SAN FRANCISCO STATE UNIVERSITY

MATHEMATICS DEPARTMENT

http://math.sfsu.edu/

Working Paper No. 06-001

Bargaining in Continuous Time:

Holding Out for Concession and for Information

by

Jean-Pierre P. Langlois

January 13, 2006

Bargaining in Continuous Time:
Holding Out for Concession and for Information

by

Jean-Pierre P. Langlois*†

January 13, 2006

Abstract

This paper studies strategic bargaining in continuous time in order to allow the timings of offers and acceptance to be true decision variables. The main technical innovation is to decouple these timing decisions in order to study how each side can use delaying tactics to wear the other down into concession. There are two main restrictive assumptions: (1) players cannot use "explosive" strategies that could generate (with positive probability) infinitely many moves in finite time; and (2) there is no arbitration for the simultaneous acceptance of mismatched offers. It is simply ruled invalid. I construct and characterize subgame perfect, perfect Bayesian, and sequential equilibria in the complete and incomplete information cases respectively. Despite the stationarity of the game parameters, and with or without an inside cost game option, I exhibit a multiplicity of equilibria that are generically inefficient whether information is complete or not. I conclude that the uniqueness and efficiency results of the standard Rubinstein model are jeopardized when the temporal monopoly constraint is lifted.

JEL codes: C70, C78

Keywords: Bellman equation, continuous time, countervailing strategy, finite state machine, survival function.

Bargaining in Continuous Time: Holding Out for Concession and for Information

1 Introduction

In the strategic approach initiated by Rubinstein (1982) and Stahl (1972) bargaining is pictured as a discounted repeated game and its solution involves an appropriate perfect equilibrium concept. This has evident conceptual advantages over Nash's original axiomatic-based solution: first, the outcome of bargaining is the result of strategic interaction between selfish and rational actors rather than that of a process akin to arbitration. Second, bargaining is often just a facet of a richer strategic interaction that may involve the power to hurt and may not always be resolved as fairly as arbitration would aim for. And third, some temporal aspects that are overlooked in the axiomatic approach are explicit in the strategic one.

However, the great benefits of the strategic approach to bargaining are mitigated by a nagging simplification of the temporal picture: in most cases the players move sequentially at fixed intervals of time, a modeling feature called "temporal monopoly". Since the value of the solution depends explicitly on the length of the interval, this temporal picture is clearly not neutral. But there is a deeper and more troublesome issue: could the very nature of the solution depend on the assumption that moves are sequential and at fixed intervals of time, therefore not subject to the players' decisions? Put another way, would theoretical predictions be substantially different should the timings of offers and acceptance become explicit decision variables? One reason this may be so is that the uniqueness and efficiency of the Rubinstein solution can be traced directly to the cost of waiting an additional *fixed* period and to the resulting incentive for immediate acceptance. So, if the delays can be manipulated strategically can uniqueness and efficiency evaporate altogether?

Interest in these issues is nothing new. Indeed, efforts at making the timing of offers or acceptance true decision variables can be found in Cramton (1992), Perry and Reny

(1993), and Sakovics (1993). More recently, Wang (2000), Smith and Stacchetti (2003), and Sandholm (2002) also frame these decisions within the time continuum. All these efforts vary in underlying assumptions and methodology. From the assumption viewpoint, for instance, Cramton allows the players to delay their offers but maintains the constraint of alternating turns. And although it is not unreasonable to expect players to take turns in a bargaining game, there is a subtle difference between allowing them and *constraining* them to do so. Smith and Stacchetti require that an offer be instantly accepted or rejected thus precluding that it remains on the table for some time. But leaving an offer on the table while attempting to wear down an opponent into acceptance may be an important strategic option that should not be precluded, especially if the players have the power to hurt each other.

From a methodology viewpoint, for instance, one may pick a unit period and obtain equilibria, then let the period length approach zero and examine the limit, as in Cramton. But the limit of the solutions of the discrete case may not yield all solutions of the limit model, as happens in the present paper. One may also frame the objective in the time continuum but restrict attention to some classes of strategies, as in Smith and Stacchetti. This may be enough for existence results but it leaves out a broader description of what behavior can be involved in equilibrium. Finally, the solution concept itself can vary from Nash, to perfect, to sequential, involving complete or incomplete information as in Wang (2000). It can also use the weaker "epsilon-equilibrium" form that has been applied more generally to games of timing, as in Laraki, Solan and Vieille (2005).

The results, primarily existence, uniqueness, and efficiency of equilibria, of course vary with the assumptions and the methodology. But, more importantly, one must wonder whether such results are only a reflection of the constraints on behavior. If one, for instance, lets the period length approach zero in the original Rubinstein model, the limit solution inherits the efficiency of the discrete one and is arguably a solution of a "continuous-time" model. But, as I show in this paper, it is not be the only one and most others are inefficient. Moreover, multiplicity or inefficiency can occur under temporal monopoly with various additional assumption: Busch and Wen (1995), for instance, show how an inside cost game can jeopardize both. And Coles and Muthoo (2003) reach similar conclusions with non-stationary game parameters.

Ideally, one would want a model as free as possible of any timing constraint if the issue is to understand whether uniqueness and efficiency are only the result of "artificial" timing assumptions. But relaxing such constraints does not come without costs. Indeed, a host of new issues arises from relaxing timing constraints: first, since offers can be accepted at any time, as long as they are still on the table, the timing of offers and acceptance must be independent decision variables. And, although offers should not be limited to the "take-it-or-leave-it" kind, such offers should still be allowed. However, if an offer is to be accepted or rejected "instantly" its recipient must make that move *immediately following* the offer. But there is no single turn following an offer in the time-continuum. Second, in a framework where an offer can remain on the table for some time, players must be allowed to withdraw it instantly from the table if it is not taken. It thus becomes necessary to clarify the relative timing of acceptance, rejection and withdrawal. Third, since alternating turns are no longer the rule, simultaneous moves must be allowed. But how should they be resolved if they are not compatible? For instance, what if the two sides simultaneously accept offers that do not add up to the whole pie?

I argue that these issues affect the very definition of the player's very decision space and I propose to extend it from simple offer, acceptance, and cost control moves to more complex logical instructions such as "take it or leave it" or "if offered at least this much then accept instantly." But if logical instructions can be the players' very decisions and if they can be simultaneous, then they must be scrutinized for consistency and legality. Conceptually, it is as if each side hands logical instructions to an agent who instantly meets the other side's agent to decide whether any deal has actually been reached according to law and logic.

As the game unfolds, various events such as offers and cost adjustments take place and become part of the history upon which the players plan their future moves. A strategy therefore expresses how the logical instructions at any time result from the prior history of the game. And it translates into a map from the time continuum, through the two sides' logical instructions, into offer, cost, and acceptance state variables. Because an infinite frequency of players' moves is ruled out by assumption, this map has isolated discontinuities. This leads to a natural definition of the players' objectives as an expectation of acceptance by one side or the other, appropriately discounted for delays, minus a possible expected cost of waiting. In the time continuum this expectation is expressed as a Lebesgue-Stieltjes integral up to the time-infinity.

Within this picture of time, action space, and objective function, I construct subgame perfect equilibria (SPE) under perfect information, and perfect-bayesian (PBE) and sequential equilibria under incomplete information. They involve "holding-out" behavior akin to that predicted in war of attrition games: each side waits strategically, applying the pressure of costs and delays, in the hope that the other side will make a better offer or will be first to concede. The optimal waiting time strategy is probabilistic and given by an exponential "survival" function with a parameter that depends on the game data and the current offer and cost conditions. But a deterministic timing structure resembling the Rubinstein framework is also possible if the players choose to adopt it. The difference is that respecting the agreed upon timing scheme becomes part of the players' strategies instead of being the fiat of the theorist. These findings suggest that Rubinstein's uniqueness and efficiency results do not survive the relaxation of its temporal monopoly assumption.

Most of the results extend to the incomplete information case with finitely many types on each side, although some results need to be qualified. For instance, a wide range of offers is permissible in holding out SPEs whereas only extreme ones are possible in PBEs as long as complete information has not yet been reached. Likewise, only a Nash equilibrium of the cost game, if it is non-trivial, can be used in a PBE. This suggests that incomplete information yields ungenerous and hurtful choices in equilibrium, and this also entails inefficient delays. However, I show that complete information is reached in finite time.

2 Model and Assumptions

In the standard framework there are two players denoted i and j with symmetric roles but possibly different characteristics. Player i makes offers $x_i \in [0,1]$ to player j, receives share $(1-x_i)$ if j accepts x_i, receives share x_j if he accept j's offer, and derives utility $u_i(x)$ if x is his share of the accepted offer. Since more is assumed better, utility u_i is a strictly increasing function of its argument. The bargaining game may also involve an inside cost game option with strategic variable $z_i \in \mathcal{C}_i$ (for i) and bounded cost flow $-c_i(z_i, z_j) \leq 0$. Acceptance by either side ends the game and the occurrence of any further costs.

2.1 Decisions as Logical Instructions

An offer by a player must stipulate its magnitude and the time when it is made. In the discrete alternating turn framework there is no ambiguity as to when that offer can be accepted or rejected since its recipient is next to move at a well specified time. But in continuous time it is less clear. One approach adopted by Smith and Stacchetti (2003) is to assume that an offer must be "instantly" accepted or rejected. This is appealing because it requires no arbitrary built-in delays and bypasses any discussion of whether a rational player might prefer to postpone acceptance. But in many bargaining situations an offer can remain available, at least for some time, perhaps with a given deadline, when it is not immediately accepted or explicitly rejected. So, requiring an immediate response for all offers is clearly restrictive. Instead, I let the players decide for themselves whether an offer they make is to be instantly dealt with or can remain available for some time. But this requires a formulation that goes beyond simply stating an offer magnitude. It implies that statements such as "here is my offer, take it or leave it *now*" be explicitly part of the player's action space.

Formally, an offer is a pair (x_i, t) of magnitude $x_i \in [0,1]$ made by player i to his counterpart j at date $t \in [0,\infty)$. An offer of magnitude zero is legal and can serve several purposes such as withdrawing a prior offer. The notation $(\emptyset_i, t)$ means that no offer is tendered by i at time t and will be used to define the outcome of strategy. All by itself (x_i, t)

doesn't say for how long x_i will remain available and whether it comes with any contingency such as "take-it-or-leave-it." It is merely an element in a universe of discourse $\mathcal{U}$ for the logical instructions that will form the players' decision spaces. So, i's counterpart j should be able to accept i's offer whenever she likes as long as it is not withdrawn. Acceptance by j is a non-empty signal denoted a_j together with the time t when it is issued, that is the pair (a_j, t). Rejection, denoted $\neg a_j$, is also a non-empty signal signifying that j makes the decision not to accept at that time, a move that must be distinguished from the lack of any acceptance decision that will be denoted $\emptyset_j$. Neither rejection nor lack of acceptance by j affects i's offer. But acceptance presumes that something is indeed available and will in fact require a careful definition of player i's "current offer." Acceptance does not spell out exactly *what* is accepted and is only, like offers, an element in the universe of discourse $\mathcal{U}$. Finally, a cost decision by i is a pair (z_i, t) made up of a cost control $z_i \in \mathcal{C}_i$ and the time t when it is implemented, and it is an element of $\mathcal{U}$. As the game begins at time $t = 0$ the two sides' offers and cost controls are assumed to be in a preexisting default state. For instance, each side may be offering zero and be currently inflicting costs on the other. A null offer, acceptance, or cost control $(\emptyset_i, t)$ has no effect on the current state of the corresponding variable which therefore remains in its prior state.

