamma)$ satisfies $\tilde{T}(\gamma) \geq T^1$, where T^1 is the time at which $\gamma^*(t)$ reaches 1. Since $u_1^*(\gamma) = \min_{t\leq T(\gamma)} e^{-rt}\underline{\gamma}(t)$, this implies that any posture γ that guarantees at least $1/(1-\log\varepsilon)$ must satisfy $\tilde{T}(\gamma) \geq T^1$; in particular, player 2 may reject $\gamma(t)$ until time T^1. But receiving the entire surplus at T^1 is worth only $1/(1-\log\varepsilon)$, so it follows that no posture guarantees more than $1/(1-\log\varepsilon)$. The appendix shows that in addition the maxmin posture is unique.

4 Extensions

This section presents three extensions of Theorem 1. Section 4.1 restricts player 1 to announcing constant postures; Section 4.2 allows for heterogeneous discounting; and Section 4.3 considers higher-order knowledge of rationality.

4.1 Constant Postures

Theorem 1 shows that the unique maxmin posture is non-constant. In this subsection, I determine how much lower a player's maxmin payoff is when she is required to announce a constant posture. This establishes that value of announcing non-constant postures, and also facilitates comparison with the existing reputational bargaining literature, in which typically players can only announce constant postures.

A posture γ is *constant* if $\gamma(t) = \gamma(0)$ for all t. If γ is constant, I slightly abuse notation by writing γ for the constant demand $\gamma(t)$ in addition to the posture itself. The constant posture γ that maximizes $u_1^*(\gamma)$ is the *maxmin constant posture*, denoted $\bar{\gamma}^*$, and the corresponding payoff is the *maxmin constant payoff*, denoted $\bar{u}_1^*$. These can be derived using Lemmas 1 through 3, leading to the following:

Proposition 2 *For all $\varepsilon < 1$, the unique maxmin constant posture is $\bar{\gamma}_\varepsilon^* = \frac{2-\log\varepsilon-\sqrt{(\log\varepsilon)^2-4\log\varepsilon}}{2}$, and the maxmin constant payoff is $\bar{u}_1^*(\varepsilon) = \exp(-(1-\bar{\gamma}_\varepsilon^*))\bar{\gamma}_\varepsilon^*$.*

Proposition 2 solves for $\bar{\gamma}_\varepsilon^*$ and $\bar{u}_1^*(\varepsilon)$, but it does not yield a clear relationship between the maxmin constant payoff, $\bar{u}_1^*(\varepsilon)$, and the (overall) maxmin payoff, $u_1^*(\varepsilon)$. Figure 2 graphs $u_1^*(\varepsilon)$ and $\bar{u}_1^*(\varepsilon)$. In addition, the following result regarding the ratio of $u_1^*(\varepsilon)$ to $\bar{u}_1^*(\varepsilon)$ is straightforward:

Corollary 1 *$u_1^*(\varepsilon)/\bar{u}_1^*(\varepsilon)$ is decreasing in ε, $\lim_{\varepsilon\to1} u_1^*(\varepsilon)/\bar{u}_1^*(\varepsilon) = 1$, and $\lim_{\varepsilon\to0} u_1^*(\varepsilon)/\bar{u}_1^*(\varepsilon) = e$.*

The most interesting part of Corollary 1 is that a player's maxmin payoff is approximately $e \approx 2.72$ times greater when she can announce non-constant postures than when she can only announce constant postures, when her commitment probability is small. Thus, there is a large advantage to announcing non-constant postures. However, a player can still guarantee herself a substantial share of the surplus when she can only announce constant postures, and her maxmin payoff goes to 0 with ε at the same rate in either case.

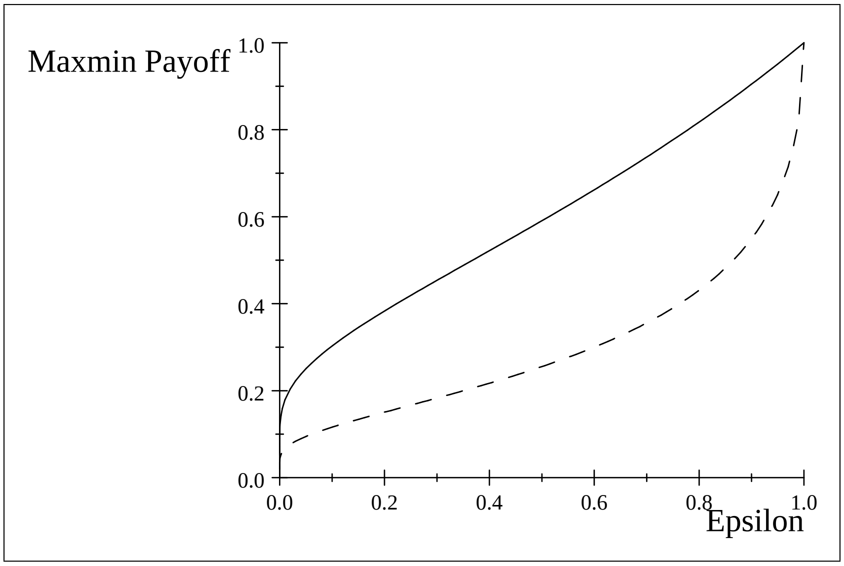

Figure 2: $u_1^*(\varepsilon)$ (solid line) and $\bar{u}_1^*(\varepsilon)$ (dashed line).

4.2 Heterogeneous Discounting

The assumption that the players have the same discount rate has simplified notation and led to simple formulas for $u_1^*(\varepsilon)$ and γ_ε^* in Theorem 1. However, it is straightforward to let player i have discount rate r_i, with $r_i \neq r_j$, and doing so yields interesting comparative statics with respect to the players' relative patience, r_1/r_2 (as will become clear, u_1^* depends on r_1 and r_2 only through r_1/r_2). First, the standard result in the reputational bargaining literature that player 1's sequential equilibrium payoff converges to 1 as r_1/r_2 converges to 0, and converges to 0 as r_1/r_2 converges to ∞, also applies to player 1's maxmin payoff. This is analogous to the finding of Watson (1993) and Battigalli and Watson (1997) that the limit uniqueness result of Fudenberg and Levine (1989) holds under knowledge of rationality. However, I also derive player 1's maxmin payoff for fixed r_1/r_2 (rather than only in the limit). This leads to a second comparative static result, which indicates that a geometric change in relative patience has a similar effect on the maxmin payoff as does an exponential change in commitment probability. An analogous result holds in equilibrium in existing reputational bargaining models.

I first present the analog of Theorem 1 for heterogeneous discount rates, and then state the two comparative statics results as corollaries, omitting their proofs.

Proposition 3 *If player i's discount rate is r_i, then player 1's maxmin payoff, $u_1^*(\varepsilon)$, is the unique number u_1^* that solves*

$$u_1^* = \frac{1}{1 - \frac{r_1}{r_2}\log\varepsilon - \left(\frac{r_1}{r_2} - 1\right)\log u_1^*}. \tag{11}$$

Corollary 2 shows that the standard limit comparative statics on r_1/r_2 in reputational bargaining models require only first-order knowledge of rationality.

Corollary 2 $\lim_{r_1/r_2\to 0} u_1^*(\varepsilon) = 1$. *If $\varepsilon < 1$, then in addition* $\lim_{r_1/r_2\to\infty} u_1^*(\varepsilon) = 0$.

Corollary 3 shows that the commitment probability ε must decrease exponentially to (approximately) offset a geometric increase in relative patience $(r_1/r_2)^{-1}$. The result is stated for the case $r_1/r_2 \le 1$, where even an exponential decrease in ε does not fully offset a geometric increase in $(r_1/r_2)^{-1}$. If $r_1/r_2 > 1$, then an exponential decrease in ε can more than offset a geometric increase in $(r_1/r_2)^{-1}$.

Corollary 3 *Suppose that $r_1/r_2 \le 1$ and that r_1/r_2 and ε both decrease while $(r_1/r_2)\log\varepsilon$ remains constant. Then $u_1^*(\varepsilon)$ increases.*

4.3 Rationalizability

Theorem 1 derives the highest payoff that player 1 can guarantee herself under first-order knowledge of rationality, the weakest epistemic assumption consistent with the possibility of reputation-building. I now show that player 1 cannot guarantee herself more than this under the much stronger assumption of normal-form rationalizability (or under any finite-order knowledge of rationality), which reinforces Theorem 1 substantially.[23]

I consider the following definition of (normal-form) rationalizability:

Definition 6 *A set of bargaining phase strategy profiles $\Omega = \Omega_1 \times \Omega_2 \subseteq \Sigma_1 \times \Sigma_2$ is* closed under rational behavior given posture γ *if for all $\sigma_1 \in \Omega_1$ there exists some belief $\pi_1 \in \Delta(\Omega_2)$ such that $\sigma_1 \in \Sigma_1^*(\pi_1)$; and for all $\sigma_2 \in \Omega_2$ there exists some belief $\pi_2 \in \Delta(\Omega_1 \cup \{\gamma\})$ such that $\pi_2(\gamma) \ge \varepsilon$, with strict inequality only if $\gamma \in \Omega_1$, and $\sigma_2 \in \Sigma_2^*(\pi_2)$.*

[23] The supplementary appendix shows that Theorem 1 also extends to iterated conditional dominance.

The set of rationalizable strategies given posture γ *is*

$$\Omega^{RAT}(\gamma) \equiv \bigcup \{\Omega : \Omega \textit{ is closed under rational behavior given posture } \gamma\}.$$

Player 1*'s* rationalizable maxmin payoff given posture γ *is* $u_1^{RAT}(\gamma) \equiv \sup_{\sigma_1} \inf_{\sigma_2 \in \Omega_2^{RAT}(\gamma)} u_1(\sigma_1, \sigma_2)$. *Player* 1*'s* rationalizable maxmin payoff *is* $u_1^{RAT} \equiv \sup_\gamma u_1^{RAT}(\gamma)$. *A posture* γ^{RAT} *is a* rationalizable maxmin posture *if there exists a sequence of postures* $\{\gamma_n\}$ *such that* $\gamma_n(t) \to \gamma^{RAT}(t)$ *for all* $t \in \mathbb{R}_+$ *and* $u_1^{RAT}(\gamma_n) \to u_1^{RAT}$.

The result is the following:

Proposition 4 *Player* 1*'s rationalizable maxmin payoff equals her maxmin payoff, and the unique rationalizable maxmin posture is the unique maxmin posture. That is,* $u_1^{RAT} = u_1^*$, *and the unique rationalizable maxmin posture is* $\gamma^{RAT} = \gamma^*$.

Any rationalizable strategy given posture γ is also rational given posture γ. Therefore, Lemma 1 applies under rationalizability. The only additional fact used in the proof of Theorem 1 is that $u_1^*(\gamma) = \min_{t \leq T(\gamma)} e^{-rt}\underline{\gamma}(t)$ for any posture γ (Lemma 3). Supposing that the analogous equation holds under rationalizability (i.e., that $u_1^{RAT}(\gamma) = \min_{t \leq T(\gamma)} e^{-rt}\underline{\gamma}(t)$), the proof of Theorem 1 goes through as written. Hence, to prove Proposition 4 it suffices to prove the following lemma, the proof of which shows that the γ-offsetting belief and strategy are not only rational but rationalizable:

Lemma 4 *For any posture* γ, $u_1^{RAT}(\gamma) = \min_{t \leq T(\gamma)} e^{-rt}\underline{\gamma}(t)$.

5 Two-Sided Commitment

This section introduces the possibility that both players may announce—and become committed to—postures prior to the start of bargaining. I show that each player i's maxmin payoff is close to that derived in Section 3 when her opponent's commitment probability, ε_j, is small in absolute terms (even if ε_j is large relative to ε_i). In addition, each player's maxmin posture is exactly as in Section 3. This shows that the analysis of Section 3 provides a two-sided theory of reputational bargaining. The results of this section contrast with the existing reputational bargaining literature, which emphasizes that *relative* commitment probabilities are crucial for determining *equilibrium* behavior and payoffs.

Formally, modify the model of Section 2 by assuming that in the announcement stage players simultaneously announce postures (γ_1, γ_2), to which they become committed with probabilities ε_1 and ε_2, respectively. The bargaining phase is unaltered. Thus, at the beginning of the bargaining phase, player i believes that player j is committed to posture γ_j with probability ε_j and is rational with probability $1-\varepsilon_j$ (though this fact is not common knowledge). The following definitions are analogs of Definitions 1 through 3 that allow for the fact that both players may become committed to the postures they announce:

Definition 7 *A belief π_i of player i's is* consistent with knowledge of rationality given postures (γ_i, γ_j) *if $\pi_i(\gamma_j) \geq \varepsilon_j$; $\pi_i(\gamma_j) > \varepsilon_j$ only if there exists π_j such that $\pi_j(\gamma_i) \geq \varepsilon_i$ and $\gamma_j \in \Sigma_j^*(\pi_j)$; and, for all $\sigma_j \neq \gamma_j$, $\sigma_j \in \operatorname{supp}(\pi_i)$ only if there exists π_j such that $\pi_j(\gamma_i) \geq \varepsilon_i$ and $\sigma_j \in \Sigma_j^*(\pi_j)$. Let $\Pi_i^{\gamma_i,\gamma_j}$ be the set of player i's beliefs that are consistent with knowledge of rationality given postures (γ_i, γ_j). Player i's* maxmin payoff given postures (γ_i, γ_j) *is*

$$u_i^*(\gamma_i, \gamma_j) \equiv \sup_{\sigma_i} \inf_{\pi_i \in \Pi_i^{\gamma_i,\gamma_j}} u_i(\sigma_i, \pi_i).$$

Player i's maxmin payoff *is*

$$u_i^* \equiv \sup_{\gamma_i} \inf_{\gamma_j} u_i^*(\gamma_i, \gamma_j).$$

A posture γ_i^* is a maxmin posture (of player i's) *if there exists a sequence of postures $\{\gamma_n\}$ such that $\gamma_n(t) \to \gamma_i^*(t)$ for all $t \in \mathbb{R}_+$ and $\inf_{\gamma_j} u_i^*(\gamma_n, \gamma_j) \to u_i^*$.*

I now show that $u_i^*(\varepsilon_i, \varepsilon_j)$ (player i's maxmin payoff with commitment probabilities ε_i and ε_j) is approximately equal to $u_i^*(\varepsilon_i)$ (her maxmin payoff in the one-sided commitment model) whenever ε_j is small, and that the maxmin posture is exactly as in the one-sided commitment model. This is simply because player i cannot guarantee herself anything in the event that player j is committed (e.g., if player j's announced posture always demands the entire surplus), which implies that player i guarantees herself as much as possible by conditioning on the event that player j is not committed. In this event, which occurs with probability $1-\varepsilon_j$, player i can guarantee herself $u_i^*(\varepsilon_i)$, and the only way she can guarantee herself this much is by announcing $\gamma_{\varepsilon_i}^*$.

Theorem 2 *Player i's maxmin payoff is* $u_i^*(\varepsilon_i,\varepsilon_j) = (1-\varepsilon_j)\,u_i^*(\varepsilon_i)$*, and player i's unique maxmin posture is* $\gamma^*_{i,(\varepsilon_i,\varepsilon_j)} = \gamma^*_{\varepsilon_i}$.

Proof. Let γ_j^0 be the posture of player j's given by $\gamma_j^0(t) = 1$ for all t. Note that $u_i\left(\sigma_i,\gamma_j^0\right) = 0$ for all σ_i. Therefore, $\inf_{\gamma_j} u_i\left(\sigma_i,\gamma_j\right) = 0$ for all σ_i.

Next, let $\Pi_i^{\gamma_i,\gamma_j}(\varepsilon_i,\varepsilon_j)$ be the set of beliefs π_i that are consistent with knowledge of rationality for commitment probabilities $(\varepsilon_i,\varepsilon_j)$, and let $\Pi_i^{\gamma_i}(\varepsilon_i)$ be the analogous set in the one-sided commitment model. I claim that if $\pi_i \in \Pi_i^{\gamma_i,\gamma_j}(\varepsilon_i,\varepsilon_j)$, then there exists $\pi_i' \in \Pi_i^{\gamma_i}(\varepsilon_i)$ such that π_i puts probability $1-\varepsilon_j$ on strategy π_i' and puts probability ε_j on strategy γ_j. To see this, note that $\pi_i\left(\gamma_j\right) \geq \varepsilon_j$, so there exists a strategy π_i' such that π_i puts probability $1-\varepsilon_j$ on π_i' and puts probability ε_j on γ_j. Furthermore, by definition of $\Pi_i^{\gamma_i,\gamma_j}(\varepsilon_i,\varepsilon_j)$, $\sigma_j \in \operatorname{supp}(\pi_i')$ only if there exists π_j such that $\pi_j(\gamma_i) \geq \varepsilon_i$ and $\sigma_j \in \Sigma_j^*(\pi_j)$ (whether or not σ_j equals γ_j). By definition of $\Pi_i^{\gamma_i}(\varepsilon_i)$, this implies that $\pi_i' \in \Pi_i^{\gamma_i}(\varepsilon_i)$.

Combining the above observations,

$$\begin{aligned}
\inf_{\gamma_j} u_i^*\left(\gamma_i,\gamma_j\right) &= \inf_{\gamma_j}\sup_{\sigma_i}\inf_{\pi_i\in\Pi_i^{\gamma_i,\gamma_j}(\varepsilon_i,\varepsilon_j)} u_i(\sigma_i,\pi_i) \\
&= \inf_{\gamma_j}\sup_{\sigma_i}\inf_{\pi_i'\in\Pi_i^{\gamma_i}(\varepsilon_i)} (1-\varepsilon_j)\,u_i(\sigma_i,\pi_i') + \varepsilon_j u_i\left(\sigma_i,\gamma_j\right). \\
&= \sup_{\sigma_i}\inf_{\pi_i'\in\Pi_i^{\gamma_i}(\varepsilon_i)} (1-\varepsilon_j)\,u_i(\sigma_i,\pi_i') + \varepsilon_j(0) \\
&= (1-\varepsilon_j)\,u_i^*(\gamma_i).
\end{aligned}$$

Therefore, the definitions of $u_i^*(\varepsilon_i,\varepsilon_j)$ and $u_i^*(\varepsilon_i)$ imply that $u_i^*(\varepsilon_i,\varepsilon_j) = \sup_{\gamma_i}(1-\varepsilon_j)\,u_i^*(\gamma_i) = (1-\varepsilon_j)\,u_i^*(\varepsilon_i)$. Similarly, the definition of a maxmin posture in the one-sided commitment model implies that $\gamma^*_{i,(\varepsilon_i,\varepsilon_j)}$ is a maxmin posture in the two-sided commitment model if and only if it is a maxmin posture in the one-sided commitment model with $\varepsilon = \varepsilon_i$. ■

Note that the assumption that both players best-respond to beliefs that are consistent with knowledge of rationality does not determine how bargaining proceeds. However, it is interesting to compare a player's opponent's worst-case conjecture (i.e., offsetting belief) about her strategy with her equilibrium strategy in, for example, Abreu and Gul (2000). Both in equilibrium and in her opponent's worst-case conjecture, a player mixes between mimicking her announced posture and (in effect) conceding. In the worst-case conjecture, a player concedes at the rate that makes her opponent indifferent between accepting and rejecting her demand when she follows her maxmin posture and he demands the entire surplus. In equilibrium, a player concedes at the rate that makes her opponent indifferent between accepting and rejecting when both players make their equilibrium demands. There is no general way to order the concession rates in equilibrium and in the worst-case conjecture, because a player's demand is often higher in equilibrium (implying a lower concession rate), while her opponent's demand is always lower in equilibrium (implying a higher concession rate). Thus, a player's concession rate in equilibrium and in her opponent's worst-case conjecture are determined by similar indifference conditions, but one cannot predict whether agreement will be reached sooner in equilibrium or under knowledge of rationality (for either the players' true strategies or for their opponents' worst-case conjectures).

6 Conclusion

This paper analyzes a model of reputational bargaining in which players initially announce postures to which they may become committed and then bargain over a unit of surplus. It shows that under first-order knowledge of rationality a player can guarantee herself a share of the surplus that is large relative to her probability of becoming committed, and that the unique bargaining posture that guarantees this much is simply demanding this share in addition to compensation for any delay in reaching agreement. These insights apply for one- or two-sided commitment, for heterogeneous discounting, for any level of knowledge of rationality or iterated conditional dominance, and for any bargaining procedure with frequent offers. In addition, if a player could only announce postures that always demand the same share of the surplus (as in most of the existing literature), her maxmin payoff would be approximately e times lower.

These results are intended to complement the existing equilibrium analysis of reputational bargaining models. Consider the fundamental question, "What posture should a bargainer stake out?" In equilibrium analysis, the answer to this question depends on her opponent's beliefs about her continuation play following every possible announcement. Yet it may be impossible for either the bargainer or an outside observer to learn these beliefs, especially when bargaining is one-shot. Hence, an appealing alternative approach is to look for a posture that guarantees a high payoff against *any* belief of one's opponent, and for the highest payoff that each player can guarantee herself. This paper shows that this alternative approach yields sharp and economically plausible results.