So, an offer (x_i, t) followed at some time $s > t$ by an acceptance (a_j, s) should unambiguously end the game with share x_i for j and $(1 - x_i)$ for i at time s. But what about *immediate* acceptance? In discrete time immediate acceptance would occur at *turn* s immediately following t. But in the time continuum there is no single time s following time t. Immediate acceptance should thus mean that x_i is considered accepted for any $s > t$ *arbitrarily close to* t. This seems only possible if the acceptance *decision* a_j is issued at precisely the same time ($s = t$) the accepted offer is tendered. But then, how does j know that an offer she hasn't yet seen will be satisfactory when she accepts it?

One can think of several avenues to deal with this issue, for instance: (1) establish some delay Δs between successive decisions. One drawback is that the value of Δs is set by the theorist and is therefore arbitrary. Indeed, it is typically made to approach zero in a limit argument that is an indirect way of adopting another option; (2) base acceptance on beliefs at information sets representing simultaneous decisions. But it seems awkward to presume that players in a bargaining game would accept *unconditionally* an offer they haven't yet seen on the basis of their beliefs about what it will turn out to be. It should be preferable to wait even a tiny instant in order to reach certainty; (3) weaken the solution concept to that of ϵ-equilibrium. Strategies can then involve waiting a tiny instant Δs after an offer is made before accepting it. The resulting cost can be made less than an arbitrary ϵ; (4) allow players to issue decisions that are explicitly conditional on what their counterpart's simultaneous decision turns out to be. In this last approach, each side instructs an *agent* to tender or accept offers on its behalf subject to contingencies. Whenever a decision is issued by either side, the two sides' agents meet and execute jointly their instructions according to law and logic.[1] Indeed, each side can issue "standing instructions" and only modify them from time to time as they see fit. This is the approach investigated in this paper.

In this conceptual framework all offer, acceptance, and cost moves can be issued at any time by players i and j and can be made conditional on each others' prior or simultaneous instructions. Player i's instructions at time t therefore can use j's own simultaneous instructions as "input" to produce i's offer, acceptance, and cost moves as output. Formally, $\mathcal{Y}_i = ([0,1] \cup \{\emptyset_i\}) \times \{a_i, \neg a_i, \emptyset_i\} \times (\mathcal{C}_i \cup \{\emptyset_i\})$ denotes the set of all possible (offer, acceptance, and cost) moves by i, any of which can be null. And if Ω_j denotes the set of all possible instructions by j, then i's instruction α_i^t at time t is a map

$$\alpha_i^t : \Omega_j \to \mathcal{Y}_i \qquad (1)$$

that, given any possible simultaneous instruction $\beta_j^t \in \Omega_j$, yields i's potential new offer x_i, acceptance a_i, and cost control decision z_i at time t, any of which may be the null $\emptyset_i$. Player i's own "action space" is then the set Ω_i of all such maps α_i^t.

A more practical description for α_i^t is that it is a "finite state machine" (i.e., a computer program) that uses j's own instructions' to issue i's moves. It is therefore convenient to use pseudocode to describe these instructions. For instance

$$\alpha_i^t := \text{ offer } x_i = \tfrac{1}{2} \tag{2}$$

means that i's decision is only to issue the unconditional offer $(\frac{1}{2})_i$ at time t regardless of j's own instructions (β_j^t is not used in (2)). The symbol " $:=$ " represents an assignment of α_i^t to the logical instructions on the right-hand side. If j wishes to instantly accept this new offer, provided it is made, she could issue the instruction

$$\alpha_j^t := \text{ input } \beta_i^t \text{ (}i\text{'s instructions); if } (\beta_i^t = \alpha_i^t) \text{ then (accept) } a_j \tag{3}$$

with α_i^t given by (2). If i does issue α_i^t the offer $(\frac{1}{2})_i$ is made and j's agent issues the decision a_j, resulting in the instant acceptance of the offer. But perhaps i only wishes to make a "take-it-or-leave-it offer. In that case, he could issue instead

$$\alpha_i^t := \text{ input } \beta_j^t\text{; if } (\beta_j^t = \alpha_j^t) \text{ then offer } x_i = \tfrac{1}{2} \tag{4}$$

If the argument α_i^t in (3) is the pseudocode in (4) and if the argument α_j^t in (4) is the pseudocode in (3), then $(\frac{1}{2})_i$ is tendered and accepted instantly with a_j. But if j deviates even slightly from (3) then i makes no offer. In fact, even if i's old standing offer has magnitude $\frac{1}{2}$, j does not accepts it according to α_j^t since it requires specifically α_i^t.

So, formulae (3) and (4) are highly restrictive since the desired outcome only occurs for a very specific pair of instructions. But these could easily be broadened. Instead of requiring i to issue just α_i^t in (4) j could draw a list $\mathcal{S}_i$ of acceptable instructions by i and replace the test $(\beta_i^t = \alpha_i^t)$ by $(\beta_i^t \in \mathcal{S}_i)$ in (4). For instance, let

$$\mathcal{S}_i = \{\alpha_i^t = \text{"input } \beta_j^t\text{; if } (\beta_j^t = \alpha_j^t) \text{ then offer } x_i\text{", such that } x_i \geq \tfrac{1}{2}\}$$

This set of instructions includes (4) and j can modify (3) into

$$\alpha_j^t := \text{ input } \beta_i^t\text{; if } (\beta_i^t \in \mathcal{S}_i) \text{ then accept } a_j \tag{5}$$

Of course, i can proceed similarly so that broader instructions on both sides can allow a successful offer and acceptance of the take-it-or-leave-it offer. In the most general case, α_j^t

can use β_i^t as a "subroutine" and implement the resulting program rather than simply read i's script and decide whether it belongs to an acceptable set. Such an approach allows more flexibility and more realism in the interaction between the two sides' respective agents but requires good programming skills since $\alpha_j^t(\alpha_i^t)$ must not "loop" and fail to yield an unambiguous outcome in $\mathcal{Y}_i$, a development that is here assumed away.

In the most general game theoretic framework, moves can be probabilistic. And finite-state machines can indeed involve probabilistic transitions between states. However, an ambiguity could arise if acceptance by one side is made contingent on a probabilistic offer by the other, since both sides' instructions must be executed simultaneously. To avoid this situation it will be assumed that any such probability is resolved by an appropriate random device *prior* to execution of the two sides' instructions. In essence, this means that several *deterministic* α_j^t's could be issued at time t, each with a probability.

When instructions result in no offer, acceptance or cost control moves, all are deemed null (denoted $\emptyset_i$ for i). It will be convenient to sometime write "offer $\emptyset_i$" to indicate that i issues no offer. But i need not issue any instructions such as α_i^t. In that case $\alpha_i^t = \emptyset_i$ by convention and this results in a null output. This is not the same as issuing "standing instructions". In that case i continuously issues the same α_i^t knowing that it will result in a non-null output only when j issues the desirable instructions. In all cases, the two sides' instructions at time t (whether null or not) read

$$\alpha^t = \left(\alpha_i^t, \alpha_j^t\right) \in \Omega = \Omega_i \times \Omega_j$$

and they result in an output denoted

$$\mathcal{L}\alpha^t \in \mathcal{Y} = \mathcal{Y}_i \times \mathcal{Y}_j$$

made up of all three possible moves by each side, any of which may be null. The "legal" operator $\mathcal{L}$ that will be described in Assumption A2 below will be necessary to resolve some potential conflicts that could arise between the various moves. For instance, the two sides could simultaneously make and accept "over-generous" offers that would add up to more than the pie. These will be ruled illegal and will result in null offers and acceptance in $\mathcal{L}\alpha^t$. It will be convenient to write $\mathcal{L}\alpha^t = \emptyset$ when the output is made up only of null moves. $\mathcal{L}\alpha^t$ is called an "event" if at least one of its moves is non-null, formally if $\mathcal{L}\alpha^t \neq \emptyset$. $\mathcal{L}\alpha^t$ is called an "i-event" if one non-null move is in $\mathcal{Y}_i$. The notation

$\alpha_i^t(\alpha_j^t) \mapsto$ offer x_i (or accept a_i, or set cost control to z_i)

will indicate that $\alpha_i^t(\alpha_j^t)$ yields the offer x_i (or acceptance or cost control) by i *before* the legal operator $\mathcal{L}$ is applied. "$\mathcal{L}\alpha^t \mapsto$ offer x_i" instead means that the offer is indeed legal.

In this perspective, the instructions contained in α^t are executed and resolved instantly as if time is suspended while the two sides' respective agents interact. One could instead assume that the interaction takes time but that the players cannot modify their instructions while it lasts and this would not change the results. If, however, the time it takes depends on the complexity of the instructions, then the results might be affected.

2.2 Strategies

$h^t = \{\alpha^s \in \Omega : s < t\}$ denotes the history *strictly* prior to time t (note that h^t gives time t). And $\mathcal{H}$ will denote the set of all possible histories. A *pure* strategy for player i is a map $\psi_i : \mathcal{H} \to \Omega_i$ that, to any prior history h^t, associates that player's contemporary instruction α_i^t. But it will be necessary to broaden the range of strategies in order to allow decisions that involve probabilities in both *what* is decided and *when* it is decided. So, if $\mathcal{R}_i$ is the set of random variables available to player i, a *randomized* strategy is a map

$$\psi_i : \mathcal{H} \times \mathcal{R}_i \to \Omega_i$$

that, to any prior history h^t, associates that player's contemporary decision α_i^t with some probability $\mathbb{P}(\alpha_i^t)$. Finitely many possible α_i^t's can be involved in such a lottery at time t, and α_i^t can be null if the lottery is about whether to issue any instruction at all.

Any random device used by i is assumed privately monitored and independent of j's devices. However, the *parameters* of the random devices used by either side can be public knowledge. There is a subtle but important difference between randomization in the *timing* of an instruction and a randomization among several deterministic instruction at a given time: in the first case, i is monitoring a private device to decide whether to issue an instruction. In the second case, i is issuing one of several non-null instructions α_i^t. For example

$$\psi_i(h_t) = \begin{cases} a_i \text{ (accept)} & \text{if } t = \theta \\ \emptyset_i & \text{otherwise} \end{cases} \tag{6a}$$

means that i accepts j's current offer at time θ. If θ is a random variable (with value true or false) in $\mathcal{R}_i$, say with exponential distribution $\mathbb{P}(\theta \geq s) = e^{-\lambda s}$ of parameter $\lambda > 0$, an i-event can only occur when the variable takes the value true. Consider instead, with fixed θ

$$\psi_i(h_t) = \begin{cases} \alpha_i^\theta & \text{if } t = \theta \\ \emptyset_i & \text{otherwise} \end{cases} \tag{6b}$$

with $\alpha_i^\theta := a_i$ (accept), with probability $\frac{1}{2}$; $\alpha_i^\theta := \neg a_i$ (reject), with probability $\frac{1}{2}$.

In that case, an event does occur at θ whether the acceptance is issued or not. Of course, other statements concerning i's offers and cost controls are generally involved in ψ_i.

As events are generated by strategies from a prior history, they become part of a longer history.[2] But in the time continuum, there is the potential for events becoming so frequent that there could be infinitely many within a finite history. Such a development is called an "event explosion." In that case, at least one side must be generating the explosion. So, a strategy ψ_i (for i) is called "explosive" if there exists a strategy ψ_j (for j) such that the pair (ψ_i, ψ_j) can generate *with positive probability* a history h^t that contains infinitely many i-events. The results of this paper require few significantly restrictive assumptions on the strategies that players can adopt. The main restriction is

***Assumption A1* (No Explosion):** A player cannot use any explosive strategy.

In practice, this means that allowable strategies should include provisions that preclude, or give probability zero, to the possibility of an event explosion. For instance, a strategy that makes only one decision at only one time, like ψ_i in (6), cannot generate any i-event explosion. But, more generally, mild statements can be included in the formulation of any strategy to ensure Assumption A1. For example, i's strategy can include the provision that it issue any decision $\alpha_i^t \in \Omega_i$ only with probability

$$\mathbb{P}\big(\alpha_i^t\big) = min\big\{1, inf\{\mathcal{O}(t-s) | \alpha^s \in h^t, \mathcal{L}\alpha^s \neq \emptyset\big\} \quad (7)$$

where $\mathcal{O}$ is such that $\lim_{\epsilon \to 0} \mathcal{O}(\epsilon) = 0$. In plain English, the probability that i issues the decision α_i^t goes to zero with the time elapsed since the last event.[3] When such a clause is used in ψ_i for the issuance of α_i^t Assumption A1 is fulfilled (see Proposition 1 in appendix).

Discrete time models of bargaining usually do not require explicit initial conditions. Conditions such as an "asking price" announced before bargaining begins may have consequences in defining the bargaining space, initial beliefs or expectations, but they are not indispensable. In the continuous time case it is not clear when and by whom an initial offer or a cost decision will be made, although it is necessary to define unambiguously the current state of the variables of the game, including each side's *current* offer. So, one can assume that, at the start of the game, each side i is offering 0 and has its cost control variable set to some default value z_i^d (if the cost game is non-trivial). Although not generated from prior history, these initial conditions form a default event $\alpha^d = (0_i, \neg a_i, z_i^d, 0_j, \neg a_j, z_j^d)$ that is included at the beginning of all histories h^t. One can now make

***Definition 1*:** Player i's current offer at time $t \geq 0$ is $\xi_i(t) = x_i$ if

$\exists s \leq t :$ ($\mathcal{L}\alpha^s \mapsto$ offer x_i) and ($\neg\exists s < \tau \leq t : \mathcal{L}\alpha^\tau \mapsto$ offer y_i);

$\xi_i(t) = 0$ otherwise.

Since there are finitely many strategy-generated events in h^t with probability one, according to A1, $\xi_i(t)$ is well defined as the most recent non-null offer. One defines similarly $\zeta_i(t) = z_i$ as the current cost control of player i. Both ξ_i and ζ_i are clearly constant by pieces in time.

Just like the offer and cost aspects of strategy yield functions $\xi_i(t)$ and $\zeta_i(t)$ giving the states of the offer and cost control variables at time t, the state of acceptance can be described by a function of time. But in continuous time the possibility of simultaneous acceptance raises new issues. Actual acceptance, just like offers and cost controls, does not just result from a one-sided strategy. Instead, actual moves result from the resolution of both sides' instructions into a possible event $\mathcal{L}\alpha^t$. It is therefore necessary to describe precisely the legal operator $\mathcal{L}$ mentioned before. It involves the cases when the players make over-generous offers and when the two sides simultaneously attempt to accept mismatched offers. In either case, arbitration would be needed. But since arbitration schemes are *arbitrary* let us make

Assumption A2 **(No Arbitration):**

$\mathcal{L}\alpha^t = \left(\alpha_i^t(\alpha_j^t), \alpha_j^t(\alpha_i^t)\right)$ with the following three exceptions:

• Simultaneous over-generous offers are null:

if ($\alpha_i^t(\alpha_j^t) \mapsto$ offer x_i) and ($\alpha_j^t(\alpha_i^t)) \mapsto$ offer x_j) and ($x_i + x_j > 1$)

then $\mathcal{L}\alpha^t \mapsto$ (no offer $\emptyset_i$) and (no offer $\emptyset_j$);

• A unilateral over-generous offer is null:[4]

if ($\alpha_i^t(\alpha_j^t) \mapsto$ offer x_i) and ($\alpha_j^t(\alpha_i^t) \mapsto$ no offer $\emptyset_j$) and ($x_i + \xi_j(t) > 1$)

then $\mathcal{L}\alpha^t \mapsto$ no offer $\emptyset_i$;

• Simultaneously acceptance of mismatched offers yields rejection by both:

if ($\alpha_i^t(\alpha_j^t) \mapsto$ accept a_i) and ($\alpha_j^t(\alpha_i^t) \mapsto$ accept a_j) and ($\xi_i(t) + \xi_j(t) \neq 1$)

then $\mathcal{L}\alpha^t \mapsto$ (reject $\neg a_i$) and (reject $\neg a_j$);

These legal restrictions are implemented (as operator $\mathcal{L}$) once both sides' instructions have been resolved into $\alpha_i^t(\alpha_j^t)$ and $\alpha_j^t(\alpha_i^t)$ respectively. It is therefore fruitless to include contingencies in either player's instructions about nullified decisions. Instead, each side should carefully write its instructions to avoid seeing any of them nullified. The third condition in A2 precludes any contentious simultaneous acceptance. Therefore, an event that results in an acceptance ends the game. There can be no further events and the history is final.

Now let $\mathcal{A}_i(I) =$ "$\exists\tau \in I: \ \alpha_i^\tau(\alpha_j^\tau)) \mapsto a_i$" mean that i issues an acceptance (legal or not) within the interval I and let $\mathcal{A}_i^{\mathcal{L}}(I) =$ "$\exists\tau \in I: \ \mathcal{L}\alpha^\tau \mapsto a_i$" mean that i legally accepts within that interval I. The probability

$$F_i(t) = \mathbb{P}\big(\mathcal{A}_i^{\mathcal{L}}[0,t)\big)$$

is a distribution function which, together with the similar F_j, describes the probabilistic state of acceptance of the game.[5] However, discrete probabilities of acceptance at specific event times are also allowed and result in right-hand side discontinuities of F_i at such times. For example, suppose that j plans to never issue any acceptance, so that $F_j \equiv 0$, and that i issues an acceptance a_i at a specific time θ. Then

$$F_i(t) = \begin{cases} 0 & \text{if } t \leq \theta \\ 1 & \text{if } t > \theta \end{cases}$$

But, instead of being a specific time θ can be a random variable, for instance with exponential distribution of parameter λ. In that case $F_i(t) = 1 - e^{-\lambda t}$.

A discontinuity $\Delta F_i(\theta) = F_i(\theta^+) - F_i(\theta) > 0$ is interpreted as a discrete probability of acceptance by i at time θ. And, although variations of F_i represent random elements in acceptance, they have different interpretations: a continuous variation is about the timing of an acceptance while the second is about the decision to accept at the time of the discontinuity. Indeed, a discrete probability of acceptance $\Delta F_i(\theta)$ at a specific θ is an i-event observable by j,[6] while a continuous $F_i(t)$ is only about whether and when an i-event may take place.

2.3 Player Objective

Any history h^t, through the induced strategy profile $\Psi^t = (\psi_i^t, \psi_j^t)$, determines further offer, cost, and acceptance behavior that translate at future time $(t+\tau)$ into current offer, cost and acceptance probability. The superscript "t" will be used to indicate developments according to this induced Ψ^t, so that $\xi_i^t(\tau) = \xi_i(t+\tau|\Psi^t)$. A "path" at time t is a map

$$\sigma^t : [0,\infty) \to [0,1] \times \mathcal{C}_i \times [0,1] \times \mathcal{C}_j$$

where $\sigma^t(\tau) = \big(\xi_i^t(\tau), \zeta_i^t(\tau), \xi_j^t(\tau), \zeta_j^t(\tau)\big)$ is the state of the game after further delay τ according to Ψ^t. A pure strategy profile defines a unique path while a randomized strategy yields a random path. But, according to Proposition 2 in appendix, the No-Explosion assumption A1 implies that, for any given time t, there is *probability one* that for some $\theta > 0$ no event occurs within the interval $(t, t+\theta)$. Such an interval will be called a "leg" of path σ^t so that a path can be viewed as a sequence of legs separated by events. Within each leg σ^t is thus constant. A leg is called "deterministic" if its length θ is not a random variable. A "subleg" is a leg that is strictly contained in another. And a leg is "maximal" if it is not a subleg. If a legal acceptance occurs at time t then (t,∞) is a trivial leg where no cost is incurred and the two sides enjoy indefinitely the accepted offer.

If $U_i(\sigma^t(\tau))d\tau$ denotes the utility enjoyed by i within the infinitesimal period at time $(t+\tau)$,[7] it is discounted by player-specific factor $e^{-r_i\tau}$ for delay τ, where $r_i > 0$ is player i's "impatience." i's expected payoff for path σ^t viewed from time t is therefore

$$E_i(\sigma^t) = \int_0^\infty r_i e^{-r_i\tau} U_i(\sigma^t(\tau))d\tau \qquad (8)$$

More generally, the magnitude and the timing of all events on σ^t can be randomized and $E_i(\sigma^t)$ should be replaced by an expectation. However, all existence results of this paper can be obtained with only the *timing* of acceptance being randomized. Other kinds of randomization are therefore left as illustrations as in section 3.3. If σ^t yields acceptance after additional time s, and if x is i's share of the accepted offer, then (8) becomes

$$E_i(\sigma^t) = \int_0^s - r_i e^{-r_i\tau} c_i\big(\zeta_i^t(\tau), \zeta_j^t(\tau)\big)d\tau + e^{-r_i s}u_i(x)$$

where $\zeta_i^t(\tau)$ denotes the evolution of i's cost control on path σ^t. So, if $\xi_j^t(s)$ denotes the evolution of j's offer on path σ^t and if $F_i^t(s) = F_i(t+s|\Psi^t)$ is the probability that i successfully accepts by time $(t+s)$ according to Ψ^t, the expected payoff of i's accepting j's offer at some point in time is given by, after integrating the first term by parts[8]

$$\int_{[0,\infty)} \Big(\int_0^s - r_i e^{-r_i\tau} c_i\big(\zeta_i^t(\tau), \zeta_j^t(\tau)\big)d\tau + e^{-r_i s}u_i(\xi_j^t(s)) \Big) dF_i^t(s) \qquad (9)$$

$$= \int_{[0,\infty)} e^{-r_i s}\Big((F_i^t(s) - F_i^t(\infty)) r_i c_i\big(\zeta_i^t(s), \zeta_j^t(s)\big) ds + u_i(\xi_j^t(s)) dF_i^t(s) \Big)$$

Formula (9) is well defined for any path as a Lebesgues-Stieltjes integral, since $e^{-r_i s} u_i(\xi_j^t(s))$ is clearly measurable. But it requires a careful interpretation at common discontinuities of ξ_j^t and F_i^t (see Section 3.1 and Proposition 3 in appendix). A symmetric formula involving the evolution of i's offer $\xi_i^t(s)$ and the distribution function $F_j^t(s)$ gives the expected payoff that j accepts. And the expected payoff that neither side ever accepts is

$$- \Phi^t(\infty) \int_{[0,\infty)} e^{-r_i s} r_i c_i\big(\zeta_i^t(s), \zeta_j^t(s)\big) ds$$

where $\Phi^t(s) = 1 - F_i^t(s) - F_j^t(s)$ is the "joint survival" function that expresses how long the game lasts before one side or the other accepts. Adding up all three cases yields

$$E_i(\sigma^t) = \int_{[0,\infty)} e^{-r_i s} dG_{ij}^t(s) \tag{10}$$

with $dG_{ij}^t(s) = u_i\big(\xi_j^t(s)\big) dF_i^t(s) + u_i\big(1 - \xi_i^t(s)\big) dF_j^t(s) - c_i\big(\zeta_i^t(s), \zeta_j^t(s)\big) \Phi^t(s) ds$