Appendix: Omitted Proofs

Proof of Proposition 1. [24]Fix a posture γ and payoff $u_1 \in [u_1^* (\gamma), 1)$. If $u_1 \neq \gamma(0)$, then let $\hat{\sigma}_2^\gamma$ be identical to the γ-offsetting strategy defined in Definition 5, with the modification that player 1's demand is accepted at any history h^t at which player 1 has demanded u_1 at all previous dates.[25] If $u_1 = \gamma(0)$, then let $\hat{\sigma}_2^\gamma$ be identical to the γ-offsetting strategy defined in Definition 5, with the modification that player 1's demand is accepted at date $(-\log(u_1)/r, -1)$ if player 1 has demanded 1 at all previous dates. In either case, let π_2^γ be as in Definition 5, and note that $\pi_2^\gamma(\gamma) \geq \varepsilon$. If $u_1 \neq \gamma(0)$, no strategy under which $u_1(0) = u_1$ is in the support of π_2^γ; similarly, if $u_1 = \gamma(0)$, no strategy under which $u_1(0) = 1$ is in the support of π_2^γ (since $u_1 < 1$). Therefore, the same argument as in the proof of Lemma 2 shows that $\hat{\sigma}_2^\gamma \in \Sigma_2^*(\pi_2^\gamma)$. Hence, the belief $\hat{\pi}_1$ given by $\hat{\pi}_1(\hat{\sigma}_2^\gamma) = 1$ is an element of Π_1^γ. Furthermore, under strategy $\hat{\sigma}_2^\gamma$, player 2 always demands 1 and only accepts player 1's demand if player 1 has either conformed to γ through time t^* (defined in Section 3.2) or has always demanded u_1 (in the $u_1 \neq \gamma(0)$ case) or 1 (in the $u_1 = \gamma(0)$ case). Note that $\exp(-r(-\log(u_1)/r)) = u_1$. Hence, in either case, $u_1(\sigma_1, \hat{\pi}_1) \in \{0, u_1^*(\gamma), u_1\}$ for every strategy σ_1. Let $\hat{\sigma}_1$ be the strategy of player 1's that always demands u_1 (if $u_1 \neq \gamma(0)$) or 1 (if $u_1 = \gamma(0)$) and never accepts player 2's demand. Then $u_1(\hat{\sigma}_1, \hat{\pi}_1) = u_1 = \max_{\sigma_1} u_1(\sigma_1, \hat{\pi}_1)$, completing the proof. ■

Proof of Lemma 1. I prove the result for pure strategies σ_2, which immediately implies the result for mixed strategies.

Fix π_2 such that $\pi_2(\gamma) \geq \varepsilon$ and pure strategy $\sigma_2 \in \Sigma_2^*(\pi_2)$. The plan of the proof is to show that if agreement is not reached by $\tilde{T}$ under strategy profile (γ, σ_2), then player 2 must be certain that player 1 is playing γ at any time $t > \tilde{T}$. This suffices to prove the lemma,

[24] The proof uses results from Section 3, and therefore should not be read before reading Section 3.

[25] That is, modify the second part of Definition 5 to include this contingency.

because $\sigma_2 \in \Sigma_2^*(\pi_2)$ implies that player 2 accepts $\underline{\gamma}(t)$ no later than time $t = T$ if at any time $t > \tilde{T}$ agreement has not been reached and he is certain that player 1 is playing γ.

I begin by introducing some notation. Let $\chi^{(\pi_2,\sigma_2)}(t)$ be the probability that player 2 assigns to player 1 *not* playing γ at date $(t,-1)$ when his initial belief is π_2 and play up until date $(t,-1)$ is given by player 1's following strategy γ and player 2's following (pure) strategy σ_2; this is determined by Bayes' rule, because $\pi_2(\gamma) \geq \varepsilon > 0$. By convention, if agreement is reached at time τ, let $\chi^{(\pi_2,\sigma_2)}(t) = \chi^{(\pi_2,\sigma_2)}(\tau)$ for all $t > \tau$. Let $t(\gamma,\sigma_2)$ be the time at which agreement is reached under strategy profile (γ,σ_2) (with the convention that $t(\gamma,\sigma_2) \equiv \infty$ if agreement is never reached under (γ,σ_2)); and let

$$\hat{t}(\gamma,\sigma_2,\pi_2) \equiv \sup\left\{t : \chi^{(\pi_2,\sigma_2)}(t) > 0\right\},$$

the latest time at which player 2 is not certain that player 1 is playing γ under strategy profile (γ,σ_2) with belief π_2. Let

$$\hat{T} \equiv \sup_{(\pi_2,\sigma_2):\pi_2(\gamma)\geq\varepsilon,\sigma_2\in\Sigma_2^*(\pi_2),t(\gamma,\sigma_2)\geq\hat{t}(\gamma,\sigma_2,\pi_2)} \hat{t}(\gamma,\sigma_2,\pi_2). \tag{12}$$

That is, $\hat{T}$ is the latest possible time t at which player 2 is not certain that player 1 is following γ and agreement is not reached by t. The remainder of the proof consists of showing that $\hat{T} = \tilde{T}$.

There are three steps involved in showing that $\hat{T} = \tilde{T}$, that is, that the value of the program (12) is $\tilde{T}$. Step 1 shows that in solving (12) one can restrict attention to a simple class of belief-strategy pairs (π_2,σ_2). Step 2 reduces the constraints that $\sigma_2 \in \Sigma_2^*(\pi_2)$ and $t(\gamma,\sigma_2) \geq \hat{t}(\gamma,\sigma_2,\pi_2)$ to an infinite system of inequalities involving player 1's concession rate and probability. Step 3 solves the reduced program.

Step 1: In the definition of $\hat{T}$ it is without loss of generality to restrict attention to (π_2,σ_2) such that σ_2 always demands 1, π_2 puts probability 1 on player 1 conceding at any history h^{t+} at which $u_1(t) \neq \gamma(t)$, and π_2 puts probability 0 on player 1 conceding at any history h^{t-}. That is, that the right-hand side of (12) continues to equal $\hat{T}$ when this additional constraint is imposed on (π_2,σ_2).

Proof. Suppose that (π_2',σ_2') satisfies $\pi_2'(\gamma) \geq \varepsilon$, $\sigma_2' \in \Sigma_2^*(\pi_2')$, and $t(\gamma,\sigma_2') \geq \hat{t}(\gamma,\sigma_2',\pi_2')$ (the constraints of (12)). Let π_2 be the belief under which player 1 demands $\gamma(t)$ for all $t \in \mathbb{R}_+$; accepts player 2's demand at every history of the form $(\gamma(\tau), 1)_{\tau \leq t}$ at the same rate and probability at which player 1 deviates from γ at time t (i.e., at date $(t, -1)$, $(t, 0)$, or $(t, 1)$) under strategy profile (π_2', σ_2'); and rejects player 2's demand at every other history. Clearly, there exists a strategy $\sigma_2 \in \Sigma_2^*(\pi_2)$ that always demands 1 and that in addition rejects player 1's demand whenever rejection is optimal under belief π_2. Note that player 1's rate and probability of deviating from γ at history $(\gamma(\tau), 1)_{\tau \leq t}$ under belief π_2 is the same as at time t under strategy profile (π_2', σ_2'), and that player 2's continuation payoff after such a deviation is weakly higher in the former case. Recall that strategy γ never accepts player 2's demand, so agreement is reached only if player 2 accepts player 1's demand or if player 1 has deviated from γ. Therefore, since rejecting player 1's demand $\gamma(t)$ under strategy profile (π_2', σ_2') is optimal for all $t < t(\gamma, \sigma_2')$, it follows that rejecting player 1's demand $\gamma(t)$ at history $(\gamma(\tau), 1)_{\tau \leq t}$ is optimal under belief π_2, for all $t < t(\gamma, \sigma_2')$. Since σ_2 prescribes rejection whenever it is optimal, this implies that $t(\gamma, \sigma_2) \geq t(\gamma, \sigma_2')$. Furthermore, $\chi^{(\pi_2, \sigma_2)}(t) = \chi^{(\pi_2', \sigma_2')}(t)$ for all $t \in \mathbb{R}_+$, so $\hat{t}(\gamma, \sigma_2, \pi_2) = \hat{t}(\gamma, \sigma_2', \pi_2')$. Hence, $t(\gamma, \sigma_2) \geq \hat{t}(\gamma, \sigma_2, \pi_2)$. Finally, $\pi_2(\gamma) \geq \varepsilon$. Therefore, (π_2, σ_2) satisfies the constraints of (12), σ_2 always demands $u_2(t) = 1$, π_2 puts probability 1 on player 1 conceding at any history h^{t+} at which $u_1(t) \neq \gamma(t)$, π_2 puts probability 0 on player 1 conceding at any history h^{t-}, and $\hat{t}(\gamma, \sigma_2, \pi_2) \geq \hat{t}(\gamma, \sigma_2', \pi_2')$. So the right-hand side of (12) continues to equal $\hat{T}$ when the additional constraint is imposed. ■

Step 2 of the proof builds on Step 1 to further simplify the constraint set of (12). For any belief π_2 satisfying the conditions of Step 1, let $\lambda^{\pi_2}(t)$ and $p^{\pi_2}(t)$ be the concession rate and probability of player 1 at history $(\gamma(\tau), 1)_{\tau \leq t}$ when her strategy is given by π_2; let S^{π_2} be the (countable) set of times s such that $p^{\pi_2}(s) > 0$; and let $\hat{t}(\pi_2) \equiv \hat{t}(\gamma, \sigma_2^0, \pi_2)$, where σ_2^0 is the strategy that always demands 1 and always rejects player 1's demand. Fixing a strategy $\sigma_2 \in \Sigma_2^*(\pi_2)$ that always demands 1 (which exists, by Step 1), note that (γ, σ_2) and (γ, σ_2^0) induce the same path of play until time $t(\gamma, \sigma_2)$ (at which point player 2 accepts under σ_2, but not under σ_2^0), and therefore $t(\gamma, \sigma_2) \geq \hat{t}(\gamma, \sigma_2, \pi_2)$ if and only if $t(\gamma, \sigma_2) \geq \hat{t}(\pi_2)$. Hence, there exists a strategy $\sigma_2 \in \Sigma_2^*(\pi_2)$ that always demands 1 and satisfies $t(\gamma, \sigma_2) \geq \hat{t}(\gamma, \sigma_2, \pi_2)$ if and only if it is optimal for player 2 to reject player 1's offer until time $\hat{t}(\pi_2)$ when he always demands 1, player 1 plays γ, and his initial belief is

π_2. I now use this observation to simplify the constraints of (12).

Step 2a: For any belief π_2 satisfying the conditions of Step 1, there exists a strategy $\sigma_2 \in \Sigma_2^*(\pi_2)$ that always demands 1 and satisfies $t(\gamma, \sigma_2) \geq \hat{t}(\gamma, \sigma_2, \pi_2)$ if and only if

$$\begin{aligned}
& 1 - \gamma(t) \\
\leq\ & \int_t^{\hat{t}(\pi_2)} \exp\left(-r(\tau - t) - \int_t^\tau \lambda^{\pi_2}(s)\, ds\right) \left(\prod_{s \in S^{\pi_2} \cap (t, \tau)} (1 - p^{\pi_2}(s))\right) \lambda^{\pi_2}(\tau)\, d\tau \\
& + \sum_{s \in S^{\pi_2} \cap \left(t, \hat{t}(\pi_2)\right)} \exp\left(-r(s - t) - \int_t^s \lambda^{\pi_2}(q)\, dq\right) \left(\prod_{q \in S^{\pi_2} \cap (t, s)} (1 - p^{\pi_2}(q))\right) p^{\pi_2}(s) \\
& + \exp\left(-r\left(\hat{t}(\pi_2) - t\right) - \int_t^{\hat{t}(\pi_2)} \lambda^{\pi_2}(s)\, ds\right) \left(\prod_{s \in S^{\pi_2} \cap \left(t, \hat{t}(\pi_2)\right)} (1 - p^{\pi_2}(s))\right) v\left(\hat{t}(\pi_2)\right)
\end{aligned}$$

for all $t < \hat{t}(\pi_2)$. (13)

Proof. By the above discussion, it suffices to show that (13) holds if and only if it is optimal for player 2 to reject player 1's offer until time $\hat{t}(\pi_2)$ when he always demands 1, player 1 plays γ, and his initial belief is π_2. The left-hand side of (13) is player 2's payoff from accepting player 1's demand at date $(t, 1)$ when $p^{\pi_2}(t) = 0$. The right-hand side of (13) is player 2's continuation payoff from rejecting player 1's demand until time $\hat{t}(\pi_2)$ when $p^{\pi_2}(t) = 0$. Thus, (13) must hold if it is optimal for player 2 to reject until time $\hat{t}(\pi_2)$, and (13) implies that it is optimal for player 2 to reject at times before $\hat{t}(\pi_2)$ where $p^{\pi_2}(t) = 0$. It remains to show that (13) implies that it is optimal for player 2 to reject at times before $\hat{t}(\pi_2)$ where $p^{\pi_2}(t) > 0$. Suppose that $p^{\pi_2}(t) > 0$. At date $(t, -1)$, the fact that S^{π_2} is countable and (13) holds at all times before t that are not in S^{π_2} implies that $\lim_{\tau \uparrow t}(1 - \gamma(\tau))$ is weakly less than player 2's continuation payoff from rejecting player 1's demand until time $\hat{t}(\pi_2)$. Furthermore, the fact that player 1 concedes with probability 0 at date $(t, -1)$ (as π_2 satisfies the conditions of Step 1) implies that $\lim_{\tau \uparrow t}(1 - \gamma(\tau))$ is indeed player 2's payoff from accepting at date $(t, -1)$. Thus, rejecting is optimal at date $(t, -1)$. At date $(t, 1)$, player 2's payoff from accepting is $(1 - p^{\pi_2}(t)/2)(1 - \gamma(t)) + (p^{\pi_2}(t)/2)(1)$, while his continuation payoff from rejecting until time $\hat{t}(\pi_2)$ is $1 - p^{\pi_2}(t)$ times the right-hand side of (13) plus $p^{\pi_2}(t)(1)$. Hence, (13) implies that rejecting is optimal at date $(t, 1)$ as well. So (13) implies that it is optimal for player 2 to reject at times before $\hat{t}(\pi_2)$ where $p^{\pi_2}(t) > 0$ (when he always demands 1, player 1 plays γ, and his initial belief is π_2). ■

Step 2b shows that one can replace $1-\gamma(t)$ with $v(t)$ in Step 2a.

Step 2b: For any belief π_2 satisfying the conditions of Step 1, there exists a strategy $\sigma_2 \in \Sigma_2^*(\pi_2)$ that always demands 1 and satisfies $t(\gamma,\sigma_2) \geq \hat{t}(\gamma,\sigma_2,\pi_2)$ if and only if

$$\begin{aligned} & v(t) \\ \leq & \int_t^{\hat{t}(\pi_2)} \exp\left(-r(\tau-t) - \int_t^{\tau} \lambda^{\pi_2}(s)\,ds\right)\left(\prod_{s\in S^{\pi_2}\cap(t,\tau)} (1-p^{\pi_2}(s))\right)\lambda^{\pi_2}(\tau)\,d\tau \\ & + \sum_{s\in S^{\pi_2}\cap\left(t,\hat{t}(\pi_2)\right)} \exp\left(-r(s-t) - \int_t^{s} \lambda^{\pi_2}(q)\,dq\right)\left(\prod_{q\in S^{\pi_2}\cap(t,s)} (1-p^{\pi_2}(q))\right)p^{\pi_2}(s) \\ & + \exp\left(-r\left(\hat{t}(\pi_2)-t\right) - \int_t^{\hat{t}(\pi_2)} \lambda^{\pi_2}(s)\,ds\right)\left(\prod_{s\in S^{\pi_2}\cap\left(t,\hat{t}(\pi_2)\right)} (1-p^{\pi_2}(s))\right)v\left(\hat{t}(\pi_2)\right) \end{aligned}$$

for all $t < \hat{t}(\pi_2)$. (14)

Proof. By Step 2a, it suffices to show that (13) is equivalent to (14). Note that (14) immediately implies (13) because $v(t) \geq 1-\gamma(t)$ for all t. For the converse, suppose that (13) holds, and that $v(t) > 1-\gamma(t)$ (as (13) and (14) are identical at t if $v(t) = 1-\gamma(t)$). Now $v(t) > 1-\gamma(t)$ implies that $v(t) = e^{-r(\tau-t)}\left(1-\underline{\gamma}(\tau)\right)$ for some $\tau > t$ such that $v(\tau,-1) = 1-\underline{\gamma}(\tau)$. Therefore, $v(\tau,-1)$ is weakly less than the limit as $s \uparrow \tau$ of the right-hand side of (14) evaluated at time s (with the convention that the right-hand side of (14) equals $v(s)$ if $s \geq \hat{t}(\pi_2)$). But the right-hand side of (14) at time t is at least $e^{-r(\tau-t)}$ times as large as this limit, which implies that the right-hand side of (14) at time t is at least $e^{-r(\tau-t)}v(\tau,-1) = v(t)$. Hence, (14) holds. ■

Let $\chi^{\pi_2}(t) \equiv \chi^{\left(\pi_2,\sigma_2^0\right)}$. By Step 1 and Step 2b (and the use of Bayes' rule to compute $\chi^{\pi_2}(t)$), (12) may be rewritten as

$$\hat{T} = \sup_{\substack{\pi_2:\pi_2(\gamma)\geq\varepsilon, \\ (14)\ holds}} \sup\left\{ t : \chi^{\pi_2}(t) = \frac{\exp\left(-\int_0^t \lambda^{\pi_2}(s)\,ds\right)\prod_{s\in S^{\pi_2}\cap[0,t)}(1-p^{\pi_2}(s)) - \varepsilon}{\exp\left(-\int_0^t \lambda^{\pi_2}(s)\,ds\right)\prod_{s\in S^{\pi_2}\cap[0,t)}(1-p^{\pi_2}(s))} > 0 \right\}. \quad (15)$$

The last step of the proof is solving this simplified program. Steps 3a and 3b show that there exists some belief π_2 that both attains the (outer) supremum in (15) (with the convention that the supremum is attained at π_2 if $\hat{t}(\pi_2) = \hat{T} = \infty$) and also maximizes $\lim_{t \uparrow \hat{T}} \chi^{\pi_2}(t)$ over all beliefs π_2 that attain the supremum (note that this limit exists for all π_2, because $\chi^{\pi_2}(t)$ is non-increasing). I then show that (14) must hold with equality (at all $t < \hat{T}$) under any such belief π_2, which implies that (15) may be solved under the additional constraint that (14) holds with equality. This final program is also solved in Step 3a, and has value $\tilde{T}$.

Step 3a: There exists a belief π_2 that attains the supremum in (15). In addition, the value of (15) under the additional constraint that (14) holds with equality equals $\tilde{T}$.

Proof. The plan of the proof is as follows. I first construct a belief π_2 such that $\pi_2(\gamma) \geq \varepsilon$ and (14) holds at all times $t < \hat{T}$ under π_2. I then show that if there exists a time $t < \hat{T}$ at which (14) holds with strict inequality under π_2, then there exists an alternative belief π_2' that attains the supremum in (15). Finally, I show that if (14) holds with equality at all times $t < \hat{T}$ under π_2, then π_2 itself attains the supremum in (15), which equals $\tilde{T}$.

Fix a sequence $\{\chi^{\pi_2^n}\}$ such that $\hat{t}(\pi_2^n) \uparrow \hat{T}$, $\pi_2^n(\gamma) \geq \varepsilon$ for all n, and (14) holds for all n. Note that $\chi^{\pi_2^n}(t)$ is non-increasing in t, for all n. Since the space of monotone functions from $\mathbb{R}_+$ to $[0,1]$ is sequentially compact (by Helly's selection theorem; see, e.g., Billingsley (1995) Theorem 25.9), there exists a subsequence $\{\chi^{\pi_2^m}\}$ that converges pointwise to some (non-increasing) function χ^{π_2}.[26] Furthermore, $\chi^{\pi_2}(0) \leq 1 - \varepsilon$, because $\chi^{\pi_2^m}(0) \leq 1 - \varepsilon$ for all m. Let $\pi_2 \in \Delta(\Sigma_1)$ be a belief such that $\pi_2(\gamma) \geq \varepsilon$ and player 1 demands $\gamma(t)$ for all t, concedes at rate $\lambda^{\pi_2}(t) = -\frac{\chi^{\pi_2\prime}(t)}{1-\chi^{\pi_2}(t)}$ if χ^{π_2} is differentiable at t and $\chi^{\pi_2}(t) > 0$ and concedes at rate $\lambda^{\pi_2}(t) = 0$ otherwise, and concedes with discrete probability $p^{\pi_2}(t) = \frac{\lim_{\tau \uparrow t} \chi^{\pi_2}(\tau) - \chi^{\pi_2}(t)}{(1-\chi^{\pi_2}(t)) \lim_{\tau \uparrow t} \chi^{\pi_2}(\tau)}$ if $\lim_{\tau \uparrow t} \chi(t) > 0$ and concedes with probability $p^{\pi_2}(t) = 0$ otherwise. Note that such a belief exists, because it can be easily verified that for any belief corresponding to concession rate

[26] Showing that the space of monotone functions from $\mathbb{R}_+ \to [0,1]$ is sequentially compact requires a slightly different version of Helly's selection theorem than that in Billingsley (1995), so here is a direct proof: If $\{f_n\}$ is a sequence of monotone functions $\mathbb{R}_+ \to [0,1]$, then there exists a subsequence $\{f_m\} \subseteq \{f_n\}$ that converges on $\mathbb{Q}_+$ to a monotone function $f : \mathbb{Q}_+ \to [0,1]$. Let $\tilde{f} : \mathbb{R}_+ \to [0,1]$ be given by $\tilde{f}(x) = \lim_{l \to \infty} f(x_l)$, where $\{x_l\}_{l=1}^{\infty} \uparrow x$ and $x_l \in \mathbb{Q}_+$ for all l. Then $\tilde{f}$ is monotone, which implies that there is a countable set $\mathcal{S}$ such that $\tilde{f}$ is continuous on $\mathbb{R}_+ \backslash \mathcal{S}$. Since $\mathcal{S}$ is countable, there exists a sub-subsequence $\{f_k\} \subseteq \{f_m\}$ such that $\{f_k\}$ converges on $\mathcal{S}$. Finally, let $\hat{f}(x) = \tilde{f}(x)$ if $x \in \mathbb{R}_+ \backslash \mathcal{S}$ and $\hat{f}(x) = \lim_{k \to \infty} f_k(x)$ if $x \in \mathcal{S}$. Then $\{f_k\} \to \hat{f}$.

and probability λ^{π_2} and p^{π_2},

$$\chi^{\pi_2}(t) = \frac{\exp\left(-\int_0^t \lambda^{\pi_2}(s)\,ds\right)\prod_{s\in S^{\pi_2}\cap[0,t)}(1-p^{\pi_2}(s)) - (1-\chi^{\pi_2}(0))}{\exp\left(-\int_0^t \lambda^{\pi_2}(s)\,ds\right)\prod_{s\in S^{\pi_2}\cap[0,t)}(1-p^{\pi_2}(s))} \text{ for all } t, \tag{16}$$

and therefore player 1 never concedes with probability at least $1-\chi^{\pi_2}(0) \geq \varepsilon$ under any belief with this concession rate and probability.

Observe that (14) holds at all times $t < \hat{T}$ under π_2. To see this, note that the fact that $\chi^{\pi_2^m}(t) \to \chi^{\pi_2}(t)$ for all t implies that

$$\exp\left(-\int_0^t \lambda^{\pi_2^m}(s)\,ds\right)\prod_{s\in S^{\pi_2}\cap[0,t)}\left(1-p^{\pi_2^m}(s)\right) \to \exp\left(-\int_0^t \lambda^{\pi_2}(s)\,ds\right)\prod_{s\in S^{\pi_2}\cap[0,t)}(1-p^{\pi_2}(s))$$

for all t. Since for all $t < \hat{T}$, there exists $M > 0$ such that (14) holds at time t under π_2^m for all $m > M$, this implies that (14) holds at all times $t < \hat{T}$ under π_2.

I now show that if there exists a time $t < \hat{T}$ at which (14) holds with strict inequality under belief π_2, then there exists an alternative belief π_2' that attains the supremum in (15). Suppose such a time t exists. I claim that there then exists a time $t_1 < \hat{t}(\pi_2)$ at which (14) holds with strict inequality and in addition either $\int_{t_1}^{t_1+\Delta}\lambda^{\pi_2}(s)\,ds > 0$ for all $\Delta > 0$ or $\sum_{s\in S^{\pi_2}\cap[t_1,t_1+\Delta)} p^{\pi_2}(s) > 0$ for all $\Delta > 0$. To see this, note that there must exist a time $t' \in \left(t, \hat{t}(\pi_2)\right)$ such that either $\int_{t'}^{t'+\Delta}\lambda^{\pi_2}(s)\,ds > 0$ for all $\Delta > 0$ or $p^{\pi_2}(t') > 0$ (because otherwise (14) could not hold with strict inequality at t). Let t_1 be the infimum of such times t', and note that either $\int_{t_1}^{t_1+\Delta}\lambda^{\pi_2}(s)\,ds > 0$ for all $\Delta > 0$ or $\sum_{s\in S^{\pi_2}\cap[t_1,t_1+\Delta)} p^{\pi_2}(s) > 0$ for all $\Delta > 0$. Then the fact that (14) holds with strict inequality at time t implies that (14) holds with strict inequality at time t_1, because otherwise the fact that $\int_t^{t_1}\lambda^{\pi_2}(s)\,ds = 0$ and $p^{\pi_2}(t'') = 0$ for all $t'' \in [t, t_1)$ would imply that (14) could not hold with strict inequality at time t. This proves the claim.

Thus, let $t_0 < \hat{t}(\pi_2)$ be such that (14) holds with strict inequality at time t_0 and in addition $\int_{t_0}^{t_0+\Delta}\lambda^{\pi_2}(s)\,ds > 0$ for all $\Delta > 0$ (the case where $\sum_{s\in S^{\pi_2}\cap[t_0,t_0+\Delta)} p^{\pi_2}(s) > 0$ is similar, and thus omitted). Since v is continuous but for downward jumps, there exist $\eta > 0$ and $\Delta > 0$ such that (14) holds with strict inequality at t for all $t \in [t_0, t_0+\Delta)$ when $\lambda^{\pi_2}(t)$ is replaced by $(1-\eta)\lambda^{\pi_2}(t)$ for all $t \in [t_0, t_0+\Delta)$. Define $\lambda^{\pi_2\prime}(t)$ by $\lambda^{\pi_2\prime}(t) \equiv \lambda^{\pi_2}(t)$ for all $t \notin [t_0, t_0+\Delta)$ and $\lambda^{\pi_2\prime}(t) \equiv (1-\eta)\lambda^{\pi_2}(t)$ for all $t \in [t_0, t_0+\Delta)$. Next, I claim that at time t_0 player 2's continuation payoff from rejecting γ until $\hat{t}(\pi_2)$ is strictly lower when player 1's concessions are given by $(\lambda^{\pi_2}(t), p^{\pi_2}(t))$ than when they are given by $(\lambda^{\pi_2\prime}(t), p^{\pi_2\prime\prime}(t))$, where $p^{\pi_2\prime\prime}(t)$ is defined by $p^{\pi_2\prime\prime}(t) \equiv p^{\pi_2}(t)$ for all $t \neq t_0$, and $p^{\pi_2\prime\prime}(t_0) \equiv 1 - \exp\left(-\eta \int_t^{t+\Delta} \lambda^{\pi_2}(s)\, ds\right)(1 - p^{\pi_2}(t_0)) > 0$. This follows because the total probability with which player 1 concedes in the interval $[t_0, t_0 + \Delta)$ is the same under $(\lambda^{\pi_2}(t), p^{\pi_2}(t))$ and under $(\lambda^{\pi_2\prime}(t), p^{\pi_2\prime\prime}(t))$, and some probability mass of concession is moved earlier to t_0 under $(\lambda^{\pi_2\prime}(t), p^{\pi_2\prime\prime}(t))$. Therefore, there exists $\zeta \in (0, p^{\pi_2\prime\prime}(t_0))$ such that at time t_0 player 2's continuation payoff from rejecting γ until $\hat{t}(\pi_2)$ is the same when player 1's concessions are given by $(\lambda^{\pi_2}(t), p^{\pi_2}(t))$ and when they are given by $(\lambda^{\pi_2\prime}(t), p^{\pi_2\prime}(t))$, where $p^{\pi_2\prime}(t)$ is defined by $p^{\pi_2\prime}(t) \equiv p^{\pi_2}(t)$ for all $t \neq t_0$, and $p^{\pi_2\prime}(t_0) \equiv p^{\pi_2\prime\prime}(t_0) - \zeta < p^{\pi_2\prime\prime}(t_0)$. The fact that (14) holds at all $t < \hat{T}$ when player 1's concessions are given by $(\lambda^{\pi_2}(t), p^{\pi_2}(t))$ now implies that (14) holds at all $t < \hat{T}$ when player 1's concessions are given by $(\lambda^{\pi_2\prime}(t), p^{\pi_2\prime}(t))$. Furthermore, $\exp\left(-\int_0^{\hat{T}} \lambda^{\pi_2\prime}(t)\, dt\right) \prod_{s \in S^{\pi_2} \cap [0, \hat{T})} (1 - p^{\pi_2\prime}(s)) > \exp\left(-\int_0^{\hat{T}} \lambda^{\pi_2}(t)\, dt\right) \prod_{s \in S^{\pi_2} \cap [0, \hat{T})} (1 - p^{\pi_2}(s)) \geq \varepsilon$. Therefore, $\sup\left\{t : \chi^{\pi_2'}(t) > 0\right\} \geq \hat{T}$, so by the definition of $\hat{T}$ it must be that $\sup\left\{t : \chi^{\pi_2'}(t) > 0\right\} = \hat{T}$. Thus, π_2' attains the supremum in (15).

Finally, suppose that (14) holds with equality at all $t < \hat{T}$ under belief π_2 (defined above). Then (14) holds with equality at all $t < \hat{t}(\pi_2)$ under belief π_2, because $\hat{t}(\pi_2) \leq \hat{T}$ (by definition of $\hat{T}$). Let $t < \hat{t}(\pi_2)$ be a time at which v is differentiable. Then the derivative of the right-hand side of (14) at t must exist and equal $v'(t)$. This implies that $p^{\pi_2}(t) = 0$, and, by Leibniz's rule, the derivative of the right-hand side of (14) equals $-\lambda^{\pi_2}(t) + (r + \lambda^{\pi_2}(t))\, v(t)$. Hence,

$$\lambda^{\pi_2}(t) = \frac{rv(t) - v'(t)}{1 - v(t)}.$$ [27]

Since v is differentiable almost everywhere, this implies that

$$\int_0^{\tau} \lambda^{\pi_2}(s)\, ds = \int_0^{\tau} \lambda(s)\, ds \text{ for all } \tau < \hat{t}(\pi_2), \tag{17}$$

[27] Note that $v(t) = 1$ is impossible. For $v(\tau) \in [0, 1]$ implies that $v'(t) = 0$ if $v'(t)$ exists and $v(t) = 1$. But then the equation $v'(t) = -\lambda^{\pi_2}(t) + (r + \lambda^{\pi_2}(t))\, v(t)$ would reduce to $0 = r$, violating the assumption that $r > 0$.

where λ is defined by (3). Similarly, if (14) holds with equality then the difference between the limit as $s \uparrow t$ of the right-hand side of (14) evaluated at s and the limit as $s \downarrow t$ of the right-hand side of (14) evaluated at s must equal $v(t,-1) - v(t)$, for all $t < \hat{t}(\pi_2)$. By inspection, this difference equals $p^{\pi_2}(t) - p^{\pi_2}(t) v(t)$. Hence,

$$p^{\pi_2}(t) = \frac{v(t,-1) - v(t)}{1 - v(t)}$$

for all $t < \hat{t}(\pi_2)$ at which v is discontinuous, and $p^{\pi_2}(t) = 0$ otherwise.[28] Therefore,

$$\prod_{s \in S^{\pi_2} \cap [0,\tau)} (1 - p^{\pi_2}(s)) = \prod_{s \in S \cap [0,\tau)} (1 - p(s)) \tag{18}$$

for all $\tau < \hat{t}(\pi_2)$, where S is the set of discontinuity points of v, and p is defined by (4). Combining (17) and (18), I conclude that if (14) holds with equality under belief π_2, then

$$\hat{t}(\pi_2) = \sup \left\{ t : \exp\left(-\int_0^t \lambda(s)\, ds \right) \prod_{s \in S \cap [0,t)} (1 - p(s)) > \varepsilon \right\},$$

which equals $\tilde{T}$. In addition, $\chi^{\pi_2}(t) \leq 0$ for all $t > \tilde{T}$ (by (16), recalling that $1 - \chi^{\pi_2}(0) \geq \varepsilon$), so $\hat{T} = \tilde{T}$ and the supremum in (15) is attained by π_2. ■

Step 3b: There exists a belief that both attains the supremum in (15) and maximizes $\lim_{t \uparrow \hat{T}} \chi^{\pi_2}(t)$ *over all beliefs* π_2 *that attain the supremum in (15).*

Proof. Let $\chi \in [0,1]$ be the supremum of $\lim_{t \uparrow \hat{T}} \chi^{\pi_2}(t)$ over all beliefs π_2 that attain the supremum in (15). If $\chi = 0$, then any belief π_2 that attains the supremum in (15) also satisfies $\lim_{t \uparrow \hat{T}} \chi^{\pi_2}(t) = \chi$. Thus, suppose that $\chi > 0$. Let $\{\pi_2^n\}$ be a sequence of beliefs that all attain the supremum in (15) such that $\lim_{t \uparrow \hat{T}} \chi^{\pi_2^n}(t) \uparrow \chi$. The sequential compactness argument in Step 3a implies that there exists a subsequence $\{\pi_2^m\} \subseteq \{\pi_2^n\}$ and a belief π_2 satisfying the constraints of (15) such that $\chi^{\pi_2^m}(t) \to \chi^{\pi_2}(t)$ for all t. Furthermore, $\chi^{\pi_2}(t)$ is non-increasing, so $\lim_{t \uparrow \hat{T}} \chi^{\pi_2}(t)$ exists. Because π_2 satisfies the constraints of (15), $\lim_{t \uparrow \hat{T}} \chi^{\pi_2}(t) \leq \chi$. Now suppose, toward a contradiction, that $\lim_{t \uparrow \hat{T}} \chi^{\pi_2}(t) < \chi$. Then there exists $\eta > 0$ and $t' \leq \hat{T}$ such that $\chi^{\pi_2}(t') < \chi - \eta$. Since $\lim_{m \to \infty} \lim_{t \uparrow \hat{T}} \chi^{\pi_2^m}(t) = \chi$, there exists $M > 0$ such that, for all $m > M$, $\lim_{t \uparrow \hat{T}} \chi^{\pi^m}(t) > \chi - \eta$. And $\chi^{\pi_2^m}(t)$ is non-increasing for all m, so this implies that $\chi^{\pi_2^m}(t') > \chi - \eta$ for all $m > M$. Now $\chi^{\pi_2^m}(t') \to \chi^{\pi_2}(t')$

[28] The fact that v cannot jump up rules out $v(t) = 1$.

implies that $\chi^{\pi_2}(t') \geq \chi - \eta$, a contradiction. Therefore, $\lim_{t \uparrow \hat{T}} \chi^{\pi_2}(t) = \chi$. Finally, the fact that $\lim_{t \uparrow \hat{T}} \chi^{\pi_2}(t) > 0$ implies that π_2 attains the supremum in (15). ■

I now complete the proof of Lemma 1. If (14) holds with strict inequality at some time $t < \hat{T}$ under a belief π_2 such that $\hat{t}(\pi_2) = \hat{T}$, then the procedure for modifying π_2 described in the fifth paragraph of the proof of Step 3a yields a belief π_2' such that $\hat{t}(\pi_2') = \hat{T}$ and $\lim_{t \uparrow \hat{T}} \chi^{\pi_2'}(t) > \lim_{t \uparrow \hat{T}} \chi^{\pi_2}(t)$. This implies that the only beliefs π_2 that both attain the supremum in (15) and maximize $\lim_{t \uparrow \hat{T}} \chi^{\pi_2}(t)$ (over all beliefs that attain the supremum in (15)) satisfy the additional constraint that (14) holds with equality. Since such a belief exists by Step 3b, the value of (15) equals the value of (15) under this additional constraint, which equals $\tilde{T}$ by Step 3a. ■

Proof of Theorem 1. Let γ_n and γ^* be defined as in Section 3.4. To show that $\lim_{n \to \infty} u_1^*(\gamma_n) = 1/(1 - \log \varepsilon)$, it remains only to show that $T_n^1 > \tilde{T}(\gamma_n)$ for all $n \in \mathbb{N}$. To see this, note that $T_n^1 = \frac{1}{r} \log\left(\frac{n+1}{n}(1 - \log \varepsilon)\right)$. Since $\gamma_n(t) = \left(\frac{n}{n+1}\right) \frac{e^{rt}}{1 - \log \varepsilon}$ for all $t \leq T_n^1$ and γ_n is non-decreasing, it follows that $v(t) = 1 - \left(\frac{n}{n+1}\right) \frac{e^{rt}}{1 - \log \varepsilon}$ for all $t \leq T_n^1$. Therefore,

$$
\begin{aligned}
& \exp\left(-\int_0^{T_n^1} \frac{rv(t) - v'(t)}{1 - v(t)} dt\right) \prod_{s \in S \cap [0, T_n^1]} \left(\frac{1 - v(s, -1)}{1 - v(s)}\right) \\
= \; & \exp\left(-\int_0^{T_n^1} r\left(\frac{n+1}{n}\right)(1 - \log \varepsilon) e^{-rt} dt\right) \\
= \; & \exp\left(-\frac{1}{n}(1 - \log \varepsilon)\right) \varepsilon \\
< \; & \varepsilon.
\end{aligned}
$$

Hence, by the definition of $\tilde{T}(\gamma_n)$, $T_n^1 \geq \tilde{T}(\gamma_n)$. Furthermore, the fact that $\exp\left(-\int_0^{\tau} \frac{rv(t) - v'(t)}{1 - v(t)} dt\right)$ is strictly decreasing in τ for all $\tau \in [0, T_n^1]$ implies that $T_n^1 > \tilde{T}(\gamma_n)$.

I now complete the proof of Theorem 1 by showing that γ^* is the unique maxmin posture.[29] That is, I show that if $\{\gamma_n\}$ is any sequence of postures converging pointwise to some posture γ satisfying $u_1^*(\gamma_n) \to u_1 \geq 1/(1 - \log \varepsilon)$, then $\gamma = \gamma^*$. There are two steps. First, letting $\{v_n\}$ be the continuation value functions corresponding to the $\{\gamma_n\}$, and letting v^* be the continuation value function corresponding to γ^* (that is,

[29] Technically, I must also show that $u_1^*(\gamma^*) \leq 1/(1 - \log \varepsilon)$. In fact, $u_1^*(\gamma^*) = 0$, by Lemma 3 and the observation that $T(\gamma^*) = \infty$ (which follows because $\gamma^*(t) = 1$ for all $t \geq \tilde{T}(\gamma^*)$).

$v^*(t) = \max\{1 - e^{rt}/(1-\log\varepsilon), 0\}$), I show that $\sup_{t\in\mathbb{R}_+} e^{-rt}|v^*(t) - v_n(t)| \to 0$. This step is dividing into showing first that $\sup_{t\le\tilde{T}(\gamma)}|v^*(t) - v_n(t)| \to 0$ (Step 1a) and then that $\sup_{t\in\mathbb{R}_+} e^{-rt}|v^*(t) - v_n(t)| \to 0$ (Step 1b). Second, I show that this implies that $\gamma = \gamma^*$.

Step 1a: For all $\delta > 0$, there exists $\zeta > 0$ such that if $u_1^*(\gamma) \ge 1/(1-\log\varepsilon) - \zeta$, then $\sup_{t\le\tilde{T}(\gamma)}|v^*(t) - v(t)| \le \delta$ (where v is the continuation value function corresponding to γ).

Proof. The plan is to first note that any posture γ that guarantees close to $1/(1-\log\varepsilon)$ must demand close to $e^{rt}(1/(1-\log\varepsilon))$ (or more) for all $t \le T(\gamma)$ and must also have $\tilde{T}(\gamma)$ close to T^1 (the time at which $\gamma^*(t)$ reaches 1), and then to show that any posture with these two properties must correspond to a continuation value function that is close to v^* until time $\tilde{T}(\gamma)$.

Formally, suppose that $u_1^*(\gamma) \ge 1/(1-\log\varepsilon) - \zeta$ for some posture γ and some $\zeta \in (0, 1/(1-\log\varepsilon))$. Let $T^1 \equiv (1/r)\log(1-\log\varepsilon)$. Then it must be that $\tilde{T}(\gamma) \le T^1 - (1/r)\log(1-\zeta(1-\log\varepsilon))$, for otherwise it would follow from $T(\gamma) \ge \tilde{T}(\gamma)$ that

$$\begin{aligned} u_1^*(\gamma) &= \min_{t\le T(\gamma)} e^{-rt}\underline{\gamma}(t) \le e^{-r\tilde{T}(\gamma)}\underline{\gamma}\left(\tilde{T}(\gamma)\right) \\ &< \exp\left(-rT^1 + \log(1-\zeta(1-\log\varepsilon))\right)(1) = \frac{1}{1-\log\varepsilon} - \zeta. \end{aligned}$$

Furthermore, if $u_1^*(\gamma) \ge 1/(1-\log\varepsilon) - \zeta$, it must also be that $\gamma(t) \ge e^{rt}(1/(1-\log\varepsilon) - \zeta)$ for all $t \le T(\gamma)$, for otherwise $\min_{t\le T(\gamma)} e^{-rt}\underline{\gamma}(t)$ would be strictly less than $1/(1-\log\varepsilon) - \zeta$. I will show that, for all $\delta > 0$, there exists $\zeta \in (0, 1/(1-\log\varepsilon))$ such that if $\gamma(t) \ge e^{rt}(1/(1-\log\varepsilon) - \zeta)$ for all $t \le T(\gamma)$ and $\sup_{t\le\tilde{T}(\gamma)}|v^*(t) - v(t)| > \delta$, then $\tilde{T}(\gamma) > T^1 - (1/r)\log(1-\zeta(1-\log\varepsilon))$. This completes the proof of Step 1a.

Fix $\delta > 0$ and $\zeta \in (0, 1/(1-\log\varepsilon))$, and suppose that $\gamma(t) \ge e^{rt}(1/(1-\log\varepsilon) - \zeta)$ for all $t \le T(\gamma)$ and $\sup_{t\le\tilde{T}(\gamma)}|v^*(t) - v(t)| > \delta$. A straightforward implication is that $v(t) \le 1 - e^{rt}(1/(1-\log\varepsilon) - \zeta)$ for all $t \le T(\gamma)$. Also, if $\tilde{T}(\gamma)$ is finite then

$$\exp\left(-\int_0^{\tilde{T}(\gamma)} \frac{rv(t) - v'(t)}{1-v(t)}dt\right) \prod_{s\in S\cap[0,\tilde{T}(\gamma)]}\left(\frac{1-v(s,-1)}{1-v(s)}\right) \le \varepsilon.$$

Thus, if $\tilde{T}(\gamma)$ is finite then it must be that

$$\inf_{\substack{v:\mathbb{R}_+\to\mathbb{R}_+: \\ v(t)\le 1-e^{rt}\left(\frac{1}{1-\log\varepsilon}-\zeta\right), \\ \sup_{t\le\tilde{T}(\gamma)}|v^*(t)-v(t)|>\delta}} \exp\left(-\int_0^{\tilde{T}(\gamma)} \frac{rv(t) - v'(t)}{1-v(t)}dt\right) \prod_{s\in S\cap[0,\tilde{T}(\gamma)]}\left(\frac{1-v(s,-1)}{1-v(s)}\right) \le \varepsilon, \quad (19)$$

where $\tilde{T}(\gamma)$ is viewed as a parameter and the infimum is taken over all functions $v : \mathbb{R}_+ \to \mathbb{R}_+$ satisfying the constraints. I will show that if $\zeta > 0$ is sufficiently small then (19) can hold only if $\tilde{T}(\gamma) > T^1 - (1/r)\log(1 - \zeta(1 - \log\varepsilon))$.