The infinitesimal $dG_{ij}^t(s)$ is simply the expected value that either side accepts minus the expected cost, all discounted appropriately in the integral. In the incomplete information case there will be two sets of types $\mathcal{I}$ and $\mathcal{J}$ and the players will entertain beliefs $b_j(t) \geq 0$ such that $\sum_{j \in \mathcal{J}} b_j(t) = 1$ (and similarly for $\mathcal{I}$). Objective (10) will then generalize to

$$E_i(\sigma^t) = \sum_{j \in \mathcal{J}} b_j(t) \int_{[0,\infty)} e^{-r_i s} dG_{ij}^t(s) \tag{11}$$

for type $i \in \mathcal{I}$. Of course, the beliefs will be appropriately updated by Bayes' law when discussing perfect Bayesian equilibria.[9]

2.4 A First Example

Let the players' strategies be such that $\xi_i^t(s) \equiv x_i^*$ (constant) and $\zeta_i^t(s) \equiv z_i^*$ (constant), and symmetrically for j (with $x_i^* + x_j^* < 1$). Further assume that i issues an acceptance at random time θ with exponential density of parameter $\mu_i > 0$, and symmetrically for j. After history h^t, since the probability of acceptance by i within the infinitesimal ds at time s is $\mu_i e^{-\mu_i s} ds$ and the (independent) probability that j has not yet accepted is $e^{-\mu_j s}$, $dF_i^t(s) = \mu_i e^{-(\mu_i + \mu_j)s} ds$. Moreover $\Phi^t(s) = e^{-(\mu_i + \mu_j)s}$. In this case

$$E_i(\sigma^t) = \int_0^\infty e^{-(r_i+\mu_i+\mu_j)s}\big(\mu_i u_i(x_j^*) + \mu_j u_i(1-x_i^*) - c_i(z_i^*, z_j^*)\big)ds$$
$$= \tfrac{\mu_i u_i(x_j^*)+\mu_j u_i(1-x_i^*)-c_i(z_i^*,z_j^*)}{r_i+\mu_i+\mu_j}$$

Now, let us define the important quantity

$$\Lambda_i(x_i, x_j, z_i, z_j) = \tfrac{c_i(z_i,z_j)+r_i u_i(x_j)}{u_i(1-x_i)-u_i(x_j)} \tag{12}$$

and let us assume that j chooses $\mu_j = \Lambda_i(x_i^*, x_j^*, z_i^*, z_j^*)$, then $E_i(\sigma^t) \equiv u_i(x_j^*)$. In that case, i is made indifferent between waiting with the given exponential distribution or accepting at any time. A symmetric argument also holds for j given the symmetric μ_i. This example will be generalized in the next section to obtain a subgame perfect equilibrium.

3. Bargaining with Perfect Information

Much of strategic bargaining theory was initially developed with perfect information. It is therefore reasonable to first extend this case to continuous time. It is also easier to discuss the characterization of optimal acceptance behavior without the added complexity arising from continuous Bayesian updating. The methodology developed in this section exploits the definition of path in dynamic programing fashion: in order to optimize their objective, the players can concentrate on optimizing their behavior within each leg knowing that the optimization problem they will face at the start of the next leg will be structurally identical. The resulting optimality condition is often called a Bellman equation. In the next section, this dynamic programming approach is made precise through Lemma 1.

3.1 Bellman Equation and Survival Functions

Because there can be no left-discontinuity in F_i^t or F_j^t one can always write[10]

$$E_i(\sigma^t) = \int_{[0,\theta)} e^{-r_i s} dG_{ij}^t(s) + \int_{[\theta,\infty)} e^{-r_i s} dG_{ij}^t(s) \tag{13}$$

And since F_i^t, and Φ^t describe conditional probabilities one can also write

$$dF_i^t(\theta+s) = \Phi^t(\theta) dF_i^{t+\theta}(s) \tag{14a}$$

and $$\Phi^t(\theta+s) = \Phi^t(\theta)\Phi^{t+\theta}(s) \tag{14b}$$

and thus $dG_{ij}^t(\theta+s) = \Phi^t(\theta) dG_{ij}^{t+\theta}(s)$. It follows from changing variables in (13) that

$$E_i(\sigma^t) = \int_{[0,\theta)} e^{-r_i s} dG_{ij}^t(s) + e^{-r_i\theta}\Phi^t(\theta)E_i(\sigma^{t+\theta}) \tag{15}$$

Formula (15) has important consequences. In particular, the so-called Bellman Equation obtains when writing that the maximum for the left-hand side of (15) is reached when maximizing the sum on the right-hand side.[11] A subgame perfect equilibrium (SPE) is a strategy profile that forms a Nash equilibrium in all subgames. But a subgame is the bargaining game itself beginning at time t with some prior history h^t. So, a SPE must be such that $E_i(\sigma^t)$, or its expected value, is maximized by i's induced strategy in response to j's induced strategy. In order to obtain such SPEs one can proceed in dynamic programming fashion: assuming that the optimization problem is "solved" for the future $\sigma^{t+\theta}$ it is enough to identify an optimal θ and the optimal behavior within $[0,\theta)$.[12]

In order to operationalize this approach, it is necessary to analyze precisely the nature of the integral over $[0,\theta)$ in (15). First, if p_i^t denotes the discrete probability that i issues an acceptance at precisely time t, and $\xi_i^t(0) + \xi_j^t(0) < 1$, then $p_i^t p_j^t$ is the probability of null acceptance. The discontinuity $\Delta F_i^t(0) = p_i^t(1-p_j^t)$ is therefore the probability of *non-null* acceptance of $\xi_j^t(0)$. With the symmetric $\Delta F_j^t(0)$ this translate$s$ at $s=0$ into a "mass"

$$\Delta G_{ij}^t(0) = p_i^t(1-p_j^t)u_i\big(\xi_j^t(0)\big) + p_j^t(1-p_i^t)u_i\big(1-\xi_i^t(0)\big) \tag{16a}$$

and, by Proposition 3 in appendix

$$\int_{[0,\theta)} e^{-r_i s} dG_{ij}^t(s) = \Delta G_{ij}^t(0) + \int_{(0,\theta)} e^{-r_i s} dG_{ij}^t(s) \tag{16b}$$

Second, let

$$\phi_i^t(s) = 1 - \mathbb{P}^t\big(\mathcal{A}_i(t,t+s)\big)$$

be the probability that i issues no acceptance within the leg $(t,t+s)$ according to Ψ^t. ϕ_i^t can be called i's "survival" function *within the leg*. It must clearly be continuous for all $s \in (0,\theta)$ since only the *timing* of acceptance is random and a discontinuity would indicate an event.[13]

Moreover, offers and costs controls are constant for $s \in (0,\theta)$ and can be denoted $\xi_i^t(s) \equiv x_i^t$, $\zeta_i^t(s) \equiv z_i^t$, and $c_i^t \equiv c_i\big(z_i^t, z_j^t\big)$.[14] With these notation one has

***Lemma 1*:** Assume that $(t,t+\theta)$ is a leg. Then

$$\int_{(0,\theta)} e^{-r_i s} dG_{ij}^t(s) = \left(\Phi^t(0^+) - e^{-r_i\theta}\Phi^t(\theta)\right) u_i(x_j^t) \quad - \Phi^t(0^+)\int_{(0,\theta)} \phi_i^t(s) d\kappa_{ij}^t(s)$$

with $$d\kappa_{ij}^t(s) = e^{-r_i s}\Big(\big(u_i(1-x_i^t) - u_i(x_j^t)\big) d\phi_j^t(s) + \big(c_i^t + r_i u_i(x_j^t)\big)\phi_j^t(s) ds\Big) \tag{17}$$

***Proof*:** Discontinuities translate into $\Phi^t(0^+) = 1 - \Delta F_i^t(0) - \Delta F_j^t(0)$ on the previously defined Φ^t. Thus, for all $s \in (0,\theta)$, $dF_i^t(s) = -\Phi^t(0^+)\phi_j^t(s) d\phi_i^t(s)$ and

$$dG_{ij}^t(s) = -\Phi^t(0^+)\Big(u_i\big(x_j^t\big)\phi_j^t(s) d\phi_i^t(s) + u_i\big(1 - x_i^t\big)\phi_i^t(s) d\phi_j^t(s) + c_i^t \phi_i^t(s)\phi_j^t(s) ds\Big)$$

Integrating by parts the term in $d\phi_i^t(s)$ (since $e^{-r_i s}\phi_j^t(s)$ is continuous) yields the result. □

Formulae (15) and (16b), together with Lemma 1, will be used to construct perfect equilibria.

3.2 A Subgame Perfect Equilibrium

The discrete-time Rubinstein theory may have promoted the belief that there is an essentially unique rational solution to the strategic bargaining problem and that it is efficient, at least in the stationary and perfect information case. However, none of this remains true in continuous time. As the results below will show, the very possibility of randomization in the *timing* of acceptance yields a vast class of SPEs that are generically inefficient.

The "countervailing" strategy defined below generalizes the example of §2.4.

***Definition 1*:** A *countervailing* strategy ψ_j (for player j) based on offer $x_j^* \in [0,1]$ and cost control $z_j^* \in \mathcal{C}_j$ is defined by $\psi_j(h_t) := \alpha_j^t$ such that:

• "Initially offer x_j^* and set cost control to z_j^* for all times." Formally:

$\alpha_j^0 := \big(\big(\text{offer } x_j^*\big) \wedge \big(\text{set cost control } z_j^*\big)\big);$ (18a)

• "Accept instantly $\big(1 - x_j^*\big)$.[15] Formally:

$\big(\xi_i(t) = 1 - x_j^*\big) \vee \big(\alpha_i^t \mapsto \text{offer } x_i = 1 - x_j^*\big) \Rightarrow \big(\alpha_j^t := a_j \text{ (accept)}\big);$ (18b)

• Otherwise, "on each leg $(t, t+\theta)$, wait with exponential distribution of parameter $\mu_j = \Lambda_i\big(x_i^t, x_j^t, z_i^t, z_j^t\big)$".[16] Formally:

$\big(x_i^t < 1 - x_j^*\big) \wedge (s = \tau) \Rightarrow (\alpha_j^t := a_j \text{ (accept)})$ (18c)

where $\tau \in (0,\theta)$ is a random variable with distribution $\mathbb{P}(\tau \geq s) = e^{-\mu_j s}$.

This last condition can also be expressed on the resulting function ϕ_j^t that must satisfy

$$d\phi_j^t + \mu_j \phi_j^t ds = 0 \qquad \text{or} \qquad d\kappa_{ij}^t(s) \equiv 0 \tag{18d}$$

By extension, the survival function ϕ_j^t will be called countervailing. The first main result is

***Theorem 1*:** Countervailing strategies form a subgame perfect equilibrium.

***Proof*:** Without loss of generality, one can assume that $x_i^* + x_j^* < 1$ since the initial offers would otherwise be null and the default $x_i^d = x_j^d = 0$ would be used instead of x_i^* and x_j^* throughout this argument.