I first show that any attainable value of the program on the left-hand side of (19) can be arbitrarily closely approximated by the value attained by a *continuous* function $v(t)$ satisfying the constraints of (19); hence, in calculating the infimum over such values, attention may be restricted to continuous functions. To see this, fix $\eta \in (0,1)$ and let

$$S^\eta \equiv \bigcup_{s \in S \cap \left[0, \tilde{T}(\gamma)\right]} [s - \eta, s].$$

Define the function $v^\eta(t)$ by $v^\eta(t) \equiv v(t)$ for all $t \notin S^\eta$, and

$$v^\eta(t) \equiv \left(1 - \frac{t-(s-\eta)}{\eta}\right) v(s-\eta) + \frac{t-(s-\eta)}{\eta} v(s) \text{ for all } t \in S^\eta.$$

Observe that v^η is continuous, and that v^η satisfies the constraints of (19) if η is sufficiently small.[30] Furthermore, for all $s \in S$,

$$\exp\left(\int_{s-\eta}^{s} \frac{v^{\eta\prime}(t)}{1 - v^\eta(t)} dt\right) = \frac{1 - v^\eta(s-\eta)}{1 - v^\eta(s)} = \frac{1 - v(s-\eta)}{1 - v(s)}.$$

Also, since $v^\eta(t) \le 1 - (1/(1-\log\varepsilon) - \zeta) < 1$ for all $t \in \left[0, \tilde{T}(\gamma)\right]$, and the measure of S^η goes to 0 as $\eta \to 0$,

$$\lim_{\eta \to 0} \exp\left(-\int_{S^\eta} \frac{r v^\eta(t)}{1 - v^\eta(t)} dt\right) = 1.$$

Therefore,

$$\begin{aligned}
& \lim_{\eta \to 0} \exp\left(-\int_0^{\tilde{T}(\gamma)} \frac{r v^\eta(t) - v^{\eta\prime}(t)}{1 - v^\eta(t)} dt\right) \\
= & \lim_{\eta \to 0} \exp\left(-\int_{\left[0, \tilde{T}(\gamma)\right] \setminus S^\eta} \frac{r v^\eta(t) - v^{\eta\prime}(t)}{1 - v^\eta(t)} dt\right) \exp\left(-\int_{S^\eta} \frac{r v^\eta(t)}{1 - v^\eta(t)} dt\right) \prod_{s \in S \cap \left[0, \tilde{T}(\gamma)\right]} \left(\frac{1 - v(s-\eta)}{1 - v(s)}\right) \\
= & \exp\left(-\int_0^{\tilde{T}(\gamma)} \frac{r v(t) - v'(t)}{1 - v(t)} dt\right) \prod_{s \in S \cap \left[0, \tilde{T}(\gamma)\right]} \left(\frac{1 - v(s,-1)}{1 - v(s)}\right).
\end{aligned}$$

[30] Proof: The constraint $v(t) \le 1 - e^{rt}\left(\frac{1}{1-\log\varepsilon} - \zeta\right)$ specifies that the graph of $v(t)$ lies in a convex subset of $\mathbb{R}^2_+$. Since the graph of $v^\eta(t)$ lies in the convex hull of the graph of $v(t)$, the graph of $v^\eta(t)$ lies in any convex subset of $\mathbb{R}^2_+$ that contains the graph of $v(t)$. Hence, $v^\eta(t) \le 1 - e^{rt}\left(\frac{1}{1-\log\varepsilon} - \zeta\right)$ (for all $\eta > 0$). In addition, the fact that $\sup_{t \le \tilde{T}(\gamma)} |v^*(t) - v(t)| > \delta$ and $v^\eta(t) \to v(t)$ for all $t \in \mathbb{R}_+$ implies that $\sup_{t \le \tilde{T}(\gamma)} |v^*(t) - v^\eta(t)| > \delta$ for sufficiently small η.

Thus, the value of the program in (19) attained by any function v is arbitrarily closely approximately by the value attained by the continuous function v^η as $\eta \to 0$, and for sufficiently small $\eta > 0$ this function also satisfies the constraints of (19).

I now derive a lower bound on the left-hand side (19) under the additional constraint that v is continuous. Using the fact that $v(s,-1) = v(s)$ for all s when v is continuous and integrating the $v'(t)/(1-v(t))$ term, this constrained program may be rewritten as

$$\inf_{\substack{v:\mathbb{R}_+\to\mathbb{R}_+ \text{ continuous:}\\ v(t)\leq 1-e^{rt}\left(\frac{1}{1-\log\varepsilon}-\zeta\right),\\ \sup_{t\leq\tilde{T}(\gamma)}|v^*(t)-v(t)|>\delta}} \exp\left(-\int_0^{\tilde{T}(\gamma)} \frac{rv(t)}{1-v(t)}dt\right)\left(\frac{1-v(0)}{1-v\left(\tilde{T}(\gamma)\right)}\right).$$

Since $v(t) \geq 0$ for all t, the value of this program is bounded from below by the value of the program:

$$\inf_{\substack{v:\mathbb{R}_+\to\mathbb{R}_+ \text{ continuous:}\\ v(t)\leq 1-e^{rt}\left(\frac{1}{1-\log\varepsilon}-\zeta\right),\\ \sup_{t\leq\tilde{T}(\gamma)}|v^*(t)-v(t)|>\delta}} \exp\left(-\int_0^{\tilde{T}(\gamma)} \frac{rv(t)}{1-v(t)}dt\right)(1-v(0)). \tag{20}$$

Note that the value of the program (20) is continuous and non-increasing in $\tilde{T}(\gamma)$.

Let $\tilde{T}_\zeta \equiv 0$ if the value of the program (20) is less than ε when $\tilde{T}(\gamma) = 0$, let $\tilde{T}_\zeta \equiv \infty$ if the value of the program is greater than ε for all $\tilde{T}(\gamma) \in \mathbb{R}_+$, and otherwise let $\tilde{T}_\zeta$ be the value of the parameter $\tilde{T}(\gamma)$ such that (20) equals ε (which exists by the Intermediate Value Theorem). I will show that $\tilde{T}_\zeta > T^1 - (1/r)\log(1-\zeta(1-\log\varepsilon))$ for sufficiently small $\zeta \in (0, 1/(1-\log\varepsilon))$.

I first show that this inequality holds for $\zeta = 0$, that is, that $\tilde{T}_0 > T^1$. To see this, note that (20) decreases whenever the value of $v(t)$ is increased on a subset of $\left[0, \tilde{T}(\gamma)\right]$ of positive measure. Hence, the unique solution to the program (20) without the constraint $\sup_{t\leq\tilde{T}(\gamma)}|v^*(t)-v(t)| > \delta$ is $v(t) = 1 - e^{rt}/(1-\log\varepsilon) = v^*(t)$ for all $t \leq \tilde{T}(\gamma)$ (when $\zeta = 0$). Using this observation, it is straightforward to check that the value of $\tilde{T}(\gamma)$ such that the value of the program (20) without the constraint $\sup_{t\leq\tilde{T}(\gamma)}|v^*(t)-v(t)| > \delta$ equals ε is T^1. Therefore, the constraint $\sup_{t\leq\tilde{T}(\gamma)}|v^*(t)-v(t)| > \delta$ binds in (20), and $\tilde{T}_0 > T^1$.

Next, $\tilde{T}_\zeta$ is continuous in ζ, by the Maximum Theorem (which implies that the value of (20) is continuous in ζ, for fixed $\tilde{T}(\gamma)$) and the Implicit Function Theorem. And $T^1 - (1/r)\log(1-\zeta(1-\log\varepsilon))$ is continuous in ζ as well. Hence, the fact that $\tilde{T}_0 > T^1$ implies that $\tilde{T}_\zeta > T^1 - (1/r)\log(1-\zeta(1-\log\varepsilon))$ for some $\zeta \in (0, 1/(1-\log\varepsilon))$.

By (19), if $\gamma(t) \geq e^{rt}(1/(1-\log\varepsilon)-\zeta)$ for all $t \leq T(\gamma)$ and $\sup_{t\leq\tilde{T}(\gamma)}|v^*(t)-v(t)| > \delta$, then $\tilde{T}_\zeta$ is a lower bound on $\tilde{T}(\gamma)$. Thus, the fact that $\tilde{T}_\zeta > T^1-(1/r)\log(1-\zeta(1-\log\varepsilon))$ for some $\zeta \in (0, 1/(1-\log\varepsilon))$ completes the proof. ■

Step 1b: For all $\delta > 0$, there exists $\zeta > 0$ such that if $u_1^*(\gamma) \geq 1/(1-\log\varepsilon)-\zeta$, then $\sup_{t\in\mathbb{R}_+} e^{-rt}|v^*(t)-v(t)| \leq \delta$.

Proof. Step 1a implies that, for any $\delta > 0$ and $K > 1$, there exists $\zeta(K) \in (0, 1/(1-\log\varepsilon))$ such that if $u_1^*(\gamma) \geq 1/(1-\log\varepsilon)-\zeta(K)$, then $\sup_{t\leq\tilde{T}(\gamma)} e^{-rt}|v^*(t)-v(t)| \leq \delta/K$. I now argue that, for K sufficiently large, there exists $\zeta' \in (0, \zeta(K))$ such that if $u_1^*(\gamma) \geq 1/(1-\log\varepsilon)-\zeta'$, then in addition $\sup_{t>\tilde{T}(\gamma)} e^{-rt}|v^*(t)-v(t)| \leq \delta$. To see this, note that as $K \to \infty$, $\tilde{T}(\gamma) \to T^1$ uniformly over all postures γ such that $\sup_{t\leq\tilde{T}(\gamma)} e^{-rt}|v^*(t)-v(t)| \leq \delta/K$. Choose $K^* > 1$ such that $\left|e^{-r\tilde{T}(\gamma)}-e^{-rT^1}\right| < \delta/2$ and $v^*\left(\tilde{T}(\gamma)\right) \leq e^{r\tilde{T}(\gamma)}\delta$ for any such posture γ, and suppose that a posture γ is such that $\sup_{t\leq\tilde{T}(\gamma)} e^{-rt}|v^*(t)-v(t)| \leq \delta/K^*$ but $e^{-rt_0}|v^*(t_0)-v(t_0)| > \delta$ for some $t_0 > \tilde{T}(\gamma)$. Now $v^*(t_0) \leq v^*\left(\tilde{T}(\gamma)\right) \leq e^{r\tilde{T}(\gamma)}\delta \leq e^{rt_0}\delta$, so it follows that $e^{-rt_0}v(t_0) > \delta + e^{-rt_0}v^*(t_0)$. Therefore,

$$\max_{t\geq\tilde{T}(\gamma)} e^{-rt}\left(1-\underline{\gamma}(t)\right) \geq e^{-rt_0}v(t_0) \geq \delta.$$

By the definition of $T(\gamma)$, this implies that there exists $t_1 \in \left[\tilde{T}(\gamma), T(\gamma)\right]$ such that $e^{-rt_1}\left(1-\underline{\gamma}(t_1)\right) \geq \delta$, or equivalently $\underline{\gamma}(t_1) \leq 1-e^{rt_1}\delta$. Hence,

$$\begin{aligned} u_1^*(\gamma) &= \min_{t\leq T(\gamma)} e^{-rt}\underline{\gamma}(t) \leq e^{-rt_1}\left(1-e^{rt_1}\delta\right) \leq e^{-r\tilde{T}(\gamma)}\left(1-e^{r\tilde{T}(\gamma)}\delta\right) \\ &= e^{-r\tilde{T}(\gamma)}-\delta < e^{-rT^1}-\delta/2 = 1/(1-\log\varepsilon)-\delta/2. \end{aligned}$$

Therefore, taking $\zeta' \equiv \min\{\zeta(K^*), \delta/2\}$, it follows that if $u_1^*(\gamma) \geq 1/(1-\log\varepsilon)-\zeta'$ then $\sup_{t\leq\tilde{T}(\gamma)} e^{-rt}|v^*(t)-v(t)| \leq \delta/K^*$ and $\sup_{t>\tilde{T}(\gamma)} e^{-rt}|v^*(t)-v(t)| \leq \delta$, and hence $\sup_{t\in\mathbb{R}_+} e^{-rt}|v^*(t)-v(t)| \leq \delta$. ■

Step 2: If $\gamma_n(t) \to \gamma(t)$ for all $t \in R_+$ for some posture γ, and $\sup_{t\in\mathbb{R}_+} e^{-rt}|v^*(t)-v_n(t)| \to 0$, then $\gamma = \gamma^*$.

Proof. First, note that if $\gamma(t) < \gamma^*(t)$ for some $t \in \mathbb{R}_+$, then there exist $N > 0$ and $\eta > 0$ such that $\gamma_n(t) < \gamma^*(t)-\eta$ for all $n > N$. Since $v_n(t) \geq 1-\gamma_n(t)$, this implies that $v_n(t) \geq 1-\gamma^*(t)+\eta = v^*(t)+\eta$ for all $n > N$, a contradiction.

It is more difficult to rule out the possibility that $\gamma(t) > \gamma^*(t)$ for some $t \in \mathbb{R}_+$. Suppose that this is so. Since γ and γ^* are right-continuous, there exist $\eta > 0$ and a non-degenerate closed interval $I_0 \subseteq \mathbb{R}_+$ such that $\gamma(t) > \gamma^*(t) + \eta$ for all $t \in I_0$. If it were the case that $\gamma_n(t) \geq \gamma^*(t) + \eta/2$ for all $t \in I_0$ and n sufficiently large, then the condition $\sup_{t\in\mathbb{R}_+} e^{-rt} |v^*(t) - v_n(t)| \to 0$ would fail, so this is not possible.[31] Hence, there exists $t_1 \in I_0$ and $n_1 \geq 0$ such that $\gamma_{n_1}(t_1) < \gamma^*(t_1) + \eta/2$. Since γ_{n_1} and γ^* are right-continuous, there exists a non-degenerate closed interval $I_1 \subseteq I_0$ such that $\gamma_{n_1}(t) < \gamma^*(t) + \eta/2$ for all $t \in I_1$. Next, it cannot be the case that $\gamma_n(t) \geq \gamma^*(t) + \eta/2$ for all $t \in I_1$ and $n > n_1$ (by the same argument as above), so there exists $t_2 \in I_1$ and $n_2 > n_1$ such that $\gamma_{n_2}(t_2) < \gamma^*(t_2) + \eta/2$. As above, this implies that there exists a non-degenerate closed $I_2 \subseteq I_1$ such that $\gamma_{n_2}(t) < \gamma^*(t) + \eta/2$ for all $t \in I_2$. Proceeding in this manner yields an infinite sequence of non-degenerate closed intervals $\{I_m\}$ and integers $\{n_m\}$ such that $I_{m+1} \subseteq I_m$, $n_{m+1} > n_m$, and $\gamma_{n_m}(t) < \gamma^*(t) + \eta/2$ for all $t \in I_m$ and $m \in \mathbb{N}$. Let $I \equiv \cap_{m\in\mathbb{N}} I_m$, a non-empty set (possibly a single point), and fix $t \in I$. Then $\gamma_{n_m}(t) < \gamma^*(t) + \eta/2$ for all $m \in \mathbb{N}$, and since $n_{m+1} > n_m$ for all $m \in \mathbb{N}$ this contradicts the assumption that $\gamma_n(t) \to \gamma(t)$. ■ ■

Proof of Proposition 2. Lemmas 1 through 3 apply to any posture, whether or not it is constant. In addition, if γ is constant then $T(\gamma) = \tilde{T}(\gamma)$. Thus, Lemma 3 implies that $u_1^*(\gamma) = \min_{t \leq T(\gamma)} e^{-rt}\gamma = e^{-r\tilde{T}(\gamma)}\gamma$. Furthermore, $\lambda(t) = r(1-\gamma)/\gamma$ and $p(t) = 0$ for all t, so it follows by the definition of $\tilde{T}(\gamma)$ that $\exp\left(-r\left(\frac{1-\gamma}{\gamma}\right)\tilde{T}(\gamma)\right) = \varepsilon$. Hence, $\tilde{T}(\gamma) = -\frac{1}{r}\left(\frac{\gamma}{1-\gamma}\right)\log\varepsilon$ if $\gamma < 1$, and $\tilde{T}(\gamma) = \infty$ if $\gamma = 1$. Therefore,

$$\bar{u}_1^* = \max_{\gamma\in[0,1]} e^{-r\tilde{T}(\gamma)}\gamma = \max_{\gamma\in[0,1)} \exp\left(\frac{\gamma}{1-\gamma}\log\varepsilon\right)\gamma. \tag{21}$$

[31] Proof: Suppose that $\sup_{t\in\mathbb{R}_+} e^{-rt}|v^*(t) - v_n(t)| \to 0$. Fix $N > 0$, suppose that $\gamma_n(t) \geq \gamma^*(t) + \eta/2$ for all $t \in I_0$ and $n > N$, and denote the length of I_0 by 2Δ and the midpoint of I_0 by t_0. Then player 2 cannot receive a payoff above $e^{-rt_0}(1 - \gamma^*(t_0) - \eta/2)$ from accepting at any time $t \in [t_0, t_0 + \Delta]$ when facing posture γ_n for any $n > N$. Hence, $v_n(t_0) \leq \max\{1 - \gamma^*(t_0) - \eta/2, e^{-r\Delta}v_n(t_0 + \Delta)\}$ for all $n > N$. In addition, noting that $\gamma^*(t_0) < 1 - \eta/2$, there exists $N' > 0$ such that $v_n(t_0 + \Delta) < v_n^*(t_0 + \Delta) + (e^{r\Delta} - 1)(1 - \gamma^*(t_0) - \eta/2) = 1 - \gamma^*(t_0 + \Delta) + (e^{r\Delta} - 1)(1 - \gamma^*(t_0) - \eta/2)$ for all $n > N'$. Therefore, $v_n(t_0) \leq \max\{1 - \gamma^*(t_0) - \eta/2, e^{-r\Delta}(1 - \gamma^*(t_0 + \Delta)) + (1 - e^{-r\Delta})(1 - \gamma^*(t_0) - \eta/2)\} \leq \max\{1 - \gamma^*(t_0) - \eta/2, 1 - \gamma^*(t_0) - (1 - e^{-r\Delta})\eta/2\} = v^*(t_0) - (1 - e^{-r\Delta})\eta/2$ for all $n > \max\{N, N'\}$, which contradicts the hypothesis that $\sup_{t\in\mathbb{R}_+} e^{-rt}|v^*(t) - v_n(t)| \to 0$.

Note that (21) is concave in γ. The first-order condition is

$$1 = -\frac{\bar{\gamma}_\varepsilon^*}{\left(1-\bar{\gamma}_\varepsilon^*\right)^2}\log\varepsilon, \tag{22}$$

which has a solution if $\varepsilon < 1$. Solving this quadratic equation yields the formula for $\bar{\gamma}_\varepsilon^*$. Finally, substituting (22) into (21) yields $\bar{u}_1^* = \exp\left(-\left(1-\bar{\gamma}_\varepsilon^*\right)\right)\bar{\gamma}_\varepsilon^*$. ■

Proof of Proposition 3. Lemmas 1 through 3 continue to hold, replacing r with r_1 or r_2 as appropriate. In particular, $\lambda(t) = \frac{r_2 v(t) - v'(t)}{1-v(t)}$, and the same argument as in the proof of Theorem 1 implies that the unique maxmin posture γ^* satisfies $\gamma^*(t) = \min\{e^{r_1 t}u_1^*, 1\}$, where u_1^* is the (unique) number such that the time at which $\gamma^*(t)$ reaches 1 equals $\tilde{T}(\gamma^*)$. Thus, given posture γ^*, it follows that $\lambda(t) = \frac{r_2\left(1-e^{r_1 t}u_1^*\right)+r_1 e^{r_1 t}u_1^*}{e^{r_1 t}u_1^*} = r_2\frac{e^{-r_1 t}}{u_1^*} + r_1 - r_2$. Now

$$\exp\left(-\int_0^{\tilde{T}(\gamma^*)}\left(r_2\frac{e^{-r_1 t}}{u_1^*}+r_1-r_2\right)dt\right) = \exp\left(-\frac{1}{u_1^*}\left(\frac{r_2}{r_1}\right)\left(1-e^{-r_1\tilde{T}(\gamma^*)}\right)+\left(r_1-r_2\right)\tilde{T}\left(\gamma^*\right)\right).$$

Setting this equal to ε and rearranging implies that $\tilde{T}(\gamma^*)$ is given by

$$e^{-r_1\tilde{T}(\gamma^*)} - \frac{r_1}{r_2}u_1^*\log\varepsilon + \left(\frac{r_1}{r_2}-1\right)u_1^* r_1\tilde{T}\left(\gamma^*\right) = 1. \tag{23}$$

Using the condition that $e^{r_1\tilde{T}(\gamma^*)}u_1^* = 1$, this can be rearranged to yield (11). Finally, there is a unique pair $\left(u_1^*, \tilde{T}(\gamma^*)\right)$ that satisfies both (23) and $e^{r_1\tilde{T}(\gamma^*)}u_1^* = 1$, because the curve in $\left(u_1^*, \tilde{T}(\gamma^*)\right)$ space defined by (23) is upward-sloping, while the curve defined by $e^{r_1\tilde{T}(\gamma^*)}u_1^* = 1$ is downward-sloping. ■

Proof of Lemma 4. The fact that $\Omega_2^{RAT}(\gamma) \subseteq \Pi_1^\gamma$ immediately implies that $u_1^{RAT}(\gamma) \geq u_1^*(\gamma) = \min_{t\leq T(\gamma)} e^{-rt}\underline{\gamma}(t)$. Therefore, it suffices to show that $u_1^{RAT}(\gamma) \leq \min_{t\leq T(\gamma)} e^{-rt}\underline{\gamma}(t)$.