Assume now that j uses her countervailing strategy. If at any time t i's current offer reaches $\left(1 - x_j^*\right)$ then j instantly (and successfully) accepts.[17] In that case $E_i(\sigma^t) = u_i(x_j^*)$. Otherwise, for any path σ^t beginning at time t, there is probability one that for some $\theta > 0$ $(t, t+\theta)$ is a leg. Within that leg, it was observed that $\phi_j^t(s) = e^{-\mu_j s}$ results in $d\kappa_{ij}^t(s) \equiv 0$. Moreover, since ψ_j results in a continuous acceptance distribution function F_j, any discontinuity at 0 can only involve i's acceptance of x_j^* and yields $\Delta G_{ij}^t(0) = \Delta F_i^t(0) u_i(x_j^*)$. And since $\Delta F_i^t(0) = 1 - \Phi^t(0^+)$, by Lemma 1 and by (15)

$$\begin{aligned} E_i(\sigma^t) &= \Big(\Delta F_i^t(0) + \Phi^t(0^+) - e^{-r_i\theta}\Phi^t(\theta)\Big) u_i(x_j^*) + e^{-r_i\theta}\Phi^t(\theta) E_i(\sigma^{t+\theta}) \\ &= \big(1 - e^{-r_i\theta}\Phi^t(\theta)\big) u_i\big(x_j^*\big) + e^{-r_i\theta}\Phi^t(\theta) E_i(\sigma^{t+\theta}) \end{aligned} \tag{19}$$

Let $(\theta_n)_{n\in\mathbb{N}}$ be the successive event times generated on path σ^t by i's strategy. If there is none or finitely many, use an arbitrary sequence of θ_n approaching ∞. Iterating (19) yields

$$E_i(\sigma^t) = \big(1 - e^{-r_i\theta_n}\Phi^t(\theta_n)\big) u_i\big(x_j^*\big) + e^{-r_i\theta_n}\Phi^t(\theta_n) E_i(\sigma^{t+\theta_n}) \tag{20}$$

By Proposition 4 in appendix, $\lim_{n\to\infty} e^{-r_i\theta_n}\Phi^t(\theta_n) = 0$. And, since E_i is clearly bounded, $E_i(\sigma^t) = u_i\big(x_j^*\big)$. Since this holds for any strategy for player i, the countervailing strategy ψ_i is a best reply after any history h^t. ☐

***Corollary 1* (Extremal Equilibria):** If $x_j^* = 0$ then the SPE of Theorem 1 is extremal for i.

***Proof*:** Player i's expected utility always reduces to $E_i(\sigma^t) = u_i(0)$ which is the minimum i can guarantee himself by accepting j's last offer at any time. ☐

Note that $x_i^* = x_j^* = 0$ provides an extremal equilibrium for *both* sides. In typical applications extremal equilibria are used as trigger threats: both sides play according to some agreed upon "scenario" and defection is punished by reversion to an extremal equilibrium of the deviating side. The next result shows how such a construction can be achieved here.

***Corollary 2* (Efficient Equilibria):** For any $x_i \in [0,1]$ and cost controls z_i^* and z_j^* the following strategy profile forms a SPE.

at $t = 0$ issue

$\alpha_i^0 :=$ input β_j^0; if $(\beta_j^0 = \alpha_j^0)$ then offer x_i else offer 0_i and set z_i^*.

$\alpha_j^0 :=$ input β_i^0; if $(\beta_i^0 = \alpha_i^0)$ then accept a_j else offer 0_j and set z_j^*.

$\forall t > 0$ use countervailing strategies based on $\xi_i(t)$, $\xi_j(t)$, $\zeta_i(t)$, and $\zeta_j(t)$.

***Proof*:** If i and j do issue α_i^0 and α_j^0 these result in the offer x_i immediately accepted. This yields immediate utilities $u_i(1 - x_i) \geq u_i(0)$ and $u_j(x_i) \geq u_j(0)$. Any other decision by i results in nonacceptance and the offer 0_j followed by countervailing behavior on j's part with expected utility $u_i(0)$ for i from then on, clearly not an improvement at $t = 0$. The argument is similar for j. And reversion to countervailing behavior provides an SPE. ☐

A consequence of Corollary 2 is that, in the continuous time framework and under the assumptions of this paper, strategic bargaining allows efficient outcomes but does not yield uniqueness since any sharing of the pie can be sustained by a SPE.

Theorem 2 has conceptual consequences as well: first, it may seem awkward that the parameter μ_j of j's survival function be given by a quantity Λ_i that only involves i's game parameters. Indeed, the denominator of Λ_i is only a measure of how far appart the two sides are, but its numerator is the total cost flow of i - his actual cost c_i of bargaining plus his opportunity cost of not accepting j's offer. One might instead expect that j should hold out as a function of *her* own costs. But that would be a misconception. The reason j's survival function is given by i's cost parameters is that her strategy is to pressure i with expected accumulating costs, by delaying her own acceptance, to wear him down into giving up.

Second, Theorem 1 holds for any pair of offers such that $x_i^* + x_j^* < 1$, whereas immediate acceptance means that $x_i^* + x_j^* = 1$. The first condition is arguably more generic than the second. If the two sides come to the bargaining table without prior exchange, and with some idea of what offer they wish to make, it is more likely that the two won't exactly add up to the whole pie. And all such pairs can be supported by a countervailing SPE. It is tempting to object that, knowing this, the two sides would be both better off choosing offers that result in an efficient outcome. But the process of choosing such offers is precisely what is modeled in this paper and fails to identify a single efficient bargain.

3.2 Optimal Acceptance Behavior

A SPE based on a threat of reversion, as in Corollary 2, need not involve immediate acceptance. Instead, some "scenario" can be implemented, until some deterministic or probabilistic end, and be sustained by an appropriate threat. Such a scenario may involve various offers and cost control adjustments, as well as a continuous or discrete-time probabilistic acceptance strategy. However, although the parameters of probabilistic behavior may be public knowledge, the random variables governing it are here assumed privately monitored. So, if i's expected acceptance behavior is given by some ϕ_i^t within some leg $(t, t+\theta)$, all that j can observe is the occurrence of an acceptance or a rejection. *A trigger threat, therefore, cannot rely on i's conformance to any specific ϕ_i^t that j can never observe directly*. This means that i will undoubtedly choose ϕ_i^t in order to optimize his expected payoff within $(t, t+\theta)$, *ceteris paribus*. It is thus useful to understand what optimal "shapes" ϕ_i^t can take. This section discusses this issue when the following holds:

***Condition AP* (Agreement is Preferred):** for both i on the leg $(t, t+\theta)$

$c_i^t + r_i u_i(x_j^t) > 0$

Condition AP holds if there is a positive cost flow ($c_i^t > 0$) to not reaching agreement, regardless of whether that cost can be manipulated strategically, or if there is an opportunity cost because the offered bargain is better than not having one ($u_i(x_j^t) > 0$).

Lemma 2 and the resulting Theorem 2 give a useful characterization of equilibrium behavior under condition AP: within each closed subleg of a deterministic leg, either both players countervail each other or they are "non-acceptant," meaning that their probability of acceptance is nil.

***Lemma 2*:** Assume that condition AP holds, that $[t, t+\theta]$, $\theta > 0$, is a closed subleg, and that *for each of* i and j the integral (here for i) $\int_0^\theta \phi_i^t(s) d\kappa_{ij}^t(s)$ reaches a minimum. Then, $\phi_i^t(s)$ and $\phi_j^t(s)$ must be *simultaneously* constant or countervailing.

***Proof*:** Let $\mu_j = \Lambda_i\left(x_i^t, x_j^t, z_i^t, z_j^t\right) > 0$, and for all $s \in [0,\theta]$ let:

$$\kappa_{ij}^t(s) = \left(u_i(1-x_i^t) - u_i(x_j^t)\right)\left(e^{-r_i s}\phi_j^t(s) + (r_i + \mu_j)\int_0^s e^{-r_i \tau}\phi_j^t(\tau) ds\right) \quad (21)$$

Clearly $\kappa_{ij}^t(s) > 0$ on $[0,\theta]$ and one easily verifies that $d\kappa_{ij}^t(s)$ is as in (17). Moreover

$$\int_0^\theta \phi_i^t(s) d\kappa_{ij}^t(s) = \phi_i^t(\theta)\kappa_{ij}^t(\theta) - \kappa_{ij}^t(0) - \int_0^\theta \kappa_{ij}^t(s) d\phi_i^t(s) \quad (22)$$

after integrating by parts (since ϕ_i^t and κ_{ij}^t are continuous). There are two cases:

(i) If ϕ_i^t is *strictly* decreasing on $[0,\theta]$ then, according to Proposition 5 in appendix, κ_{ij}^t must be constant for the last integral in (22) to reach an extremum given $\phi_i^t(\theta)$. This implies that $d\kappa_{ij}^t(s) = 0$ and $\phi_j^t(s) = e^{-\mu_j s}$ is countervailing. Since ϕ_j^t is strictly decreasing, κ_{ji}^t must be constant for an extremum to be reached in (22) written for j, *mutatis mutandis* and $\phi_i^t(s) = e^{-\mu_i s}$ is countervailing.

(ii) If ϕ_j^t is *not strictly* decreasing on $[0,\theta]$ then there must be some subleg, call it $[0,\theta]$ again, with $\phi_j^t \equiv 1$. But $d\kappa_{ij}^t(s) = e^{-r_i s}\left(c_i^t + r_i u_i(x_j^t)\right) ds$ and the integral

$$\int_0^\theta \phi_i^t(s) d\kappa_{ij}^t(s) = \left(c_i^t + r_i u_i(x_j^t)\right)\int_0^\theta e^{-r_i s}\phi_i^t(s) ds \quad (23)$$

can only reach an extremum when ϕ_i^t is constant by Proposition 5 (with $df(s) = e^{-r_i s} ds$). But $\phi_i^t \equiv 1$ since $\phi_i^t(0) = 1$. The same holds when exchanging i and j. ☐

The integral that is minimized in Lemma 2 is only one term in the Bellman equation (15). But knowing the optimal shapes of ϕ_i^t on $(0,\theta)$ reduces the optimization problem to a choice of boundary value at θ. Theorem 2 shows that, within a maximal leg, countervailing and non-acceptant behaviors can only cohabit in a very precise way and that they correspond to different types of payoff expectations. One has

***Theorem 2*:** Assume condition AP and that $(t, t+\theta)$ is a maximal deterministic leg of a SPE-generated path σ^t. Then, there exists $0 \le \nu \le \theta$ such that *for both* i and j, *mutatis mutandis*:

(i) $E_i(\sigma^{t+s}) \equiv u_i(x_j^t)$ for $s \in (0, \nu)$ and both sides are countervailing;

and (ii) $E_i(\sigma^{t+s}) > u_i(x_j^t)$ for $s \in (\nu, \theta)$ and both sides are non-acceptant.[18]

***Proof*:** In any SPE one must have $E_i(\sigma^{t+s}) \ge u_i(x_j^t)$ for all $s \in (0, \theta)$ since i can always accept x_j^t at any time. And since there can be no mass within $(0, \theta)$ by definition of a leg one always has, by (15) and by Lemma 1, for any $0 < s < s+\nu < \theta$:

$$E_i(\sigma^{t+s}) = u_i(x_j^t) + e^{-r_i\nu}\phi_i^{t+s}(\nu)\phi_j^{t+s}(\nu)\big(E_i(\sigma^{t+s+\nu}) - u_i(x_j^t)\big) \\ - \int_0^{\nu}\phi_i^{t+s}(\tau)d\kappa_{ij}^{t+s}(\tau) \qquad (24)$$

Since j does not *observe* ϕ_i^{t+s}, its very shape does not influence the parameters of the optimization of $E_i(\sigma^{t+s+\nu})$.[19] Therefore, the maximization by i of $E_i(\sigma^{t+s})$ always involves the minimization over ϕ_i^{t+s} in $[0, \nu]$ of the integral in (24) given any $\phi_i^{t+s}(\nu)$. By Lemma 2, this yields two cases:

(i) If ϕ_j^{t+s} is countervailing then the integral in (24) is nil. And if $E_i(\sigma^{t+s+\nu}) > u_i(x_j^t)$ then (24) is maximum for $\phi_i^{t+s}(\nu) = 1$. But $\phi_i^{t+s} \equiv 1$ is not countervailing and neither can ϕ_j^{t+s} be by Lemma 2. So, $E_i(\sigma^{t+s+\nu}) = u_i(x_j^t)$ and $E_i(\sigma^{t+s}) \equiv u_i(x_j^t)$ for $s \in [0, \nu]$.