Let $\dot{T} \equiv \min\operatorname{argmax}_t e^{-rt}\left(1-\underline{\gamma}(t)\right)$. Note that $\dot{T}$ is well-defined and finite because $\underline{\gamma}(t)$ is lower semi-continuous and $\lim_{t\to\infty} e^{-rt}\left(1-\underline{\gamma}(t)\right) = 0$. Let $\dot{\sigma}_2 \in \Sigma_2$ be the strategy that always demands 0, rejects up to time $\dot{T}$, accepts at date $\left(\dot{T},-1\right)$ if and only if $\lim_{\tau\uparrow\dot{T}}\gamma(\tau) \leq \gamma\left(\dot{T}\right)$, and accepts at all dates $\left(\dot{T},1\right)$ and later (for all histories). Let $\dot{\pi}_2^\gamma \in \Sigma_1$ be identical to the γ-offsetting belief π_2^γ, with the modification that $\dot{\pi}_2^\gamma$ always accepts demands of 0. Since $p(0) = 0$ for any posture γ, it follows that $u_1\left(\dot{\pi}_2^\gamma, \dot{\sigma}_2\right) = 1$, and therefore $\dot{\pi}_2^\gamma \in \Sigma_1^*\left(\dot{\sigma}_2\right)$. In addition, it is clear that $u_2\left(\pi_2^\gamma, \sigma_2\right) \geq u_2\left(\dot{\pi}_2^\gamma, \sigma_2\right)$ for all $\sigma_2 \in \Sigma_2$, so the observations that $\sigma_2^\gamma \in \Sigma_2^*\left(\pi_2^\gamma\right)$ (by Lemma 2) and $u_2\left(\pi_2^\gamma, \sigma_2^\gamma\right) = u_2\left(\dot{\pi}_2^\gamma, \sigma_2^\gamma\right)$ imply that $\sigma_2^\gamma \in \Sigma_2^*\left(\dot{\pi}_2^\gamma\right)$. Finally, it is clear that $\dot{\sigma}_2 \in \Sigma_2^* (\gamma)$, and $\gamma \in \Sigma_1^* (\sigma_2^\gamma)$ by Lemma 3. Summarizing, I have established that the arrows in the following diagram may be read as "is a best-response to" :

$$\begin{array}{ccc} \gamma & \rightarrow & \sigma_2^\gamma \\ \uparrow & & \downarrow \\ \dot{\sigma}_2 & \leftarrow & \dot{\pi}_2^\gamma \end{array}.$$

Therefore, the set $\{\gamma, \dot{\pi}_2^\gamma\} \times \{\sigma_2^\gamma, \dot{\sigma}_2\}$ is closed under rational behavior given posture γ, which implies that $\{\sigma_2^\gamma, \dot{\sigma}_2\} \subseteq \Omega_2^{RAT} (\gamma)$. Hence, $u_1^{RAT} (\gamma) \leq \sup_{\sigma_1} u_1 (\sigma_1, \sigma_2^\gamma) = u_1 (\gamma, \sigma_2^\gamma) = \min_{t \leq T(\gamma)} e^{-rt} \underline{\gamma} (t)$. ■

References

[1] Abreu, D., and Gul, F. (2000), "Bargaining and Reputation," *Econometrica*, 2000, 85-117.

[2] Abreu, D., and Pearce, D. (2007), "Bargaining, Reputation, and Equilibrium Selection in Repeated Games with Contracts," *Econometrica*, 75, 653-710.

[3] Battigalli, P., and Watson, J. (1997), "On 'Reputation' Refinements with Heterogeneous Beliefs," *Econometrica*, 65, 369-374.

[4] Bergin, J., and MacLeod, W.B. (1993), "Continuous Time Repeated Games," *International Economic Review*, 34, 21-37.

[5] Billingsley, P. (1995), *Probability and Measure*. New York, NY: John Wiley & Sons, Inc..

[6] Chatterjee, K., and Samuelson, L. (1987), "Bargaining with Two-Sided Incomplete Information: An Infinite-Horizon Model with Alternating Offers," *Review of Economic Studies*, 59, 175-192.

[7] Chatterjee, K., and Samuelson, L. (1988), "Bargaining Under Two-Sided Incomplete Information: The Unrestricted Offers Case," *Operations Research*, 36, 605-618.

[8] Cho, I.-K. (1994), "Stationarity, Rationalizability, and Bargaining," *Review of Economic Studies*, 61, 357-374.

[9] Compte, O., and Jehiel, P. (2002), "On the Role of Outside Options in Bargaining with Obstinate Parties," *Econometrica*, 70, 1477-1517.

[10] Compte, O., and Jehiel, P. (2004), "Gradualism in Bargaining and Contribution Games," *Review of Economic Studies*, 71, 975-1000.

[11] Crawford, V. (1982), "A Theory of Disagreement in Bargaining," *Econometrica*, 50, 607-637.

[12] Ellingsen, T., and Miettinen, T. (2008), "Commitment and Conflict in Bilateral Bargaining," *American Economic Review*, 98, 1629–1635.

[13] Feinberg, Y., and Skrzypacz, A. (2005), "Uncertainty about Uncertainty and Delay in Bargaining," *Econometrica*, 73, 69-91.

[14] Fershtman, C., and Seidmann, D.J. (1993), "Deadline Effects and Inefficient Delay in Bargaining with Endogenous Commitment," *Journal of Economic Theory*, 60, 306-321.

[15] Fudenberg, D., and Levine, D.K. (1989), "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, 57, 759-778.

[16] Kambe, S. (1999), "Bargaining with Imperfect Commitment," *Games and Economic Behavior*, 28, 217-237.

[17] Kreps, D.M. and Wilson, R. (1982), "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253-279.

[18] Knoll, M.S. (1996), "A Primer on Prejudgment Interest," *Texas Law Review*, 75, 293-374.

[19] Milgrom, P., and Roberts, J. (1982), "Predation, Reputation, and Entry Deterrence," *Journal of Economic Theory*, 27, 280-312.

[20] Muthoo, A. (1996), "A Bargaining Model Based on the Commitment Tactic," *Journal of Economic Theory*, 69, 134-152.

[21] Myerson, R. (1991), *Game Theory: Analysis of Conflict*. Cambridge, MA: Harvard University Press.

[22] Royden, H.L. (1988), *Real Analysis*, Third Edition. Upper Saddle River, NJ: Prentice Hall.

[23] Rubinstein, A. (1982), "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, 97-109.

[24] Schelling, T.C. (1956), "An Essay on Bargaining," *American Economic Review*, 46, 281-306.

[25] Simon, L.K., and Stinchcombe, M.B. (1989), "Extensive Form Games in Continuous time: Pure Strategies," *Econometrica*, 57, 1171-1214.

[26] Watson, J. (1993), "A 'Reputation' Refinement without Equilibrium," *Econometrica*, 61, 199-205.

[27] Watson, J. (1998), "Alternating-Offer Bargaining with Two-Sided Incomplete Information," *Review of Economic Studies*, 65, 573-594.

[28] Wolitzky, A. (2011), "Indeterminacy of Reputation Effects in Repeated Games with Contracts," *Games and Economic Behavior*, 73, 595-607.

[29] Yildiz, M. (2003), "Bargaining without a Common Prior: An Immediate Agreement Theorem," *Econometrica*, 71, 793-811.

[30] Yildiz, M. (2004), "Waiting to Persuade," *Quarterly Journal of Economics*, 119, 223-248.

Supplementary Appendix (NOT FOR PUBLICATION)

This appendix shows that the characterization of the maxmin payoff and posture (Theorem 1) continues to apply when the solution concept is strengthened from first-order knowledge of rationality to iterated conditional dominance, or when the continuous-time bargaining protocol of the text is replaced by any discrete-time bargaining protocol with sufficiently frequent offers. However, the characterization does not apply with *both* iterated conditional dominance and discrete-time bargaining, as the fact that (complete-information) discrete-time bargaining is solvable by iterated conditional dominance implies that the predictions of the model with iterated conditional dominance and discrete-time bargaining depend on the order and relative frequency of offers.[32]

Iterated Conditional Dominance

This section shows that Theorem 1 continues to hold under a natural notion of iterated conditional dominance. Because the model has incomplete information and is not a multistage game with observed actions (as players do not observe each other's choice of demand paths on the integers), no off-the-shelf version of iterated conditional dominance is applicable, and even the simplest version that is applicable requires some new notation.

For integer t, let $\sigma_i(h^t)$ be the element of $\Delta(\mathcal{U}^t)$ prescribed by strategy σ_i at date $(t,0)$ and history h^t. I first introduce the idea that a triple $\left(h^t, u_1^{\lfloor t\rfloor}, u_2^{\lfloor t\rfloor}\right)$ is "σ_i-coherent" if $u_i^{\lfloor t\rfloor} \in \operatorname{supp}\sigma_i\left(h^{\lfloor t\rfloor}\right)$ and at h^t the path of realized demands between $\lfloor t\rfloor$ and t coincides with $\left(u_1^{\lfloor t\rfloor}, u_2^{\lfloor t\rfloor}\right)$.

Definition 8 *A triple* $\left(h^t, u_1^{\lfloor t\rfloor}, u_2^{\lfloor t\rfloor}\right)$ *is* σ_i-coherent *if* $u_i^{\lfloor t\rfloor} \in \operatorname{supp}\sigma_i\left(h^{\lfloor t\rfloor}\right)$ *and* $\left(u_1^{\lfloor t\rfloor}(\tau), u_2^{\lfloor t\rfloor}(\tau)\right) =$

[32] Informally, there is a race between the number of rounds of iterated conditional dominance and the frequency of offers. I conjecture that, for any number of rounds of iterated conditional dominance, the maxmin payoff and posture in discrete-time bargaining converge (in the sense of Proposition 6) to the maxmin payoff and posture in continuous-time bargaining as offers become frequent. This is consistent with Rubinstein bargaining, where the round at which any demand other than 0 or 1 is deleted goes to infinity as the time between offers vanishes (so that iterated conditional dominance has no "bite" in the continuous-time limit). I thank Jeff Ely for helpful comments on this point.

$(u_1(\tau), u_2(\tau))$ *for all* $\tau \in [\lfloor t \rfloor, t]$, *where* $h^t = (u_1(\tau), u_2(\tau))_{\tau \le t}$. *A history* h^t *is* σ_i-coherent *if there exist demand paths* $\left(u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ *such that* $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ *is* σ_i*-coherent.*

For any strategy profile (σ_1, σ_2) and any triple $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ such that $\left(u_1^{\lfloor t \rfloor}(\tau), u_2^{\lfloor t \rfloor}(\tau)\right) = (u_1(\tau), u_2(\tau))$ for all $\tau \in [\lfloor t \rfloor, t]$, where $h^t = (u_1(\tau), u_2(\tau))_{\tau \le t}$, each player i's expected payoff under strategy profile (σ_i, σ_j) conditional on reaching the triple $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ is well-defined, and is denoted

$$u_i\left(\sigma_i, \sigma_j | h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right).$$

I also write $u_i(\sigma_i, \sigma_j | h^t)$ for player i's expected payoff conditional on reaching h^t at date $(t, 0)$ for integer t.

I now define iterated conditional dominance. Informally, the idea is that a strategy is conditionally dominated if it is either strictly dominated or is "conditionally weakly dominated." The difference between the definition of iterated conditional dominance for the two players reflects the fact that player 1 is committed to strategy γ with probability ε, and therefore that player 2 is restricted to assigning probability at least ε to strategy γ at histories that are consistent with γ. Note that the support of the γ-offsetting belief π_2^γ includes strategies that are iteratively conditionally dominated, as it is easy to verify that any strategy of player 1's that ever accepts a demand of 1 is iteratively conditionally dominated.

Definition 9 *For any posture* γ *and set of bargaining phase strategy profiles* $\Omega = \Omega_1 \times \Omega_2 \subseteq \Sigma_1 \times \Sigma_2$, *a strategy* $\sigma_1 \in \Sigma_1$ *is* conditionally dominated with respect to (γ, Ω) *if either of the following conditions hold*

1. *There exists a strategy* $\sigma_1' \in \Sigma_1$ *such that*

$$u_1(\sigma_1', \pi_1) > u_1(\sigma_1, \pi_1)$$

   for all beliefs $\pi_1 \in \Delta(\Omega_2)$.

2. *There exists a strategy* $\sigma_1' \in \Sigma_1$ *such that*

$$u_1\left(\sigma_1', \pi_1 | h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right) \ge u_1\left(\sigma_1, \pi_1 | h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$$

   for all σ_1*-coherent* $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ *and all beliefs* $\pi_1 \in \Delta(\Omega_2)$, *with strict inequality for some* σ_1*-coherent* $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ *and some belief* $\pi_1 \in \Delta(\Omega_2)$.

A strategy $\sigma_2 \in \Sigma_2$ *is* conditionally dominated with respect to (γ, Ω) *if either of the following conditions hold*

1. *There exists a strategy* $\sigma_2' \in \Sigma_2$ *such that*

$$u_2(\sigma_2', \pi_2) > u_2(\sigma_2, \pi_2)$$

for all beliefs $\pi_2 \in \Delta(\Omega_1 \cup \{\gamma\})$ *such that* $\pi_2(\gamma) \geq \varepsilon$ *with strict inequality only if* $\gamma \in \Omega_1$.

2. *There exists a strategy* $\sigma_2' \in \Sigma_2$ *such that*

$$u_2\left(\sigma_2', \pi_2 | h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right) \geq u_2\left(\sigma_2, \pi_2 | h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$$

for all σ_2*-coherent* $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ *that are inconsistent with* γ *and all beliefs* $\pi_2 \in \Delta(\Omega_1)$, *with strict inequality for some* σ_2*-coherent* $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ *that is inconsistent with* γ *and some belief* $\pi_2 \in \Delta(\Omega_1)$.

A set of bargaining phase strategy profiles $\Omega = \Omega_1 \times \Omega_2 \subseteq \Sigma_1 \times \Sigma_2$ *is* closed under conditional dominance given posture γ *if every* $\sigma_i \in \Omega_i$ *is conditionally undominated (i.e., not conditionally dominated) with respect to* (γ, Ω). *The set of* iteratively conditionally undominated strategies given posture γ *is*

$$\Omega^{ICD}(\gamma) \equiv \bigcup \{\Omega : \Omega \textit{ is closed under conditional dominance given posture } \gamma\}.$$

Player 1*'s* maxmin payoff under iterated conditional dominance given posture γ *is*

$$u_1^{ICD}(\gamma) \equiv \sup_{\sigma_1} \inf_{\sigma_2 \in \Omega_2^{ICD}(\gamma)} u_1(\sigma_1, \sigma_2).$$

Player 1*'s* maxmin payoff under iterated conditional dominance *is*

$$u_1^{ICD} \equiv \sup_{\gamma} u_1^{ICD}(\gamma).$$

A posture γ^{ICD} *is a* maxmin posture under iterated conditional dominance *if there exists a sequence of postures* $\{\gamma_n\}$ *such that* $\gamma_n(t) \to \gamma^{ICD}(t)$ *for all* $t \in \mathbb{R}_+$ *and* $u_1^{ICD}(\gamma_n) \to u_1^{ICD}$.

This version of iterated conditional dominance is stronger than rationalizability, in that $\Omega^{ICD}(\gamma) \subseteq \Omega^{RAT}(\gamma)$ for any posture γ. This can be seen by noting that every set Ω that is closed under conditional dominance is also closed under rationalizability, because rationalizability is equivalent to imposing only the first of the two conditions in the definition of conditional dominance (for both player 1 and player 2). An immediate consequence of this observation is that the maxmin payoff under iterated conditional dominance is weakly greater than the maxmin payoff (under first-order knowledge of rationality), that is, $u_1^{ICD} \geq u_1^*$. In fact, the two payoffs are equal, as are the corresponding maxmin postures.

Proposition 5 *Player 1's maxmin payoff under iterated conditional dominance equals her maxmin payoff, and the unique maxmin posture under iterated conditional dominance is the unique maxmin posture. That is, $u_1^{ICD} = u_1^*$, and the unique maxmin posture under iterated conditional dominance is $\gamma^{ICD} = \gamma^*$.*

The rest of this section is devoted to proving Proposition 5. The proof builds on that of Proposition 4. This is because it can be shown that the set of iteratively conditionally undominated strategies and the set of rationalizable strategies are identical up to strategies that are "exceptional" in the following sense.

Definition 10 *A strategy $\sigma_i \in \Sigma_i$ is* exceptional given posture γ *if either of the following conditions hold:*

- *$i \in \{1, 2\}$ and σ_i ever accepts a demand of 1, rejects a demand of 0, makes a demand of 0 or a path of demands converging to 0 (i.e., $\lim_{\tau \uparrow t} u_i(\tau) = 0$), or makes a demand of 1 at every successor of some history h^t.*

- *$i = 1$ and σ_i ever accepts a demand $u_2(t) \geq 1 - e^{-r(t^*-t)}\underline{\gamma}(t^*) > 0$ at any history h^t consistent with γ with $t \leq t^*$, rejects a demand $u_2(t) \leq 1 - e^{-r(t^*-t)}\underline{\gamma}(t^*) < 1$ at any history h^t consistent with γ with $t \leq t^*$, or demands $\gamma(t^*)$ at any history h^{t^*} consistent with γ.*

The relationship between iterated conditional dominance and rationalizability is formalized in the following lemma, which is a key step in the proof of Proposition 5.

Lemma 5 *For any posture γ, every strategy that is rationalizable and non-exceptional given posture γ is also iteratively conditionally undominated given posture γ.*

The proof of the lemma uses the concept of a unique optimal action: an action (accepting, rejecting, or choosing a demand path for the next integer) is the *unique optimal action* at a triple $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ under a belief π_i if every strategy σ_i that maximizes $u_i\left(\sigma_i, \pi_i | h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ prescribes that action at history h^t (where the arguments $\left(u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ are omitted in the case of choosing a demand path for the next integer).

Proof of Lemma 5. Fix a posture γ; for the duration of the proof, I omit the modifier "given posture γ." To prove the lemma, I show that for every non-exceptional strategy σ_i and every σ_i-coherent history h^t that is inconsistent with γ, there exist demand paths $\left(u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ and belief π_i with support on strategies that are rationalizable and non-exceptional such that $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ is σ_i-coherent and σ_i prescribes the unique optimal action at $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ under belief π_2. If $i = 1$, this conclusion also holds at σ_i-coherent histories that are consistent with γ. This implies that the second of the two conditions in the definition of conditional dominance can never hold if σ_i is non-exceptional, for $i = 1, 2$. Therefore, every non-exceptional strategy that is conditionally dominated is also strictly dominated, and hence every non-exceptional strategy that is rationalizable is also iteratively conditionally undominated.

I start by establishing a statement with the important implication that, starting from a history that is inconsistent with γ, any continuation strategy is part of a rationalizable strategy.

Step 1: Any strategy $\sigma_1 \in \Sigma_1$ that demands $\gamma(t)$ and rejects player 2's demand at every history h^t that is consistent with γ is rationalizable. Any strategy $\sigma_2 \in \Sigma_2$ that demands 1 *and accepts at (and not before) the more favorable of dates $(t^*, -1)$ and $(t^*, 1)$ if player 1 follows γ until time t^* is rationalizable.*

Proof. By the proof of Lemma 4, strategy $\dot{\pi}_2^{\gamma}$ is rationalizable for player 1, and strategy σ_2^{γ} is rationalizable for player 2. Now if strategies σ_1 and σ_2 are as in the statement, then $\sigma_1 \in \Sigma_1^*(\sigma_2^{\gamma})$ and $\sigma_2 \in \Sigma_2^*(\dot{\pi}_2^{\gamma})$, so σ_1 and σ_2 are rationalizable as well. ■

Step 2: For $i = 1, 2$, if a strategy σ_i is non-exceptional and a history h^t is σ_i-coherent and inconsistent with γ, then there exist demand paths $\left(u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ and a belief π_i with *support on strategies that are rationalizable and non-exceptional such that* $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ *is* σ_i*-coherent and* σ_i *prescribes the unique optimal action at* $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ *under belief* π_i*.*

Proof. Fix a non-exceptional strategy σ_i and a history h^t that is σ_i-coherent and inconsistent with γ. Step 1 implies that any continuation strategy of player j's is part of a rationalizable strategy. Hence, the restriction that π_i has support on strategies that are rationalizable and non-exceptional implies only that continuation strategies are non-exceptional.

Suppose that σ_i accepts at h^t. Then the fact that σ_i is non-exceptional and h^t is σ_i-coherent imply that $u_i(t) > 0$ and $u_j(t) < 1$. Let $\left(u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ specify that the players continue to demand $u_i(t)$ and $u_j(t)$ until $\lceil t \rceil$, and let π_i assign probability 1 to a rationalizable strategy under which at every successor history of h^t player j demands $\frac{1+u_j(t)}{2}$ (after time $\lceil t \rceil$) and rejects any strictly positive demand; such a strategy exists by the previous paragraph, and is clearly non-exceptional (in particular, player j always chooses demand paths that always make demands $u_2(\tau) \in (0,1)$). Then it is clear that accepting at h^t is the optimal action at $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ under belief π_i.[33]

Suppose that σ_i rejects at h^t. Then the fact that σ_i is non-exceptional and h^t is σ_i-coherent imply that $u_i(t) > 0$ and $u_j(t) > 0$. Let π_i assign probability 1 to a rationalizable and non-exceptional strategy under which player j reduces his demand to $u_j(t)/2$ by some time τ such that

$$e^{-r(\tau-t)}\left(1 - \frac{u_j(t)}{2}\right) > 1 - u_j(t),$$

subsequently demands $u_j(t)/2$ forever, and rejects player i's demand at every successor history of h^t unless player i demands 0 (or let $\left(u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ specify that player j's demand follows such a path, in case there is no integer between t and τ). Choose any $\left(u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ such that $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ is σ_i-coherent and player j's demands follow such a path. Now rejecting until time τ and then accepting (while never demanding 0) is strictly better for player i under belief π_i than is accepting at h^t, so rejecting is the unique optimal action at $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ under belief π_i.