Conversely, if $E_i(\sigma^{t+s}) \equiv u_i(x_j^t)$ then $\int_0^{\nu}\phi_i^{t+s}(\tau)d\kappa_{ij}^{t+s}(\tau) = 0$. And since ϕ_i^{t+s} is continuous positive on $[0, \nu]$, $d\kappa_{ij}^{t+s}(\tau) \equiv 0$ (a.e.) and the continuous ϕ_j^{t+s} must satisfy the countervailing condition (18d). By Lemma 2, ϕ_i^{t+s} must also be countervailing.

(ii) If $\phi_j^{t+s} \equiv 1$ in $[0, \nu]$ then (23) and (24) yield

$$E_i(\sigma^{t+s}) = u_i(x_j^t) + e^{-r_i\nu}\phi_i^{t+s}(\nu)\big(E_i(\sigma^{t+s+\nu}) - u_i(x_j^t)\big) \\ - \big(c_i^t + r_i u_i(x_j^t)\big)\int_0^{\nu} e^{-r_i s}\phi_i^{t+s}(\tau)d\tau \qquad (25)$$

Again, for any $\phi_i^{t+s}(\nu)$ the integral in (25) must be minimized but this can only occur with a constant ϕ_i^{t+s} which must be identically 1 by continuity. Since $c_i^t + r_i u_i(x_j^t) > 0$ by the assumed AP, one must have $E_i(\sigma^{t+s+\nu}) > u_i(x_j^t)$. And by continuity in ν, this last strict inequality must hold within some neighborhood of ν. Conversely, this inequality precludes a countervailing ϕ_i^{t+s} by (i) and implies $\phi_i^{t+s} \equiv 1$ in that neighborhood.

Finally, assume the strict inequality $E_i(\sigma^{t+s+\tau}) > u_i(x_j^t)$ for $\tau \in (0, \nu]$ for simplicity. Replacing $\phi_i^{t+s} \equiv 1$ in (25) yields:

$$0 < c_i^t + r_i u_i(x_j^t) \leq c_i^t + r_i E_i(\sigma^{t+s}) = e^{-r_i \tau}(c_i^t + r_i E_i(\sigma^{t+s+\tau})) \tag{26}$$

It follows that $E_i(\sigma^{t+s+\tau})$ must be *increasing* with τ. It can therefore never decrease back to $u_i(x_j^t)$. So, countervailing behavior cannot follow non-acceptant behavior within a same leg, but it can precede it as stated. ☐

In the continuous time bargaining framework, perfect equilibria can only involve two possible kinds of acceptance behavior on a deterministic leg according to Theorem 2: both sides are either countervailing or non-acceptant. These are two forms of "holding-out" behavior: in the first case, both sides hold-out in the hope that the other will be first to concede; in the second, they hold-out in the expectation of a better future. And under condition AP, if the players are not countervailing on a leg their expected payoffs must be *increasing* with time. This is not surprising since non-acceptant behavior yields either accumulating cost $c_i^t > 0$ or the opportunity cost of not accepting $u_i(x_j^t) > 0$ sooner. And these are not compensated by some probability that the other side will accept along the way according to countervailing. So, each side must be incurring these costs in the expectation of a better future at the end of the leg, perhaps as a better offer or as a discrete probability of the other side's accepting. As this better future draws nearer, one's expected payoff increases. If instead the players are countervailing then their expected payoffs are held to the constant utility of the other side's current offer.

3.3 A Second Example

Consider the following Rubinstein-like negotiating framework: the two sides "agree" that they will alternately make offers at regular intervals of time of given length $\theta > 0$, perhaps with i going first. They also agree that when i makes an offer to j, j instantly withdraws his own last offer to i (i.e. adopts offer magnitude 0). Although i's offer remains on the table until j's next turn to make one, neither side is expected to accept or even speak between turns. Indeed, if anyone deviates from this "scenario", for instance by not speaking at their turn or not making the expected offer, they both instantly revert to countervailing behavior on the basis of their last offers and cost controls, a well known SPE.

There are many possibilities regarding the choice of θ and the two sides' successive offers. For simplicity, let us assume here that the two sides stick stubbornly to offers $x_i^* > 0$ and $x_j^* > 0$ such that $x_i^* + x_j^* < 1$ and to controls z_i^* and z_j^* resulting in costs $c_i^* \geq 0$ that they impose on each other from the very start. Finally, assume that each side accepts the other side's offer at each of his turns with probability p_j (for j). So, if t is a time at which i receives j's offer x_j^* and may accept it rationally with probability $p_i > 0$, one needs by (15)

$$\begin{aligned} E_i(\sigma^t) &= u_i(x_j^*) \\ &= -k_i + e^{-r_i\theta}\Big(p_j u_i(1-x_i^*) + (1-p_j)\big(-k_i + e^{-r_i\theta}u_i(x_j^*)\big)\Big) \end{aligned} \tag{27}$$

with $k_i = (1-e^{-r_i\theta})\frac{c_i^*}{r_i}$ since $E_i(\sigma^{t+2\theta}) = u_i(x_j^*)$. This requires that

$$p_j = \frac{(e^{r_i\theta}-e^{-r_i\theta})u_i(x_j^*)+(1+e^{r_i\theta})k_i}{u_i(1-x_i^*)-e^{-r_i\theta}u_i(x_j^*)+k_i} \tag{28}$$

which is a true probability provided that

$$0 \leq k_i < e^{-r_i\theta}u_i(1-x_i^*) - u_i(x_j^*)$$

a condition that holds provided θ is small enough.

After time t, and until time $(t+\theta)$, $E_i(\sigma^{t+s}) > u_i(x_j^*)$ and it is not rational for i to accept x_j^*. Before time t, from $(t-\theta)$ on, x_j^* is not available and i can only accept the lesser offer 0_j. So, provided θ is small enough, and since for $s \in [0,\theta)$

$$E_i(\sigma^{t-\theta+s}) \geq E_i(\sigma^{t-\theta}) = -k_i + e^{-r_i\theta}u_i(x_j^*) > u_i(0) \tag{29}$$

it is never rational for i to accept at such times. So, it is only at precisely time t (plus multiples of 2θ) that i finds it rational to accept x_j^* with any probability p_i. Of course, the situation is entirely symmetrical for j who can thus use probability p_j optimally at time $(t+\theta)$. Deviation in any way only brings countervailing behavior on the basis of the other side's current offer (x_j^* or 0_j) a shift that can never improve one's expected payoff. This scheme thus forms a SPE. The probability that the other side accepts at the next turn is what provides the brighter future that makes it possible to wait for time θ. In the discrete-time Rubinstein framework the above scheme would not provide an SPE if utilities are weakly concave. What makes it possible in this continuous time framework is the fact that the two sides can move simultaneously (when one side makes its offer the other withdraws its own) and can revert instantly to a countervailing SPE if either fails to carry out the agreed upon scenario. This suggests that it is only by depriving the players of the full control of the timing of their moves that the above scheme can be ruled out in the Rubinstein framework.

3.4 A Third Example

An interesting scheme that is worked out by Smith & Stacchetti (2003) makes the delay θ be a random variable and the successive offers approach full agreement. They also assume that an offer must be instantly accepted or rejected. However, it is possible to recast their model of "aspirational bargaining" in the terms of this paper where the timing of offers and acceptance are independent. Let us assume that x_j^t is j's current offer at time t and that j intends to make a new offer denoted $x_j^{t+\theta}$ after some delay θ. But, instead of θ being deterministic, let it be a random variable of distribution function $f_j^t(\theta)$. Let us further assume that j countervails i at all times on the basis of the current offer x_i^t while it lasts, and on the basis of $x_j^{t+\theta}$ once it is made. Equation (15) together with the countervailing conditions $E_i(\sigma^t) = u_i(x_j^t)$ and $E_i(\sigma^{t+\theta}) = u_i(x_j^{t+\theta})$ yields an expected value

$$u_i(x_j^t) = \int_0^\infty \left(\int_0^\theta e^{-r_i s} dG_{ij}^t(s) + e^{-r_i \theta} \Phi^t(\theta) u_i(x_j^{t+\theta}) \right) df_j^t(\theta) \qquad (30)$$

Following Smith and Stacchetti let us assume, for example, that $f_j^t(\theta) = 1 - e^{-\rho_j \theta}$ with some $\rho_j > 0$. One can still anticipate that countervailing should involve exponential densities with parameters that should account for the distribution of θ. So, let us assume that $\phi_i^t(s) = e^{-\mu_i s}$ for $s \in [0, \theta)$ and symmetrically for j. Equation (30) then becomes

$$u_i(x_j^t) = \int_0^\infty \Big(\int_0^\theta e^{-(r_i+\mu_i+\mu_j)s} (\mu_i u_i(x_j^t) + \mu_j u_i(1 - x_i^t) - c_i^t) ds + e^{-(r_i+\mu_i+\mu_j)\theta} u_i(x_j^{t+\theta}) \Big) \rho_j e^{-\rho_j \theta} d\theta \quad (31)$$

or, after integration

$$u_i(x_j^t) = \frac{\mu_i u_i(x_j^t) + \mu_j u_i(1-x_i^t) - c_i^t + \rho_j u_i(x_j^{t+\theta})}{r_i + \mu_i + \mu_j + \rho_j}$$

which yields

$$\mu_j = \frac{c_i^t + r_i u_i(x_j^t) - \rho_j \big(u_i(x_j^{t+\theta}) - u_i(x_j^t) \big)}{u_i(1-x_i^t) - u_i(x_j^t)} \quad (32)$$

As $\rho_j \to 0$ the offer $x_j^{t+\theta}$ is delayed indefinitely and formula (32) reduces to the standard $\mu_j = \Lambda_i(x_i^t, x_j^t, z_i^t, z_j^t)$. But a positive ρ_j introduces a correction in the numerator of Λ_i. If the offer $x_j^{t+\theta}$ is expected to be better for i than the current x_j^t then the correction is negative and this decreases μ_j and lowers the chances of acceptance by j while x_j^t is still on the table. Intuitively, $x_j^{t+\theta}$ will be more attractive and more likely to be accepted by i and this favors some temporizing by j. If instead $x_j^{t+\theta}$ is expected to be worse than the current x_j^t, j is more likely to accept. This corresponds to a threat of withdrawing a current "good" offer which would yield more temporizing. Also note that the condition

$$\rho_j \big(u_i(x_j^{t+\theta}) - u_i(x_j^t) \big) \leq c_i^t + r_i u_i(x_j^t)$$

must hold for μ_j to be non-negative as required. So, the expected time $\frac{1}{\rho_j}$ it will take for a generous offer $x_j^{t+\theta}$ to be made may have to be higher than for a less generous one.