Finally, suppose that t is an integer and that σ_i chooses demand path u_i^t at h^t. The fact that σ_i is non-exceptional implies that $u_i^t(\tau) > 0$ for all $\tau \in [t, t+1)$ and that

[33] Note that the possibility that player i could reject at h^t but accept "immediately" after h^t is ruled out by the assumption that the probability that a player accepts by date $(t,1)$ is right-continuous in t.

$\lim_{\tau \uparrow t+1} u_i^t(\tau) > 0$. Let π_i assign probability 1 to a rationalizable and non-exceptional strategy under which player j demands $1 - \frac{e^{-r} \lim_{\tau \uparrow t+1} u_i^t(\tau)}{2}$ at h^t and at every successor history of h^t; accepts at date $(t+1, -1)$ if $u_i(\tau) = u_i^t(\tau)$ for all $\tau \in [t, t+1)$; and otherwise rejects any strictly positive demand at every successor history of h^t. Now choosing demand path u_i^t at h^t and rejecting player j's demand until time $t+1$ yields payoff $e^{-r} \lim_{\tau \uparrow t+1} u_i^t(\tau)$ under belief π_i, while every other continuation strategy yields payoff at most $\frac{e^{-r} \lim_{\tau \uparrow t+1} u_i^t(\tau)}{2}$ under belief π_i, so choosing demand path u_i^t is the unique optimal action at h^t under belief π_i. ■

Step 3: If strategy $\sigma_1 \in \Sigma_1$ is non-exceptional and a history h^t is σ_1-coherent and consistent with γ, then there exist demand paths $\left(u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ and a belief π_1 with support on strategies that are rationalizable and non-exceptional such that $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ is σ_1-coherent and σ_1 prescribes the unique optimal action at $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ under belief π_1.

Proof. If $t > t^*$, then if player 2 plays a rationalizable and non-exceptional strategy σ_2 that accepts at time t^* under strategy profile (γ, σ_2) (which exists), then player 2's continuation play starting from h^t is restricted only by the requirement that it is non-exceptional. Hence, the proof in this case is just like the proof of Step 2. I therefore assume that $t \leq t^*$.

Suppose that σ_1 accepts at h^t. Then the fact that σ_1 is non-exceptional, h^t is σ_1-coherent and consistent with γ, and $t \leq t^*$ implies that $u_1(t) > 0$ and $u_2(t) < 1 - e^{-r(t^*-t)} \underline{\gamma}(t^*)$. Define the strategy $\ddot{\sigma}_2 \in \Sigma_2$ as follows:

- If h^τ is consistent with γ, then demand 1 until time $\lceil t^* + 1 \rceil$, subsequently demand $\frac{1}{2}$, reject all positive demands until the more favorable of dates $(t^*, -1)$ and $(t^*, 1)$, and subsequently accept all demands of less than 1.

- If h^τ is inconsistent with γ, then demand $\frac{1+u_2(t)}{2}$ and reject all positive demands.

Note that $\ddot{\sigma}_2 \in \Sigma_2^*(\gamma)$, so $\ddot{\sigma}_2$ is rationalizable. In addition, $\ddot{\sigma}_2$ is clearly non-exceptional. Let π_1 assign probability 1 to $\ddot{\sigma}_2$, and let $\left(u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ specify that player 1 demands $u_1(\tau) = \gamma(\tau)$ for all $\tau \in [t, \lceil t \rceil)$ and that player 2 continues to demand $u_2(t)$ until $\lceil t \rceil$. Then accepting at $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ yields payoff $1 - u_2(t)$ under belief π_1, while any strategy that rejects at $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ yields strictly less. So accepting is the unique optimal action at $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ under belief π_1.

Suppose that σ_1 rejects at h^t. Then the fact that σ_1 is non-exceptional, h^t is σ_1-coherent and consistent with γ, and $t \leq t^*$ implies that $u_1(t) > 0$ and $u_2(t) > 1 - e^{-r(t^*-t)}\underline{\gamma}(t^*)$. Let $\ddot{\sigma}_2$, π_1, and $\left(u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ be as above, with the modification that $\ddot{\sigma}_2$ demands $1 - \frac{e^{-r(t^*-t)}\underline{\gamma}(t^*)}{2}$ rather than $\frac{1+u_2(t)}{2}$ at histories h^τ that are inconsistent with γ.[34] Then rejecting and following strategy γ at $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ yields payoff $e^{-r(t^*-t)}\underline{\gamma}(t^*)$ under belief π_1, while any strategy that rejects at $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ yields strictly less. So rejecting is the unique optimal action at $\left(h^t, u_1^{\lfloor t \rfloor}, u_2^{\lfloor t \rfloor}\right)$ under belief π_1.

Finally, suppose that t is an integer and that σ_1 chooses demand path u_1^t at h^t. Since u_1^t and γ are continuous on $[t, t+1)$, there are three cases.

1. $u_1^t(\tau) = \gamma(\tau)$ for all $\tau \in [t, t+1)$.

2. $u_1^t(t) = \gamma(t)$ but $u_1^t(\tau) \neq \gamma(\tau)$ for some $\tau \in [t, \min\{t+1, t^*\})$.

3. $u_1^t(t) \neq \gamma(t)$.

Start with case 1. Here, the fact that σ_1 is non-exceptional, h^t is consistent with γ, and $t \leq t^*$ implies that in fact $t+1 \leq t^*$, as σ_1 never demands $\gamma(t^*)$ at a history h^{t^*} consistent with γ. Now for all $\eta > 0$, Step 1 implies that there exists a rationalizable and non-exceptional strategy σ_2 that demands 1 at all times τ such that $e^{-r(\tau-(t+1))} \geq \eta$, accepts at (but not before) time t^* under strategy profile (γ, σ_2), rejects all strictly positive demands at dates $(t+1, -1)$ and earlier, rejects all strictly positive demands at dates after $(t+1, 0)$ such that $u_1(\tau) \neq \gamma(\tau)$ for some $\tau < t+1$, and accepts all demands less than 1 at dates after $(t+1, 0)$ such that $u_1(\tau) = \gamma(\tau)$ for all $\tau < t+1$ but $u_1(t+1) \neq \gamma(t+1)$. Since $e^{-r(t^*-(t+1))}\underline{\gamma}(t^*) < 1$ (which follows from the definition of t^*), choosing demand path u_1^t and then deviating from γ to a demand close to 1 at time $t+1$ yields a strictly higher payoff under any belief that assigns probability 1 to such a strategy than does choosing any other demand path, for η sufficiently small.

For case 2, the fact that σ_1 is non-exceptional, h^t is consistent with γ, and $t \leq t^*$ implies that $t < t^*$. Let τ_0 be the infimum over times $\tau \in [t, \min\{t+1, t^*\})$ such that $u_1^t(\tau) \neq \gamma(\tau)$.

[34] This modification serves only to ensure that $\ddot{\sigma}_2$ does not demand 1 forever after some history, and thus that $\ddot{\sigma}_2$ is non-exceptional.

Now for all $\eta > 0$, Step 1 implies that there exists a rationalizable and non-exceptional strategy σ_2 that demands 1 at all times τ such that $e^{-r(\tau - \max\{t+1,t^*\})} \geq \eta$, accepts at (but not before) time t^* under strategy profile (γ, σ_2), rejects all strictly positive demands at all histories that are inconsistent with either γ or u_1^t, and, for all $k \in \left\{0, 1, \ldots, \left\lfloor \frac{t+1-\tau_0}{\eta} \right\rfloor - 1\right\}$, accepts at time $\tau_0 + k\eta$ with probability $\eta^k (1-\eta)$ if $u_1(\tau) = u_1^t(\tau)$ for all $\tau \in [t, \tau_0 + k\eta]$, and accepts at date $(t+1, -1)$ with probability $\eta^{\left\lfloor \frac{t+1-\tau_0}{\eta} \right\rfloor}$ if player 1's demands are consistent with u_1^t on $[t, t+1)$. Since u_1^t is continuous, player 1 receives a strictly higher payoff from choosing u_1^t and then rejecting until time $t+1$ than from mimicking γ, under the belief that player 2 plays such a strategy for sufficiently small η (as player 1 strictly prefers to have her demand accepted at any time prior to t^* than at t^*, by definition of t^*). In addition, player 1 receives a strictly higher payoff from choosing u_1^t than from choosing any demand path that coincides with u_1^t until some time $\tau \in [t, t+1)$ and then diverges from u_1^t. Therefore, choosing demand path u_1^t is player 1's unique optimal action at h^t under the belief that player 2 plays such a strategy for sufficiently small η.

For case 3, Step 1 implies that for all $\eta > 0$ there exists a rationalizable and non-exceptional strategy σ_2 that accepts at (but not before) time t^* under strategy profile (γ, t^*), rejects all strictly positive demands at all histories that are inconsistent with γ, and, for all $k \in \left\{0, 1, \ldots, \left\lfloor \frac{1}{\eta} \right\rfloor - 2\right\}$, with probability $\eta^k (1-\eta)$ demands 1 until time $t + k\eta$ and reduces its demand to η by time $t + (k+1)\eta$, and with probability $\eta^{\left\lfloor \frac{1}{\eta} \right\rfloor}$ demands 1 until time $t + 1 - \eta$ and reduces its demand to η by time $t+1$. Step 1 also implies that there exists a rationalizable and non-exceptional strategy that accepts at (but not before) time t^* under strategy profile (γ, t^*), demands $1 - \left(\frac{e^{-r} \lim_{\tau \uparrow t+1} u_i^t(\tau)}{2}\right)\eta$ at h^t and at every successor history of h^t; accepts at date $(t+1, -1)$ if $u_i(\tau) = u_i^t(\tau)$ for all $\tau \in [t, t+1)$; and otherwise rejects any strictly positive demand at every successor history of h^t (as in the last part of the proof of Step 2). Let π_1 assign probability $1 - \eta$ to player 2's playing a strategy of the first kind and assign probability η to player 2's playing a strategy of the second kind. I claim that, for η sufficiently small, u_1^t is player 1's unique optimal action at h^t under belief π_1. To see this, first note that $e^{-r(t^*-t)}\underline{\gamma}(t^*) < 1$ implies that choosing u_1^t is strictly better than choosing any demand path that coincides with γ until time t^*, for η sufficiently small. Finally, any strategy that chooses a demand path other than u_1^t that diverges from γ before time t^* does no better than choosing demand path u_1^t and rejecting until time $t+1$ in the event that player 2 plays a strategy of the first kind, and does strictly worse in the event that player 2 plays a strategy of the second kind. Therefore, choosing demand path u_1^t is player 1's unique optimal action at h^t under belief π_1, for $\eta > 0$ sufficiently small. ■

I now complete the proof of the lemma. By the definition of conditional dominance and Step 2, strategy σ_2 can be conditionally dominated (with respect to some (γ, Ω)) by strategy σ_2' only if either σ_2' strictly dominates σ_2 (with respect to (γ, Ω); i.e., if the first condition in the definition of conditionally dominance with respect to (γ, Ω) holds) or σ_2' agrees with σ_2 at all σ_2-coherent histories that are inconsistent with γ. But, again by the definition of conditional dominance, if σ_2' conditionally dominates σ_2 and agrees with σ_2 at all σ_2-coherent histories that are inconsistent with γ, then σ_2' must strictly dominate σ_2. The same argument applies for player 1, noting that Steps 2 and 3 imply that a strategy σ_1 can be conditionally dominated by a strategy σ_1' only if σ_1' strictly dominates σ_1 or if σ_1' agrees with σ_1 at all σ_1-coherent histories h^t, whether or not h^t is consistent with γ (which is needed for the argument given the difference in the definitions of conditional dominance for players 1 and 2). Therefore, for $i = 1, 2$, if σ_i is non-exceptional then it cannot be conditionally dominated unless it is also strictly dominated. Finally, if σ_i is rationalizable and non-exceptional, then it is not strictly dominated with respect to $\left(\gamma, \Omega^{RAT}\right)$, hence not conditionally dominated with respect to $\left(\gamma, \Omega^{RAT}\right)$, and hence also not conditionally dominated with respect to the smaller set $\left(\gamma, \Omega^{ICD}\right)$. This proves that every rationalizable and non-exceptional strategy is iteratively conditionally undominated. ■

I now prove Proposition 5.

Proof of Proposition 5. As was the case for Proposition 4, it suffices to show that $u_1^{ICD}(\gamma) = \underline{\gamma}(t^*)$ for every posture γ. The proof proceeds by approximating the γ-offsetting belief π_2^γ with beliefs $\{\pi_2^\gamma(\eta)\}_{\eta>0}$ that have support on rationalizable and non-exceptional strategies only (unlike the offsetting belief π_2^γ itself, which assigns positive probability to the exceptional strategy $\tilde{\gamma}$), and by approximating the γ-offsetting strategy σ_2^γ with non-exceptional strategies $\{\sigma_2^\gamma(\eta)\}_{\eta>0}$ such that $\sigma_2^\gamma(\eta) \in \Sigma_2^*(\pi_2^\gamma(\eta))$ for all $\eta > 0$. Lemma 5 then implies that the strategy $\sigma_2^\gamma(\eta)$ is iteratively conditionally undominated for all $\eta > 0$.

Finally, as $\eta \to 0$,

$$\sup_{\sigma_1} u_1\left(\sigma_1, \sigma_2^{\gamma}(\eta)\right) \to u_1^*(\gamma),$$

which implies that $u_1^{ICD}(\gamma) \leq u_1^*(\gamma)$. Since $u_1^{ICD}(\gamma) \geq u_1^*(\gamma)$ is immediate because $\Omega^{ICD}(\gamma) \subseteq \Omega^{RAT}(\gamma)$, this shows that $u_1^{ICD}(\gamma) = u_1^*(\gamma) = \underline{\gamma}(t^*)$, completing the proof of the proposition.

I now present an argument leading to the construction of the beliefs $\{\pi_2^{\gamma}(\eta)\}_{\eta>0}$ and strategies $\{\sigma_2^{\gamma}(\eta)\}_{\eta>0}$. I start by defining the strategies of player 1's that receive positive weight under belief $\pi_2^{\gamma}(\eta)$. Fix $t \in (0, t^*)$ and $\eta \in \left(0, \frac{1}{2}\right)$. Let

$$\eta' \equiv \begin{cases} \min\left\{\eta, \frac{r(t^*-t)}{3}, \frac{\gamma(t^*)}{2}\right\} & \text{if } \gamma(t^*) > 0 \\ \min\left\{\eta, \frac{r(t^*-t)}{3}\right\} & \text{if } \gamma(t^*) = 0 \end{cases},$$

and let $\tilde{\gamma}(t, \eta)$ be the strategy that demands $u_1(\tau) = \gamma(\tau)$ for all $\tau \in [0, t)$; demands

$$u_1(\tau) = \frac{t + \eta'/r - \tau}{\eta'/r}\gamma(t) + \left(1 - \frac{t + \eta'/r - \tau}{\eta'/r}\right)(1 - \eta')$$

for all $\tau \in [t, t + \eta'/r]$; demands

$$u_1(\tau) = \frac{t + 2\eta'/r - \tau}{\eta'/r}(1 - \eta') + \left(1 - \frac{t + 2\eta'/r - \tau}{\eta'/r}\right)\eta'$$

for all $\tau \in [t + \eta'/r, t + 2\eta'/r]$; demands $u_1(\tau) = \eta'$ if $\tau > t + 2\eta'/r$; and accepts a demand of player 2's if and only if it equals 0. Intuitively, $\tilde{\gamma}(t, \eta)$ mimics γ until time t and then quickly rises to almost one before quickly falling to almost zero, where "quickly" and "almost zero" are both measured by η (the point of having η' rather than η in the formulas will become clear shortly).

I claim that $\tilde{\gamma}(t, \eta)$ is iteratively conditionally undominated. To see this, observe that $\tilde{\gamma}(t, \eta)$ is a best response to any strategy σ_2 with the following properties:

- σ_2 demands 1 and accepts at (and not before) date $(t^*, -1)$ if player 1 follows γ until time t^*.
- If h^{τ} is inconsistent with γ but consistent with $\tilde{\gamma}(t, \eta)$, then σ_2 demands 1 and accepts if and only if $\tau \geq t + \eta'/r$.

- If h^τ is inconsistent with both γ and $\tilde{\gamma}(t,\eta)$, then σ_2 demands 1 and rejects player 1's demand.

This follows because playing $\tilde{\gamma}(t,\eta)$ against such a strategy σ_2 yields payoff

$$e^{-rt-\eta'}(1-\eta'),$$

while the only other positive payoff that can be obtained against strategy σ_2 is

$$e^{-rt^*}\gamma(t^*) \leq e^{-rt^*},$$

and $e^{-rt-\eta'}(1-\eta') \geq e^{-rt^*}$ because $\eta' \leq \min\left\{\frac{1}{2}, \frac{r(t^*-t)}{3}\right\}$ (as can be easily checked). Now, by Step 1 of the proof of Lemma 5, there exists a rationalizable strategy σ_2 of this form, so $\tilde{\gamma}(t,\eta)$ is rationalizable. In addition, $\tilde{\gamma}(t,\eta)$ is non-exceptional, because $\gamma(t) > 0$ for all $t \in [0,t^*)$ (recalling the definition of γ^*) and $\tilde{\gamma}(t,\eta)$ always demands $\eta' \neq \gamma(t^*)$ at time t^*, so Lemma 5 implies that $\tilde{\gamma}(t,\eta)$ is iteratively conditionally undominated.

I now introduce versions of some of the key objects of Section 3.2, indexed by η. Let

$$\lambda(t,\eta) = \frac{rv(t) - v'(t)}{e^{-2\eta'}(1-\eta') - v(t)}$$

if v is differentiable at t and $v(t) < e^{-2\eta'}(1-\eta')$, and let $\lambda(t,\eta) = 0$ otherwise; and let

$$p(t,\eta) = \frac{v(t,-1) - v(t)}{e^{-2\eta'}(1-\eta') - v(t)}$$

if $v(t) < v(t,-1) \leq e^{-2\eta'}(1-\eta')$, and let $p(t,\eta) = 0$ otherwise. Define $\tilde{T}(\eta)$, $T(\eta)$, $t^*(\eta)$, $\hat{\lambda}(t,\eta)$, and $\hat{p}(t,\eta)$ as in Section 3.2, with $\lambda(t,\eta)$ and $p(t,\eta)$ replacing $\lambda(t)$ and $p(t)$ in the definitions. Note that, as $\eta \to 0$, $\lambda(t,\eta) \downarrow \lambda(t)$ and $p(t,\eta) \downarrow p(t)$ for all $t \in \mathbb{R}_+$. Hence, $\hat{\lambda}(t,\eta) \downarrow \lambda(t)$, $\hat{p}(t,\eta) \downarrow p(t)$, $\tilde{T}(\eta) \uparrow \tilde{T}$, $T(\eta) \uparrow T$, and $t^*(\eta) \uparrow t^*$.

Let $\mu^\gamma(\eta)$ be the belief that player 1 rejects all non-zero demands of player 2's and that her path of demands begins by following $\gamma(t)$ and then switches to following $\tilde{\gamma}(t,\eta)$ at time t with hazard rate $\hat{\lambda}(t,\eta)$ and discrete probability $\hat{p}(t,\eta)$, for all $t < t^*(\eta)$.[35] Let $\mu^{\gamma,t}(\eta)$ be the belief that coincides with $\mu^\gamma(\eta)$ until date $(t,-1)$ and subsequently coincides with the

[35] There is a technical problem here because it is not clear that $\mu^\gamma(\eta)$ can be written as a *finite*-dimensional distribution over Σ_1, that is, as an element of $\Delta(\Sigma_1)$. However, it should be clear that $\mu^\gamma(\eta)$ can in turn be approximated by a finite-dimensional distribution over Σ_1 in a way that suffices for the proof.

belief that player 1 follows γ. Let $\pi_2^{\gamma}(\eta)$ put probability ε on strategy γ and put probability $1-\varepsilon$ on strategy $\mu^{\gamma,t^*}(\eta)$. Let $\sigma_2^{\gamma}(\eta)$ be some best response to $\pi_2^{\gamma}(\eta)$ with the following properties:

- $\sigma_2^{\gamma}(\eta)$ demands 1 at all times t such that $t \leq t^*$ and $e^{-rt} \geq \eta$.
- $\sigma_2^{\gamma}(\eta)$ rejects player 1's demand at those histories that are consistent with $\pi_2^{\gamma}(\eta)$ where accepting and rejecting are both optimal actions.
- $\sigma_2^{\gamma}(\eta)$ rejects all positive demands at histories that are inconsistent with γ.
- $\sigma_2^{\gamma}(\eta)$ is non-exceptional.

It is clear that such a strategy exists. Furthermore, any such strategy is rationalizable, by Step 1 of the proof of Lemma 5, and hence any such strategy is iteratively conditionally undominated, by Lemma 5.