Of course, $x_j^{t+\theta}$ can be followed by further offers with randomized timing and this process can be bilateral and involve changes in costs as well. But if both sides countervail each other in this fashion they are in SPE. This suggests that Smith and Stacchetti's theory of aspirational bargaining is not limited to the "take-it-or-leave-it" offer framework.

4. Bargaining with Incomplete Information

The goal in this section is to extend the existence of equilibrium result of the perfect information case to incomplete information. The natural solution concepts are the perfect Bayesian (PBE) and the sequential equilibria.

4.1 Beliefs Updating

Instead of two individuals i and j bargaining with each other, let us now assumes two *finite* sets of types, with $i \in \mathcal{I}$ and $j \in \mathcal{J}$. Each type i is described by a specific utility function u_i, cost function c_i, and discount rate r_i. But, since neither side can observe who among the other side's types makes decisions on offer, cost, and acceptance, these will be denoted by $x_{\mathcal{I}}$, $z_{\mathcal{I}}$, $x_{\mathcal{J}}$, and $z_{\mathcal{I}}$. At each time t each side has beliefs $\left(b_j(t)\right)_{j\in\mathcal{J}}$ about the other side's type (i.e., $b_j(t) \geq 0$ and $\sum\limits_{j\in\mathcal{J}} b_j(t) = 1$) that will be updated by Bayes' Law.

The natural generalization of player i's objective was given by (11). But belief updating must be formulated in the time continuum, both at and between event times. Initial beliefs at $t = 0$ can be arbitrary, although the initial choices of offers and cost controls may immediately result in an adjustment of beliefs. In general, Bayesian updating reads for $s \geq 0$

$$b_j(t+s) = \frac{b_j(t)\mathbb{P}(S_j^{t+s})}{\sum\limits_{k\in\mathcal{J}} b_k(t)\mathbb{P}(S_k^{t+s})} \qquad (33)$$

where S_k^{t+s} is the state in which type $k \in \mathcal{J}$ will find itself at time $(t+s)$ according to strategy. A standard component of S_k^{t+s} is that k has not yet successfully accepted but it may in general involve offer and cost decisions. However, this section only aims to establish an existence of equilibrium result that extends the concept of holding-out behavior. In particular, it does not explore cases where cost control adjustments would reveal information about one's type. Let us therefore make

***Assumption A3* (Incentive Compatibility):** The cost game, if it is non-trivial, has a Nash equilibrium $(z_{\mathcal{I}}^*, z_{\mathcal{J}}^*)$ that is common to all types.

With this assumption, it is possible to focus on PBEs where offer and cost control decisions are pooling and only acceptance behavior can be separating. In addition, the equilibria obtained in Theorem 3 rely on behavior that rules out null acceptance events. This is guaranteed when the following holds:

***Condition CA* (Continuous Acceptance):** Side $\mathcal{J}$'s strategy profile is "continuously acceptant" (CA) if for all $k \in \mathcal{J}$ and all t, F_k^t is a continuous function.

Clearly, a continuous F_k^t involves zero probability of acceptance by k at any one time and therefore cannot (with any positive probability) lead to any simultaneous attempt that could generate a null acceptance event. In all the results of this section at least one side will always be CA and notation will be simpler if the survival function is now defined by $\varphi_i^t(s) = 1 - \mathbb{P}^t(\mathcal{A}_i[t, t+s))$.[20] Whenever at least one side satisfies CA, the Bayesian updating formula (33) simplifies into[21]

$$b_j(t+s) = \frac{b_j(t)\varphi_j^t(s)}{\varphi_{\mathcal{J}}^t(s)} \tag{34}$$

where $$\varphi_{\mathcal{J}}^t(s) = \sum_{k \in \mathcal{J}} b_k(t)\varphi_k^t(s) \tag{35}$$

is the probability that no type in $\mathcal{J}$ has accepted by time $(t+s)$. One further consequence of condition CA is that any mass at t can only result from a one-sided discrete probability of acceptance. If all types in $\mathcal{J}$ are CA then (16a) is replaced by

$$\Delta G_{i\mathcal{J}}^t(0) = \sum_{j \in \mathcal{J}} b_j(t)\Delta G_{ij}^t(0) = -\Delta\varphi_i^t(0)u_i\left(\xi_{\mathcal{J}}^t(0)\right) \tag{36}$$

4.2 A Second Bellman Equation

***Lemma 3*:** If (34) holds then for any path σ^t

$$E_i(\sigma^t) = \sum_{j \in \mathcal{J}} b_j(t)\int_{[0,\theta)} e^{-r_i s} dG_{ij}^t(s) + e^{-r_i\theta}\varphi_i^t(\theta)\varphi_{\mathcal{J}}^t(\theta)E_i(\sigma^{t+\theta}) \tag{37}$$

***Proof*:** Since our distribution functions are *left*-continuous one can always write

$$E_i(\sigma^t) = \sum_{j \in \mathcal{J}} b_j(t)\int_{[0,\theta)} e^{-r_i s} dG_{ij}^t(s) + \sum_{j \in \mathcal{J}} b_j(t)\int_{[\theta,\infty)} e^{-r_i s} dG_{ij}^t(s) \tag{38}$$

But, just as with (14a) and (14b), $dG_{ij}^t(\theta+s) = \varphi_i^t(\theta)\varphi_j^t(\theta)dG_{ij}^{t+\theta}(s)$ since $\varphi_i^t(\theta)\varphi_j^t(\theta)$ is the joint survival function of i and j under condition CA. The second sum in (38) thus reads

$$\sum_{j \in \mathcal{J}} b_j(t)\int_{[0,\infty)} e^{-r_i(\theta+s)} dG_{ij}^t(\theta+s) = e^{-r_i\theta}\varphi_i^t(\theta)\sum_{j \in \mathcal{J}} b_j(t)\varphi_j^t(\theta)\int_{[0,\infty)} e^{-r_i s} dG_{ij}^{t+\theta}(s)$$

$$= e^{-r_i\theta}\varphi_i^t(\theta)\varphi_{\mathcal{J}}^t(\theta)\sum_{j\in\mathcal{J}} b_j(t+\theta)\int_{[0,\infty)} e^{-r_i s} dG_{ij}^{t+\theta}(s) = e^{-r_i\theta}\varphi_i^t(\theta)\varphi_{\mathcal{J}}^t(\theta)E_i(\sigma^{t+\theta})$$

by (34). □

Formula (37) can also be called a Bellman equation since it breaks the optimization problem into one over the interval $[0,\theta)$ together with the choice of an optimal θ and the result of the same optimization problem shifted by θ into the future.[22] Indeed, the addition of beliefs is the only significant difference with the perfect information case in the optimization over $[0,\theta)$. This suggests that optimal behavior should still be either countervailing or non-acceptant. But, although this could be formally discussed, it is enough here to describe an equilibrium that uses the observation.

Countervailing behavior is meant to hold one side's expected payoff to its utility of the other side's current offer. It is a razor-thin condition: any slight adjustment in i's utility, cost, or impatience parameters should jeopardize that balance. So, if i is countervailed it is unlikely that any other type l of its side, with different parameters, also is. And if l is not countervailed, his optimal behavior, short of instant acceptance, should be non-acceptant. This suggests a construction that is in fact quite standard in incomplete information games of attrition: side $\mathcal{J}$ (all types together) will countervail a specific type $i \in \mathcal{I}$ while all other types $l \in \mathcal{I}$ with non-nil beliefs will be made non-acceptant, and symmetrically.

4.3 The Case of a Single Countervailed Type

This section establishes some essential lemmata for the case when only one type on each side is countervailed while all other types are either given nil beliefs or remain non-acceptant. Each side's countervailed type is also the only one who "actively" countervails the other. These lemmata lead to an existence of PBE theorem in the next section. Using the notation $x_{\mathcal{I}}^t$, $x_{\mathcal{J}}^t$, $c_{\mathcal{I}}^t$, and $c_{\mathcal{J}}^t$ for constant decisions within legs, Lemma 1 becomes

***Lemma 4*:** Assume that all $k \in \mathcal{J}$ satisfy condition CA and that $(t, t+\theta)$ is a leg. Further assume for some $j \in \mathcal{J}$ that $b_j(t) \neq 0$, and that $\varphi_k^t \equiv 1$ for all $k \in \mathcal{J}$ such that $k \neq j$ and $b_k(t) \neq 0$. Then for all $l \in \mathcal{I}$:

$$\sum_{k\in\mathcal{J}} b_k(t)\int_{(0,\theta)} e^{-r_l s} dG_{lk}^t(s)$$

$$= \left(\varphi_l^t(0^+) - e^{-r_l\theta}\varphi_l^t(\theta)\varphi_{\mathcal{J}}^t(\theta)\right)u_l(x_{\mathcal{J}}^t) - \int_{(0,\theta)} \varphi_l^t(s)d\kappa_{l\mathcal{J}}^t(s) \tag{39}$$

with $\quad d\kappa_{l\mathcal{J}}^t(s) = e^{-r_l s}\Big(\big(u_l(1-x_{\mathcal{I}}^t) - u_l(x_{\mathcal{J}}^t)\big)d\varphi_{\mathcal{J}}^t(s) + \big(c_l^t + r_l u_l(x_{\mathcal{J}}^t)\big)\varphi_{\mathcal{J}}^t(s)ds\Big)$

***Proof*:** Similarly to Lemma 1, integrate by parts the term $-e^{-r_l s}\varphi_k^t(s)d\varphi_l^t(s)$ in $e^{-r_l s}dG_{lk}^t(s)$ (integrating $d\varphi_l^t(s)$). Then use the definition of $\varphi_{\mathcal{J}}^t$. ☐

A countervailing profile for $\mathcal{J}$ will again involve constant offer and cost control. But most importantly, $\varphi_{\mathcal{J}}^t$ will also satisfy a condition akin to (18d) that reads here

$$d\kappa_{i\mathcal{J}}^t(s) \equiv 0 \tag{40}$$

Together with $\varphi_k^t \equiv 1$, for $k \neq j$ and $b_k(t) \neq 0$, this translates into a condition on φ_j^t:

***Lemma 5*:** Under the conditions of Lemma 4, formula (40) holds if and only if

$$\varphi_j^t(s) = \frac{e^{-\mu_j s} + b_j(t) - 1}{b_j(t)} \tag{41}$$

with $\mu_j = \Lambda_i\big(x_{\mathcal{I}}^t, x_{\mathcal{J}}^t, z_{\mathcal{I}}^t, z_{\mathcal{J}}^t\big) > 0$ and $\theta \leq \frac{-ln\left(1-b_j(t)\right)}{\mu_j}$.