I claim that as $\eta \to 0$, the time at which agreement is reached under strategy profile $(\gamma, \sigma_2^{\gamma}(\eta))$ converges to t^*, uniformly over possible choices of $\sigma_2^{\gamma}(\eta)$ satisfying the above properties. To see this, observe that, as in Section 3.2, if $v(t) < e^{-2\eta'}(1-\eta')$ then $\lambda(t,\eta)$ and $p(t,\eta)$ are the rate and probability of player 1's switching to $\tilde{\gamma}(t,\eta)$ that make player 2 indifferent between accepting and rejecting γ. And, under belief $\pi_2^{\gamma}(\eta)$, player 2 believes that player 1 switches to $\tilde{\gamma}(t,\eta)$ with rate and probability $\lambda(t,\eta)$ and $p(t,\eta)$ if $v(t) < e^{-2\eta'}(1-\eta')$ and $t < t^*(\eta)$. Furthermore, since $\gamma(t)$ is positive and continuous on $[0,t^*)$, it follows that

$$\liminf_{\eta \to 0} \left\{ t : v(t) \geq e^{-2\eta'}(1-\eta') \right\} = t^*.$$

Hence, for small η player 2 is indifferent between accepting and rejecting γ until close to time $\min\{t^*, t^*(\eta)\}$, and therefore $\sigma_2^{\gamma}(\eta)$ specifies that he rejects until close to time $\min\{t^*, t^*(\eta)\}$. Since $t^*(\eta) \to t^*$, this shows that the time at which agreement is reached under strategy profile $(\gamma, \sigma_2^{\gamma}(\eta))$ converges to t^*.

The proof is nearly complete. Strategy $\sigma_2^{\gamma}(\eta)$ is iteratively conditionally undominated for all $\eta > 0$. When facing strategy $\sigma_2^{\gamma}(\eta)$, the highest payoff that player 1 can receive when player 2 accepts at a history that is consistent with γ converges to $\underline{\gamma}(t^*)$ as $\eta \to 0$. Furthermore, for any $\eta > 0$, the most player 2 accepts at a history that is inconsistent with

γ is η'; and the highest payoff player 1 can receive by accepting a demand of player 2's is η (since $\sigma_2^\gamma(\eta)$ demands 1 at all times t such that $e^{-rt} \geq \eta$). It follows that

$$u_1^{ICD}(\gamma) \leq \lim_{\eta\to 0} \sup_{\sigma_1} u_1\left(\sigma_1, \sigma_2^\gamma(\eta)\right) = \max\left\{\underline{\gamma}(t^*), \lim_{\eta\to 0}\eta', \lim_{\eta\to 0}\eta\right\} = \max\left\{\underline{\gamma}(t^*), 0, 0\right\} = \underline{\gamma}(t^*),$$

completing the proof. ■

Discrete-Time Bargaining with Frequent Offers

This section shows that Theorem 1 continues to hold when the continuous-time bargaining protocol of the text is replaced by any discrete-time bargaining protocol with sufficiently frequent offers. More precisely, for any sequence of discrete-time bargaining games that converges to continuous time (in that each player may make an offer close to any given time), the corresponding sequence of maxmin payoffs and postures converges to the continuous-time maxmin payoff and posture given by Theorem 1. Abreu and Gul (2000) provide a similar independence-of-procedures result for sequential equilibrium outcomes of reputational bargaining. Because my result concerns maxmin payoffs and postures rather than equilibria, my proof is very different from Abreu and Gul's.

Formally, replace the (continuous time) bargaining phase of Section 2 with the following procedure: There is a (commonly known) function $g : \mathbb{R}_+ \to \{0,1,2\}$ that specifies who makes an offer at each time. If $g(t) = 0$, no player takes an action at time t. If $g(t) = i \in \{1,2\}$, then player i makes a demand $u_i(t) \in [0,1]$ at time t, and player j immediately accepts or rejects. If player j accepts, the game ends with payoffs $\left(e^{-rt}u_i(t), e^{-rt}(1-u_i(t))\right)$; if player j rejects, the game continues. Let $I_i^g = \{t : g(t) = i\}$, and assume that $I_i^g \cap [0,t]$ is finite for all t and that I_i^g is infinite. The announcement phase is correspondingly modified so that player 1 announces a posture $\gamma : I_i^g \to [0,1]$, and if player 1 becomes committed to posture γ (which continues to occur with probability ε), she demands $\gamma(t)$ at time t and rejects all of player 2's demands. I refer to the function g as a *discrete-time bargaining game*.

I now define convergence to continuous time. This definition is very similar to that of Abreu and Gul (2000), as is the above model of discrete-time bargaining and the corresponding notation.

Definition 11 *A sequence of discrete-time bargaining games* $\{g_n\}$ converges to continuous time *if for all* $\Delta > 0$*, there exists* N *such that for all* $n \geq N$*,* $t \in \mathbb{R}_+$*, and* $i \in \{1, 2\}$*,* $I_i^{g_n} \cap [t, t + \Delta] \neq \emptyset$*.*

The maxmin payoff and posture in a discrete-time bargaining game are defined exactly as in Section 2. Let $u_1^{*,g}$ be player 1's maxmin payoff in discrete-time bargaining game g, and let $u_1^{*,g}(\gamma)$ be player 1's maxmin payoff given posture γ in g. The independence-of-procedures result states that, for any sequence of discrete-time bargaining games converging to continuous time, the corresponding sequence of maxmin payoffs $\{u_1^{*,g_n}\}$ converges to u_1^*, and any corresponding sequence of postures $\{\gamma^{g_n}\}$ such that $u_1^{*,g_n}(\gamma^{g_n}) \to u_1^*$ "converges" to γ^*, where u_1^* and γ^* are the maxmin payoff and posture identified in Theorem 1. The nature of the convergence of the sequence $\{\gamma^{g_n}\}$ to γ^* is slightly delicate. For example, there may be (infinitely many) times $t \in \mathbb{R}_+$ such that $\lim_{n\to\infty} \gamma^{g_n}(t)$ exists and is greater than $\gamma^*(t)$, because these demands may be "non-serious" (in that they are followed immediately by lower demands).[36] Thus, rather than stating the convergence in terms of $\{\gamma^{g_n}\}$ and γ^*, I state it in terms of the corresponding continuation values of player 2, which are the economically more important variables. Formally, given a posture γ^{g_n} in discrete-time bargaining game g^n, let

$$v^{g_n}(t) \equiv \max_{\tau \geq t : \tau \in I_1^{g_n}} e^{-r(\tau - t)} \left(1 - \gamma^{g_n}(\tau)\right).$$

Let $v^*(t) = \max\{1 - e^{rt}/(1 - \log \varepsilon), 0\}$, the continuation value corresponding to γ^* in the continuous-time model of Section 2. The independence-of-procedures result is as follows:

Proposition 6 *Let* $\{g_n\}$ *be a sequence of discrete-time bargaining games converging to continuous time. Then* $u_1^{*,g_n} \to u_1^*$*, and if* $\{\gamma^{g_n}\}$ *is a sequence of postures with* γ^{g_n} *a posture in* g_n *and* $u_1^{*,g_n}(\gamma^{g_n}) \to u_1^*$*, then* $v^{g_n}(t) \to v^*(t)$ *for all* $t \in \mathbb{R}_+$*.*

The key fact behind the proof of Proposition 6 is that for any sequence of discrete-time postures $\{\gamma^{g_n}\}$ converging to some continuous-time posture γ, $\lim_{n\to\infty} u_1^{*,g_n}(\gamma^{g_n}) = \lim_{n\to\infty} u_1^*(\gamma^{g_n})$ (where $u_1^*(\gamma^{g_n})$ is the maxmin payoff given a natural embedding of γ^{g_n} in

[36] The reason this complication does not arise in Theorem 1 is that the assumption that γ is continuous at non-integer times rules out "non-serious" demands.

continuous time, defined formally in the proof). This fact is proved by constructing a belief that is similar to the γ^{g_n}-offsetting belief in each discrete-time game g_n and then showing that these beliefs converge to the γ-offsetting belief in the limiting continuous-time game.

Proof of Proposition 6. Observe that a posture γ in discrete-time bargaining game g induces a "continuous-time posture" $\hat{\gamma}$ (i.e., a map from $\mathbb{R}_+ \to [0,1]$) according to $\hat{\gamma}(t) = \gamma(\min\{\tau \geq t : \tau \in I_i^g\})$. That is, $\hat{\gamma}$'s time-t demand is simply γ's next demand in g. I henceforth refer to a posture γ in g as also being a continuous-time posture, with the understanding that I mean the posture $\hat{\gamma}$ defined above.

However, γ may *not* be a posture in the continuous-time bargaining game of Section 2, because it may be discontinuous at a non-integer time. To avoid this problem, I now introduce a modified version of the continuous-time bargaining game of Section 2. Formally, let the continuous-time bargaining game g^{cts} be defined as in Section 2, with the following modifications: Most importantly, omit the requirement that player i's demand path $u_i^t : [t, t+1) \to [0,1]$ (which is still chosen at integer times t) is continuous. Second, specify that the payoffs if player i accepts player j's offer at date $(t,-1)$ are $(e^{-rt}(1-\liminf_{\tau\uparrow t} u_j(\tau)), e^{-rt}\liminf_{\tau\uparrow t} u_j(\tau))$ (because $\lim_{\tau\uparrow t} u_j(\tau)$ may now fail to exist). Third, add a fourth date, $(t,2)$ to each instant of time t. At date $(t,2)$, each player i announces *accept* or *reject*, and, if player i accepts player j's offer at date $(t,2)$, the game ends with payoffs $(e^{-rt}(1-\liminf_{\tau\downarrow t} u_j(\tau)), e^{-rt}\liminf_{\tau\downarrow t} u_j(\tau))$. Adding the date $(t,2)$ ensures that each player has a well-defined best-response to her belief, even though $u_j(t)$ may now fail to be right-continuous. One can check that the analysis of Sections 3 and 4, including Lemmas 1 through 3 and Theorem 1, continue to apply to the game g^{cts}, with the exception that in g^{cts} the maxmin posture γ^* is not in fact unique; however, every maxmin posture corresponds to the continuation value function v^* (by the same argument as in Step 1 of the proof of Theorem 1).[37] Because of this, for the remainder of the proof I slightly abuse notation by writing $u_1^*(\gamma)$ for player 1's maxmin payoff given posture γ *in the game*

[37] The reason I did not use the game g^{cts} in Sections 3 and 4 is that it is difficult to interpret the assumption that player i can accept the demand $\liminf_{\tau\downarrow t} u_j(\tau)$ at time t, since the demand $u_j(\tau)$ has not yet been made at time t for all $\tau > t$. Thus, I view the game g^{cts} as a technical construct for analyzing the limit of discrete-time games, and not as an appealing model of continuous-time bargaining in its own right.

g^{cts}, rather than in the model of Section 2. Importantly, $u_1^*(\gamma)$ equals player 1's maxmin payoff given γ in both g^{cts} and in the model of Section 2 when γ is a posture in the model of Section 2, but $u_1^*(\gamma)$ is well-defined for all $\gamma : \mathbb{R}_+ \to [0,1]$. Similarly, I write $u_1^*(v)$ for player 1's maxmin payoff given continuation value function $v : \mathbb{R}_+ \to [0,1]$. This is well-defined because $u_1^*(\gamma) = \min_{t \le T} e^{-rt}\underline{\gamma}(t)$ by Lemma 3, T depends on γ only through v (by Lemma 1), and it can be easily verified that $\min_{t\le T} e^{-rt}\underline{\gamma}(t) = \min_{t \le T} e^{-rt}(1 - v(t))$ (and thus depends on γ only through v). A similar argument, which I omit, implies that one may write $u_1^{*,g_n}(v^{g_n})$ for player 1's maxmin payoff given continuation value function v^{g_n} in discrete-time bargaining game g_n.

I now establish two lemmas, from which Proposition 6 follows. Their proofs require some additional notation. Let Σ_i^g be the set of player i's strategies in g^{cts} with the property that player i's demand only changes at times $t \in I_i^g$, player i only accepts player j's offer at times $t \in I_j^g$, and player i's action at time t only depends on past play at times $\tau \in I_i^g \cup I_j^g$. One can equivalently view Σ_i^g as player i's strategy set in g itself. Thus, any belief π_2 in g may also be viewed as a belief in g^{cts} (with $\operatorname{supp}(\pi_2) \subseteq \Sigma_i^g$).

Lemma 6 *For any sequence of discrete-time bargaining games converging to continuous time, $\{g_n\}$, there exists a sequence of postures $\{\gamma^{g_n\prime}\}$ with $\gamma^{g_n\prime}$ a posture in g_n and $\lim_{n\to\infty} u_1^{*,g_n}(\gamma^{g_n\prime}) \ge u_1^*$.*

Proof. Let $\gamma^{g_n\prime}$ be given by $\gamma^{g_n\prime}(t) = \left(\frac{n}{n+1}\right)\gamma^*\left(\max\{\tau \le t : \tau \in I_1^{g_n}\}\right)$ for all $t \in \mathbb{R}_+$, with the convention that $\max\{\tau \le t : \tau \in I_1^{g_n}\} \equiv 0$ if the set $\{\tau \le t : \tau \in I_1^{g_n}\}$ is empty. I first claim that $\lim_{n\to\infty} u_1^*(\gamma^{g_n\prime}) \ge u_1^*$.[38] To show this, I first establish that $\tilde{T}(\gamma^{g_n\prime}) \le \min\{\tau > T^1 : \tau \in I_1^{g_n}\}$ for all n, where T^1 is defined as in the proof of Theorem 1. Since γ^* (and thus $\gamma^{g_n\prime}$) are non-decreasing, $\sup_{\tau \ge t} e^{-r(\tau - t)}(1 - \gamma^{g_n\prime}(\tau)) = 1 - \gamma^{g_n\prime}(t)$. Therefore, by Lemma 1, $\tilde{T}(\gamma^{g_n\prime})$ satisfies

$$\exp\left(-\int_0^{\tilde{T}(\gamma^{g_n\prime})} \frac{r\left(\frac{n+1}{n} - \gamma^*\left(\max\{\tau \le t : \tau \in I_1^{g_n}\}\right)\right)}{\gamma^*\left(\max\{\tau \le t : \tau \in I_1^{g_n}\}\right)} dt\right) \prod_{t \in I_1^{g_n} \cap \left[0, \tilde{T}(\gamma^{g_n\prime})\right)} \frac{\gamma^*\left(\max\{\tau < t : \tau \in I_1^{g_n}\}\right)}{\gamma^*(t)} \ge \varepsilon. \tag{24}$$

[38] Theorem 1 implies that $\lim_{n\to\infty} u_1^*(\gamma^{g_n\prime}) \le u_1^*$, so this inequality must hold with equality. But only the inequality is needed for the proof.

Now

$$\exp\left(-\int_0^{\tilde{T}(\gamma^{g_n\prime})}\frac{r\left(\frac{n+1}{n}-\gamma^*\left(\max\left\{\tau\le t:\tau\in I_1^{g_n}\right\}\right)\right)}{\gamma^*\left(\max\left\{\tau\le t:\tau\in I_1^{g_n}\right\}\right)}dt\right)\prod_{t\in I_1^{g_n}\cap\left[0,\tilde{T}(\gamma^{g_n\prime})\right)}\frac{\gamma^*\left(\max\left\{\tau<t:\tau\in I_1^{g_n}\right\}\right)}{\gamma^*(t)}$$

$$\le \exp\left(-\int_0^{\tilde{T}(\gamma^{g_n\prime})}\frac{r(1-\gamma^*(t))}{\gamma^*(t)}dt\right)\frac{\gamma^*(0)}{\gamma^*\left(\max\left\{\tau<\tilde{T}\left(\gamma^{g_n\prime}\right):\tau\in I_1^{g_n}\right\}\right)}$$

$$\le \exp\left(-\int_0^{\max\left\{\tau<\tilde{T}(\gamma^{g_n\prime}):\tau\in I_1^{g_n}\right\}}\frac{r(1-\gamma^*(t))+\gamma^{*\prime}(t)}{\gamma^*(t)}dt\right). \qquad (25)$$

Observe that if $\tilde{T}\left(\gamma^{g_n\prime}\right)>\min\left\{\tau>T^1:\tau\in I_1^{g_n}\right\}$ then $\max\left\{\tau<\tilde{T}\left(\gamma^{g_n\prime}\right):\tau\in I_1^{g_n}\right\}>T^1$, and therefore (25) is less than ε, which contradicts (24). Hence, $\tilde{T}\left(\gamma^{g_n\prime}\right)\le\min\left\{\tau>T^1:\tau\in I_1^{g_n}\right\}$ for all n. In addition, $\gamma^{g_n\prime}(t)$ is non-decreasing and $\gamma^{g_n\prime}(t)<1$ for all t, which implies that $T\left(\gamma^{g_n\prime}\right)=\tilde{T}\left(\gamma^{g_n\prime}\right)$. Therefore, by Lemma 3, $u_1^*\left(\gamma^{g_n\prime}\right)=\min_{t\le\tilde{T}(\gamma^{g_n\prime})}e^{-rt}\gamma^{g_n\prime}(t)$. Since $\tilde{T}\left(\gamma^{g_n\prime}\right)\le\min\left\{\tau>T^1:\tau\in I_1^{g_n}\right\}$ for all n, and $\{g_n\}$ converges to continuous time, $\lim_{n\to\infty}\tilde{T}\left(\gamma^{g_n\prime}\right)\le T^1$. In addition, $\lim_{n\to\infty}\sup_{t\in\mathbb{R}_+}\left|\gamma^{g_n\prime}(t)-\gamma^*(t)\right|=0$, so it follows that

$$\lim_{n\to\infty}u_1^*\left(\gamma^{g_n\prime}\right)=\lim_{n\to\infty}\min_{t\le\tilde{T}(\gamma^{g_n\prime})}e^{-rt}\gamma^{g_n\prime}(t)\ge\lim_{n\to\infty}\min_{t\le\tilde{T}(\gamma^{g_n\prime})}e^{-rt}\gamma^*(t)\ge\min_{t\le T^1}e^{-rt}\gamma^*(t)=u_1^*.$$

Next, I claim that $u_1^{*,g_n}\left(\gamma^{g_n}\right)\ge u_1^*\left(\gamma^{g_n}\right)$ for *any* posture γ^{g_n} in discrete-time bargaining game g^n. To see this, note that if $\operatorname{supp}(\pi_2)\subseteq\Sigma_1^{g_n}$ and $\sigma_2\in\Sigma_2^{*,g_n}(\pi_2)$, then $\sigma_2\in\Sigma_2^*(\pi_2)$ as well (i.e., there is no benefit to responding to a strategy in $\Delta\left(\Sigma_1^{g_n}\right)$ with a strategy outside of $\Sigma_2^{g_n}$). Therefore, if $\pi_1\in\Pi_1^{\gamma^{g_n},g_n}$ (i.e., if π_1 is consistent with knowledge of rationality in g_n), then $\pi_1\in\Pi_1^{\gamma^{g_n},g^{cts}}$; that is, $\Pi_1^{\gamma^{g_n},g_n}\subseteq\Pi_1^{\gamma^{g_n},g^{cts}}$. Now

$$\begin{aligned}
u_1^{*,g_n}\left(\gamma^{g_n}\right) &= \sup_{\sigma_1\in\Sigma_1^{g_n}}\inf_{\pi_1\in\Pi_1^{\gamma^{g_n},g_n}}u_1(\sigma_1,\pi_1)\\
&\ge \sup_{\sigma_1\in\Sigma_1^{g_n}}\inf_{\pi_1\in\Pi_1^{\gamma^{g_n},g^{cts}}}u_1(\sigma_1,\pi_1)\\
&= u_1\left(\gamma^{g_n},\sigma^{\gamma^{g_n}}\right)\\
&= u_1^*\left(\gamma^{g_n}\right),
\end{aligned}$$

where $\sigma^{\gamma^{g_n}}$ is as in Definition 5, and the second line follows because $\Pi_1^{\gamma^{g_n},g_n}\subseteq\Pi_1^{\gamma^{g_n},g^{cts}}$; the third line follows because $u_1\left(\gamma^{g_n},\sigma^{\gamma^{g_n}}\right)=\sup_{\sigma_1\in\Sigma_1^{g^{cts}}}\inf_{\pi_1\in\Pi_1^{\gamma^{g_n},g^{cts}}}u_1(\sigma_1,\pi_1)$ by Lemma 3, and $\gamma^{g_n}\in\Sigma_1^{g_n}\subseteq\Sigma_1^{g^{cts}}$; and the fourth line follows by Lemma 3.

Combining the above claims, it follows that $\lim_{n\to\infty}u_1^{*,g_n}\left(\gamma^{g_n\prime}\right)\ge\lim_{n\to\infty}u_1^*\left(\gamma^{g_n\prime}\right)\ge u_1^*$.

■

Lemma 7 *For any sequence of discrete-time bargaining games converging to continuous time, $\{g_n\}$ and any sequence of functions $\{v^{g_n}\}$ such that v^{g_n} is a continuation value function in g_n and $\lim_{n\to\infty} v^{g_n}(t)$ exists for all $t \in \mathbb{R}_+$, it follows that $\lim_{n\to\infty} u_1^{*,g_n}(v^{g_n})$ exists and equals $\lim_{n\to\infty} u_1^*(v^{g_n})$.*

Proof. Fix a sequence of continuation value functions $\{v^{g_n}\}$ (with v^{g_n} a continuation value function in discrete-time game g_n) converging pointwise to some function $v : \mathbb{R}_+ \to [0,1]$. I have already shown that $u_1^{*,g_n}(\gamma^{g_n}) \geq u_1^*(\gamma^{g_n})$ for any posture γ^{g_n} in game g_n, or equivalently $u_1^{*,g_n}(v^{g_n}) \geq u_1^*(v^{g_n})$. This immediately implies that $\lim_{n\to\infty} u_1^{*,g_n}(v^{g_n}) \geq \limsup_{n\to\infty} u_1^*(v^{g_n})$ for every convergent subsequence of $\{u_1^{*,g_n}(v^{g_n})\}$. Hence, I must show that $\lim_{n\to\infty} u_1^{*,g_n}(v^{g_n}) \leq \liminf_{n\to\infty} u_1^*(v^{g_n})$ for every convergent subsequence of $\{u_1^{*,g_n}(v^{g_n})\}$. I establish this inequality by assuming that there exists $\eta > 0$ such that $\lim_{n\to\infty} u_1^{*,g_n}(v^{g_n}) > \liminf_{n\to\infty} u_1^*(v^{g_n}) + \eta$ for some convergent subsequence of $\{u_1^{*,g_n}(v^{g_n})\}$ and then deriving a contradiction. The approach is to first define analogs of the continuous-time γ-offsetting belief and the time $\tilde{T}$ (defined in Section 3.2) for game g_n, denoted $\pi_2^n \in \Sigma_2^{g_n}$ and $\tilde{T}^n \in \mathbb{R}_+$, and then show that $\tilde{T}^n \to \tilde{T}$.

I must introduce some additional notation before defining the belief π_2^n. Let $t_{g_n}^{next}(i) = \min\{\tau > t : \tau \in I_i^{g_n}\}$ be the time of player i's next demand at t. Given continuation value function v^{g_n} and any corresponding posture γ^{g_n}, let $\widetilde{\gamma^{g_n}}^n$ be defined as follows: First, $\widetilde{\gamma^{g_n}}^n$ demands $\widetilde{\gamma^{g_n}}^n(h^t) = \gamma^{g_n}(h^t)$ for all $t \in I_1^{g_n}$. Second, $\widetilde{\gamma^{g_n}}^n$ accepts player 2's demand at time $t \in I_2^{g_n}$ with probability

$$\hat{p}^n(t) \equiv \min\left\{\frac{p^n(t)}{\chi^n(t)}, 1\right\},$$

where

$$p^n(t) \equiv \max_{\substack{\tau<t: \\ \tau \in I_1^{g_n}, \tau_{g_n}^{next}(2)=t}} \frac{e^{r(t-\tau)} v^{g_n}(\tau) - v^{g_n}(t)}{1 - v^{g_n}(t)}$$

if $\left\{\tau < t : \tau \in I_1^{g_n}, \tau_{g_n}^{next}(2) = t\right\}$ is non-empty and $v^{g_n}(\tau) < 1$ for all time τ in this set, and $p^n(t) \equiv 0$ otherwise; and

$$\chi^n(t) \equiv \max\left\{\frac{\Pi_{\tau<t:\tau\in I_2^{g_n}}(1 - p^n(\tau)) - \varepsilon}{\Pi_{\tau<t:\tau\in I_2^{g_n}}(1 - p^n(\tau))}, 0\right\}.$$

Let $\tilde{T}^n$ be the supremum over times t at which $\chi^n\left(t_{g_n}^{next}(2)\right)\hat{p}^n\left(t_{g_n}^{next}(2)\right) = p^n\left(t_{g_n}^{next}(2)\right)$, and let

$$T^n \equiv \sup \underset{\substack{t \geq \tilde{T}^n: \\ t \in I_1^{g_n}}}{\operatorname{argmax}} e^{-rt} v^{g_n}(t).$$

By an argument similar to the proof of Lemma 2, if $\gamma^{g_n}(t) < \eta$ for some $t \leq T^n$, then there exists a belief $\pi_2 \in \Delta\left(\Sigma_1^{g_n}\right)$ and strategy $\sigma_2 \in \Sigma_2^{g_n}$ such that $\pi_2\left(\gamma^{g_n}\right) \geq \varepsilon$, $\sigma_2 \in \Sigma_2^{*,g_n}(\pi_2)$, and the demand $\gamma^{g_n}(t)$ is accepted under strategy profile (γ^{g_n}, σ_2). In particular, $u_1^{g_n}(\gamma^{g_n}, \sigma_2) < \eta$. Thus, by the hypothesis that $\lim_{n\to\infty} u_1^{*,g_n}(v^{g_n}) > \liminf_{n\to\infty} u_1^*(v^{g_n}) + \eta$, there must exist $N > 0$ such that $\gamma^{g_n}(t) \geq \eta$ for all $t \leq T^n$ and all $n > N$, and hence $v^{g_n}(t) \leq 1 - \eta$ for all $t \leq T^n$ and all $n > N$.

Let π_2^n assign probability ε to γ^{g_n} and probability $1-\varepsilon$ to $\widetilde{\gamma^{g_n}}^n$, and fix $\sigma_2^n \in \Sigma_2^{*,g_n}(\pi_2^n)$ with the property that σ_2^n always demands 1 and rejects player 1's demand at any history at which player 1 has deviated from γ^{g_n} (which is possible because π_2^n assigns probability 0 to such histories, except for terminal histories), as well as at any history at which player 2 is indifferent between accepting and rejecting player 1's demand under belief π_2^n. Note that γ^{g_n} is a best-response to σ_2^n in g_n. This implies that $u_1^{*,g_n}(\gamma^{g_n}) \leq u_1^{g_n}(\gamma^{g_n}, \sigma_2^n)$ for all n. Thus, to show that $\lim_{n\to\infty} u_1^{*,g_n}(v^{g_n}) \leq \liminf_{n\to\infty} u_1^*(v^{g_n}) + \eta$ (the desired contradiction), it suffices to show that $\lim_{n\to\infty} u_1^{g_n}(\gamma^{g_n}, \sigma_2^n) \leq \liminf_{n\to\infty} u_1^*(v^{g_n}) + \eta$.

Observe that $p^n(t)$ satisfies

$$\exp(-r(t-\tau))\left(p^n(t)(1) + (1-p^n(t))\, v^{g_n}(t)\right) \geq v^{g_n}(\tau)$$

for all $\tau \leq t$ such that $\tau \in I_1^{g_n}$ and $\tau_{g_n}^{next}(2) = t$. Hence, it is optimal for player 2 to reject player 1's demand γ at any time τ at which $\chi^n\left(\tau_{g_n}^{next}(2)\right)\hat{p}^n\left(\tau_{g_n}^{next}(2)\right) = p^n\left(\tau_{g_n}^{next}(2)\right)$ (under belief π_2^n). Therefore, $u_1^{g_n}(\gamma^{g_n}, \sigma_2^n) \leq \min_{t \leq T^n} e^{-rt}(1 - v^{g_n}(t))$. Now $u_1^*(v^{g_n}) = \min_{t \leq T(v^{g_n})} e^{-rt}(1 - v^{g_n}(t))$, and $\lim_{n\to\infty} \tilde{T}(v^{g_n}) = \tilde{T}(v)$. Hence, showing that $\lim_{n\to\infty} \tilde{T}^n = \tilde{T}(v) \equiv \tilde{T}$ would imply that $\lim_{n\to\infty} u_1^{g_n}(\gamma^{g_n}, \sigma_2^n) \leq \liminf_{n\to\infty} u_1^*(v^{g_n})$, yielding the desired contradiction. The remainder of the proof shows that $\lim_{n\to\infty} \tilde{T}^n = \tilde{T}$.

To see that $\lim_{n\to\infty} \tilde{T}^n = \tilde{T}$, first fix $t_0 \leq \tilde{T}$ and note that for all $\delta > 0$ there exists $N' > 0$ such that, for all $t \leq t_0$ and all $n \geq N'$, if $g_n(t) = 2$ then $\min\left\{\tau \leq t : \tau \in I_1^{g_n}, \tau_{g_n}^{next}(2) = t\right\} \geq t - \delta$ (if this set is non-empty). Next, since both $e^{-rt} v(t)$ and $e^{-rt} v^{g_n}(t)$ are non-increasing (as is easily checked) and $v^{g_n}(t) \to v(t)$ for all $t \in \mathbb{R}_+$, it follows that for all $\delta' > 0$ there exists $\delta > 0$ such that $t \leq t_0$ and $\tau \in [t-\delta, t]$ implies that $\left|e^{r(\tau - t)} v^{g_n}(\tau) - v(t,-1)\right| < \delta'$. Since $1 - v(t) \geq \eta$ for all $t \leq \tilde{T}$, combining these observations and letting S be the (countable) set of discontinuity points of $v(t)$, for all $\delta' > 0$ there exists N'' such that if $t = s^{next}_{g_n}(2)$ for some $s \in S \cap [0, t_0]$, and $n \geq N''$, then $\left|p^n(t) - \frac{v(t,-1)-v(t)}{1-v(t)}\right| < \delta'$.[39] Hence,

$$\lim_{n\to\infty} \prod_{s\in S\cap[0,t_0]} \left(1 - p^n\left(s^{next}_{g_n}(2)\right)\right) = \prod_{s\in S\cap[0,t_0]} \left(1 - p(s)\right) \tag{26}$$

for all $t_0 \leq \tilde{T}$, where p is as in Section 3.2.

Finally, I establish that whenever v is continuous on an interval $[t_0, t_\infty]$ with $t_\infty \leq \tilde{T}$,

$$\lim_{n\to\infty} \prod_{t\in I_2^{g_n}\cap[t_0,t_\infty]} \left(1 - p^n(t)\right) = \exp\left(-\int_{t_0}^{t_\infty} \frac{rv(t) - v'(t)}{1 - v(t)} dt\right) = \exp\left(-\int_{t_0}^{t_\infty} \lambda(t)\, dt\right), \tag{27}$$

where λ is as in Section 3.2. I will prove this fact by showing that the limit as $n \to \infty$ of a first-order approximation of the logarithm of $\prod_{t\in I_2^{g_n}\cap[t_0,t_\infty]} (1 - p^n(t))$ equals $-\int_{t_0}^{t_\infty} \frac{rv(t)-v'(t)}{1-v(t)}$.

Let $\left\{t_{1,g_n}, t_{2,g_n}, \ldots, t_{K(n),g_n}\right\} = \{t \in [t_0, t_\infty] : p^n(t) > 0\}$, with $t_{k,g_n} < t_{k+1,g_n}$ for all $k \in \{1, \ldots, K(n) - 1\}$ and all $n \in \mathbb{N}$, and let $t_{0,g_n} = \max\left\{\tau : \tau \in I_1^{g_n}, \tau^{next}_{g_n}(2) = t_{1,g_n}\right\}$. Note that $K(n)$ is finite because $I_2^{g_n} \cap [t_0, t_\infty]$ is finite, and that in addition $t^{next}_{k,g_n}(1) < t_{k+1,g_n}$ for all k (where $t^{next}_{k,g_n}(1) \equiv t^{next}_{k,g_n,g_n}(1)$ to avoid redundant notation). Furthermore, since $e^{-r\tau} v^{g_n}(\tau)$ is non-increasing,

$$t^{next}_{k,g_n}(1) \in \underset{\substack{\tau < t_{k+1,g_n}: \\ \tau\in I_1^{g_n}, \tau^{next}_{g_n}(2) = t_{k+1,g_n}}}{\operatorname{argmax}} e^{r\left(t_{k+1,g_n} - \tau\right)} v^{g_n}(\tau)$$

for all $k \in \{0, 1, \ldots, K(n) - 1\}$. Therefore,

$$\begin{aligned}
\prod_{k=1}^{K(n)} \left(1 - p^n(t_{k,g_n})\right) &= \prod_{k=1}^{K(n)} \min_{\substack{\tau < t_{k,g_n}: \\ \tau\in I_1^{g_n}, \tau^{next}_{g_n}(2) = t_{k,g_n}}} \frac{1 - e^{r\left(t_{k,g_n} - \tau\right)} v^{g_n}(\tau)}{1 - v^{g_n}(t_{k,g_n})} \\
&= \prod_{k=1}^{K(n)} \min_{\substack{\tau < t_{k,g_n}: \\ \tau\in I_1^{g_n}, \tau^{next}_{g_n}(2) = t_{k,g_n}}} \frac{1 - e^{r\left(t_{k,g_n} - \tau\right)} v^{g_n}(\tau)}{1 - e^{-r\left(t^{next}_{k,g_n}(1) - t_k\right)} v^{g_n}\left(t^{next}_{k,g_n}(1)\right)} \\
&= \left(\prod_{k=1}^{K(n)-1} \frac{1 - e^{r\left(t_{k+1,g_n} - t^{next}_{k,g_n}(1)\right)} v^{g_n}\left(t^{next}_{k,g_n}(1)\right)}{1 - e^{-r\left(t^{next}_{k,g_n}(1) - t_{k,g_n}\right)} v^{g_n}\left(t^{next}_{k,g_n}(1)\right)}\right) \frac{1 - e^{r\left(t_{1,g_n} - t^{next}_{0,g_n}(1)\right)} v^{g_n}\left(t^{next}_{0,g_n}(1)\right)}{1 - e^{-r\left(t^{next}_{K(n),g_n}(1) - t_{K(n),g_n}\right)} v^{g_n}\left(t^{next}_{K(n),g_n}(1)\right)}
\end{aligned} \tag{28}$$

[39] S is countable because $e^{-rt} v(t)$ is non-increasing, and monotone functions have at most countably many discontinuities. Unlike in Section 3, S need not be a subset of $\mathbb{N}$ here.

Next, taking a first-order Taylor approximation of $\log\left(1-e^{rx}v^{g_n}(t)\right)$ at $x=0$ yields

$$\log\left(1-e^{rx}v^{g_n}(t)\right)=\log\left(1-v^{g_n}(t)\right)-\frac{rxv^{g_n}(t)}{1-v^{g_n}(t)}+O\left(x^2\right).$$

Therefore, a first-order approximation of the logarithm of (28) equals

$$\left(\sum_{k=1}^{K(n)-1}-\left(t_{k+1,g_n}-t_{k,g_n}\right)\frac{rv^{g_n}\left(t_{k,g_n}^{next}(1)\right)}{1-v^{g_n}\left(t_{k,g_n}^{next}(1)\right)}\right)$$
$$+\log\left(1-e^{r\left(t_{1,g_n}-t_{0,g_n}^{next}(1)\right)}v^{g_n}\left(t_{0,g_n}^{next}(1)\right)\right)-\log\left(1-e^{-r\left(t_{K(n),g_n}^{next}(1)-t_{K(n),g_n}\right)}v^{g_n}\left(t_{K(n),g_n}^{next}(1)\right)\right).$$

I now show that

$$\lim_{n\to\infty}\sum_{k=1}^{K(n)-1}-\left(t_{k+1,g_n}-t_{k,g_n}\right)\frac{rv^{g_n}\left(t_{k,g_n}^{next}(1)\right)}{1-v^{g_n}\left(t_{k,g_n}^{next}(1)\right)}=-\int_{t_0}^{t_\infty}\frac{rv(t)}{1-v(t)}dt \tag{29}$$

and

$$\begin{aligned}&\lim_{n\to\infty}\left(\log\left(1-e^{r\left(t_{1,g_n}-t_{0,g_n}^{next}(1)\right)}v^{g_n}\left(t_{0,g_n}^{next}(1)\right)\right)-\log\left(1-e^{-r\left(t_{K(n),g_n}^{next}(1)-t_{K(n),g_n}\right)}v^{g_n}\left(t_{K(n),g_n}^{next}(1)\right)\right)\right)\\ =\;&\int_{t_0}^{t_\infty}\frac{v'(t)}{1-v(t)}dt,\end{aligned} \tag{30}$$

which completes the proof of (27). Equation (30) is immediate, because, since v is continuous on $[t_0,t_\infty]$, both the left- and right-hand sides equal

$$\log\left(1-v(t_0)\right)-\log\left(1-v(t_\infty)\right).$$

To establish (29), let

$$f^n(t)\equiv\exp\left(-r\left(\frac{1+\eta}{\eta}\right)t\right)\frac{rv^{g_n}(t)}{1-v^{g_n}(t)}$$

and let

$$f(t)\equiv\exp\left(-r\left(\frac{1+\eta}{\eta}\right)t\right)\frac{rv(t)}{1-v(t)}.$$

For all $n>N$, it can be verified that both $f^n(t)$ and $f(t)$ are non-increasing on the interval $[t_0,t_\infty]$, using the facts that $e^{-rt}v^{g_n}(t)$ and $e^{-rt}v(t)$ are non-increasing and that $v^{g_n}(t)\leq 1-\eta$ for all $n>N$ and $t\leq t_\infty\leq\tilde{T}$. Fix $\zeta>0$ and $m\in\mathbb{N}$. Because $v^{g_n}(t)\to v(t)$ for all $t\in\mathbb{R}_+$, there exists $N'''\geq N$ such that, for all $n>N'''$, $|f^n(t)-f(t)|<\zeta$ for all t in the set

$$\left\{t_0,\frac{(m-1)t_0+t_\infty}{m},\frac{(m-2)t_0+2t_\infty}{m},\ldots,t_\infty\right\}.$$

Since both f^n and f are non-increasing on $[t_0, t_\infty]$, this implies that

$$|f^n(t) - f(t)| < \zeta + \max_{k \in \{1,\dots,K(n)-1\}} \left(f\left(\frac{(m-k)\,t_0 + kt_\infty}{m}\right) - f\left(\frac{(m-k-1)\,t_0 + (k+1)\,t_\infty}{m}\right)\right)$$

for all $t \in [t_0, t_\infty]$. Since f is continuous on $[t_0, t_\infty]$, taking $m \to \infty$ implies that $|f^n(t) - f(t)| < 2\zeta$ for all $t \in [t_0, t_\infty]$. Therefore $\left|\frac{rv^{g_n}(t)}{1-v^{g_n}(t)} - \frac{rv(t)}{1-v(t)}\right| \leq 2\zeta \exp\left(r\left(\frac{1+\eta}{\eta}\right)t_\infty\right)$ for all $t \in [t_0, t_\infty]$. Hence,

$$\begin{aligned}
\lim_{n\to\infty} \sum_{k=1}^{K(n)-1} -\left(t_{k+1,g_n} - t_{k,g_n}\right) \frac{rv^{g_n}\left(t^{next}_{k,g_n}(1)\right)}{1 - v^{g_n}\left(t^{next}_{k,g_n}(1)\right)} &= \lim_{n\to\infty} \sum_{k=1}^{K(n)-1} -\left(t_{k+1,g_n} - t_{k,g_n}\right) \frac{rv\left(t^{next}_{k,g_n}(1)\right)}{1 - v\left(t^{next}_{k,g_n}(1)\right)} \\
&= \lim_{n\to\infty} \sum_{k=1}^{K(n)-1} -\left(t_{k+1,g_n} - t_{k,g_n}\right) \frac{rv\left(t_{k,g_n}\right)}{1 - v\left(t_{k,g_n}\right)} \\
&= -\int_{t_0}^{t_\infty} \frac{rv(t)}{1-v(t)}\,dt,
\end{aligned}$$

where the first equality follows because $\sum_{k=1}^{K(n)-1}\left(t_{k+1,g_n} - t_{k,g_n}\right) \leq t_\infty - t_0$ for all $n \in \mathbb{N}$, the second follows because $t^{next}_{k,g_n}(1) \in [t_{k,g_n}, t_{k+1,g_n}]$ and v is continuous on $[t_0, t_\infty]$, and the third follows by definition of the (Riemann) integral.

Combining (26) and (27), it follows that

$$\lim_{n\to\infty} \prod_{s \in I_2^{g_n} \cap [0,t]} (1 - p^n(s)) = \exp\left(-\int_0^t \lambda(s)\,ds\right) \prod_{s \in S \cap [0,t]} (1 - p(s))$$

for all $t \leq \tilde{T}$. This implies that $\lim_{n\to\infty} \tilde{T}^n = \tilde{T}$, completing the proof of the lemma. ■

I now complete the proof of Proposition 6.

Let $\{g_n\}$ be a sequence of discrete-time bargaining games converging to continuous time. Recall that $u_1^{*,g_n} = \sup_{\gamma^{g_n}} u_1^{*,g_n}(\gamma^{g_n})$. Thus, there exists a sequence of postures $\{\gamma^{g_n}\}$, with γ^{g_n} a posture in g_n, such that $\lim_{n\to\infty} |u_1^{*,g_n} - u_1^{*,g_n}(\gamma^{g_n})| = 0$. Let $\{v^{g_n}\}$ be the corresponding sequence of continuation value functions. Because $e^{-rt}v^{g_n}(t)$ is non-increasing and the space of monotone functions from $\mathbb{R}_+ \to [0,1]$ is sequentially compact (by Helly's selection theorem or footnote 26), this sequence has a convergent subsequence $\{v^{g_k}\}$ converging to some v on $\mathbb{R}_+$.

I claim that $v = v^*$. Toward a contradiction, suppose not. Since v^* is the unique maxmin continuation value function in g^{cts}, there exists $\eta > 0$ such that $u_1^* > \lim_{k\to\infty} u_1^*(v^{g_k}) + \eta$. By Lemma 7, $\lim_{k\to\infty} u_1^{*,g_k}(v^{g_k}) = \lim_{k\to\infty} u_1^*(v^{g_k})$. Finally, by Lemma 6, there exists an alternative sequence of postures $\{\gamma^{g_k\prime}\}$ such that $\lim_{k\to\infty} u_1^{*,g_k}(\gamma^{g_k\prime}) \geq u_1^*$. Combining these observations implies that there exists $K > 0$ such that, for all $k \geq K$,

$$u_1^{*,g_k}(\gamma^{g_k\prime}) > u_1^* - \eta/3 > u_1^*(v^{g_k}) + 2\eta/3 > u_1^{*,g_k}(v^{g_k}) + \eta/3,$$

which contradicts the fact that $\lim_{k\to\infty} |u_1^{*,g_k} - u_1^{*,g_k}(\gamma^{g_k})| = 0$. Therefore, $v = v^*$. In addition, since this argument applies to any convergent subsequence of $\{v^{g_n}\}$, and every subsequence of $\{v^{g_n}\}$ has a convergent sub-subsequence, this implies that $v^{g_n} \to v^*$ pointwise.

A similar contradiction argument shows that $\lim_{k\to\infty} u_1^{*,g_k}(v^{g_k}) = u_1^*$, for any convergent subsequence $\{v^{g_k}\} \subseteq \{v^{g_n}\}$. Since $\lim_{k\to\infty} |u_1^{*,g_k} - u_1^{*,g_k}(\gamma^{g_k})| = 0$, it follows that $u_1^{*,g_k} \to u_1^*$. And, since this argument applies to any convergent subsequence of $\{v^{g_n}\}$, this implies that $u_1^{*,g_n} \to u_1^*$. ■