***Proof*:** Rewrite $d\kappa_{i\mathcal{J}}^t(s) \equiv 0$ as $d\varphi_{\mathcal{J}}^t + \mu_j\varphi_{\mathcal{J}}^t ds = 0$, so that $\varphi_{\mathcal{J}}^t(s) = e^{-\mu_j s}$. Then use the definition of $\varphi_{\mathcal{J}}^t$ to obtain φ_j^t. Finally solve $\varphi_j^t(\theta) = 0$ to obtain the bound on θ. ☐

Note that, by Bayes Law (34), $b_j(t+\theta) = 0$ since $\varphi_j^t(\theta) = 0$. And since θ is finite belief about type j's reaches zero at the end of the leg where she is actively countervailing. So, in order to continue countervailing the other side some new type from side $\mathcal{J}$ must become involved. In fact, it is the careful organizing of this succession that provides a PBE in Theorem 3. Also, Corollary 3 will show that, unless acceptance has been reached, each side must learn the other side's exact type in finite time. Lemma 6 is the final key to Theorem 3. It uses the notation $\lambda_l = \Lambda_l\big(0, 0, z_{\mathcal{I}}^*, z_{\mathcal{J}}^*\big)$.

***Lemma 6*:** Under the conditions of Lemma 4, assume further that side $\mathcal{J}$ uses $x_{\mathcal{J}}^t = 0$ and $z_{\mathcal{J}}^t = z_{\mathcal{J}}^*$ and that $\mu_j = \lambda_i$ for some $i \in \mathcal{I}$. Then, for all $x_{\mathcal{I}}^t$, $z_{\mathcal{I}}^t$, and $l \in \mathcal{I}$

$$d\kappa_{l\mathcal{J}}^t(s) = -e^{-(r_l+\lambda_i)s}\rho_l\big(x_{\mathcal{I}}^t, z_{\mathcal{I}}^t\big)ds \tag{42a}$$

with $\quad \rho_l\big(x_{\mathcal{I}}^t, z_{\mathcal{I}}^t\big) = \lambda_i\big(u_l(1-x_{\mathcal{I}}^t) - u_l(0)\big) - \big(c_l(z_{\mathcal{I}}^t, z_{\mathcal{J}}^*) + r_l u_l(0)\big) \qquad (42b)$

Moreover, for all $\big(x_{\mathcal{I}}^t, z_{\mathcal{I}}^t\big)$

$$\rho_l(x^t_{\mathcal{I}}, z^t_{\mathcal{I}}) \leq \rho_l(0, z^*_{\mathcal{I}}) = (\lambda_i - \lambda_l)\big(u_l(1) - u_l(0)\big) \tag{42c}$$

***Proof*:** Replace $x^t_{\mathcal{J}} = 0$, $z^t_{\mathcal{J}} = z^*_{\mathcal{J}}$, and $\varphi^t_{\mathcal{J}} = e^{-\lambda_i s}$ in (39) to obtain (42a) and (42b). The first inequality in (42c) holds because u_l is increasing and $z^*_{\mathcal{I}}$ is a best reply to $z^*_{\mathcal{J}}$ in the cost game. Finally one can factor $\big(u_l(1) - u_l(0)\big)$. □

4.4 The Existence of Perfect Bayesian Equilibria

To exploit the above ideas in the construction of a PBE, one needs to describe which type on each side is countervailed and is *actively* countervailing the other. At all times during the game, each set of types is divided into three subsets: (1) those types on which belief is nil and who are *inactive*; (2) the single type who is *active*; and (3) those types that are non-acceptant and are *waiting* to possibly become active. As the game unfolds, the beliefs about the current active types on each side decrease to zero in finite time. When either belief reaches zero, that type joins the subset of inactive types and is expected to play no further role. It is then instantly replaced by a previously waiting type who thus becomes active.

As the process develops in time, each side actually gains information about the other since it successively eliminates some of the types as their possible actual counterpart. Of course the current active type on the other side might be the actual opponent and could thus accept at any time. This combination of countervailing and non-acceptant behavior on each side is what is meant by "holding-out for concession and for information."

The set of waiting and active types, those with positive beliefs at time t, will be denoted $\mathcal{I}'(t) \subset \mathcal{I}$ and $\mathcal{J}'(t) \subset \mathcal{J}$ respectively. $z^*_{\mathcal{J}}$ denotes the Nash equilibrium cost control of side $\mathcal{J}$. With the notation λ_k introduced before Lemma 6, Definition 1 is generalizes to

***Definition 2*:** A *holding-out* strategy profile $\Psi_{\mathcal{J}} = (\psi_k)_{k \in \mathcal{J}}$ for side $\mathcal{J}$ is defined by $\psi_k(h_t) := \alpha^t_k$ such that:

• "Initially offer $x_{\mathcal{J}} = 0$ and set cost control to $z^*_{\mathcal{J}}$ for all times." Formally, $\forall k \in \mathcal{J}$:

$$\alpha^0_k := ((\text{offer } x_{\mathcal{J}} = 0) \wedge (\text{set cost control } z^*_{\mathcal{J}} \in \mathcal{C}_j)); \tag{43a}$$

• "Accept instantly the whole pie $x_I = 1$. Formally, $\forall k \in \mathcal{J}$:

$(\xi_{\mathcal{I}}(t) = 1) \vee (\alpha^t_{\mathcal{I}} \mapsto (\text{offer } x_{\mathcal{I}} = 1)) \Rightarrow (\alpha^t_k := a_j \text{ (accept)});$ (43b)

• Otherwise, at each time t let the active types be $j = \text{argmax}\{\lambda_k | k \in \mathcal{J}'(t)\}$ and $i = \text{argmax}\{\lambda_l | l \in \mathcal{I}'(t)\}$;[23] Further let $\theta = \min\{\frac{-ln(1-b_j(t))}{\lambda_i}, \frac{-ln(1-b_i(t))}{\lambda_j}\}$.[24]

And for all $s \in [0, \theta)$:

(a) For all inactive types accept instantly. Formally:

$\big(k \in \mathcal{J} - \mathcal{J}'(t)\big) \Rightarrow (\alpha^{t+s}_k := a_k \text{ (accept)});$

(b) For all waiting types do not accept. Formally:

$\big(k \in \mathcal{J}'(t) - \{j\}\big) \Rightarrow (\alpha^{t+s}_k := \emptyset_k);$

(b) For the active type j on leg $(0, \theta)$ wait with probability given below. Formally:

$\big(x^t_{\mathcal{I}} < 1\big) \wedge (s = \tau) \Rightarrow (\alpha^{t+s}_j := a_j \text{ (accept)})$ (43c)

where $\tau \in (0, \theta)$ is a random variable with distribution given by

$\mathbb{P}(\tau \geq s) = \varphi^t_j(s) = \frac{e^{-\lambda_i s} + b_j(t) - 1}{b_j(t)}$ (43d)

One has

Theorem 3**:** The holding-out strategy profile $(\Psi_{\mathcal{I}}, \Psi_{\mathcal{J}})$ together with continuous Bayesian updating (34) forms a PBE.

Proof**:** Assume that after any history h^t side $\mathcal{J}$ plays according to the induced profile $\Psi^t_{\mathcal{J}}$ which clearly satisfies condition CA since all φ^t_k are continuous for $k \in \mathcal{J}$. For any $l \in \mathcal{I}$, any strategy ψ^t_l induced by any history h^t, and any leg $(t, t+\theta)$ of any path σ^t, by lemmata 3, 4, and 6, one has[25]

$$E_l(\sigma^t) = \big(1 - e^{-r_l\theta}\varphi^t_l(\theta)\varphi^t_{\mathcal{J}}(\theta)\big)u_l(0) + \rho_l(x^t_{\mathcal{I}}, z^t_{\mathcal{I}})\int_0^\theta e^{-(r_l+\lambda_i)s}\varphi^t_l(s)ds$$
$$+ e^{-r_l\theta}\varphi^t_l(\theta)\varphi^t_{\mathcal{J}}(\theta)E_l(\sigma^{t+\theta}) \quad (44)$$

Since $\int_0^\theta e^{-(r_l+\lambda_i)s}\varphi^t_l(s)ds \geq 0$ it is always best for l to maximize $\rho_l(x^t_{\mathcal{I}}, z^t_{\mathcal{I}})$ by the choice $(x^t_{\mathcal{I}}, z^t_{\mathcal{I}}) = (0, z^*_{\mathcal{I}})$ according to Lemma 6. Let us first characterize $E_l(\sigma^t)$ when l plays according to ψ^t_l in $\Psi^t_{\mathcal{I}}$. Second, let us show that no "single-leg" deviation from strategy ψ^t_l can be beneficial for $l \in \mathcal{I}$. And third, let us extend this last claim to any multi-leg deviation in dynamic programming fashion.

Step 1: Assume that type $k \in \mathcal{J}$ plays according to ψ_k^t in profile $\Psi_{\mathcal{I}}^t$. Since there can thus be no event[26] $(0, \infty)$ is the only maximal leg and one can let θ be the supremum of all τ such that $b_l(t+\tau) > 0$ (so θ may be ∞). This implies that $e^{-r_l\theta}\varphi_l^t(\theta) = 0$ and by (44)

$$E_l(\sigma^t) = u_l(0) + (\lambda_i - \lambda_l)\big(u_l(1) - u_l(0)\big)\int_0^\theta e^{-(r_l+\lambda_i)s}\varphi_l^t(s)ds \tag{45}$$

If $l \notin \mathcal{I}'(t)$ then $E_l(\sigma^t) = u_l(0)$ since $\theta = 0$. If $l = i$ then $E_l(\sigma^t) = u_l(0)$ since $\lambda_l = \lambda_i$. And if $l \in \mathcal{I}'(t) - \{i\}$ then $E_l(\sigma^t) \geq u_l(0)$ since i denotes the active type in $\mathcal{I}$ so that $\lambda_i \geq \lambda_l$.

Step 2 (First leg deviation)**:** Consider any deviation $\widetilde{\psi}_l^t$ by any type $l \in \mathcal{I}$, after any history h^t, from his assigned strategy $\psi_l^t \in \Psi_{\mathcal{I}}^t$ with the following restriction: $\widetilde{\psi}_l^t$ may involve any $x_{\mathcal{I}}^t$, $z_{\mathcal{I}}^t$, and $\widetilde{\varphi}_l^t$ on a first leg $(t, t+\theta)$ of any path σ^t. But $\widetilde{\psi}_l^t \equiv \psi_l^t$ on $[t+\theta, \infty)$ (i.e., l reverts to his strategy in $\Psi_{\mathcal{I}}$ at the first event). There are three cases

If $l \notin \mathcal{I}'(t)$ then $E_l(\sigma^{t+\theta}) = u_l(0)$, $\lambda_i \leq \lambda_l$ and $\rho_l(0, z_{\mathcal{I}}^*) \leq 0$ so that by (44)

$$E_l(\sigma^t) \leq \big(1 - e^{-r_l\theta}\varphi_l^t(\theta)\varphi_{\mathcal{J}}^t(\theta)\big)u_l(0) \tag{46}$$

and it is best for l to accept instantly. So $\widetilde{\psi}_l^t \equiv \psi_l^t$ is best.

If $l = i$ then $E_i(\sigma^{t+\theta}) = u_i(0)$ since i will be either active or inactive at time $(t+\theta)$. Thus $E_i(\sigma^t) = u_l(0)$ and $\widetilde{\psi}_i^t \equiv \psi_i^t$ is best.

If $l \in \mathcal{I}'(t) - \{i\}$ then $E_l(\sigma^{t+\theta}) \geq u_l(0)$ since l will be either waiting or active at time $(t+\theta)$. Moreover, $\lambda_i \geq \lambda_l$ so that $\rho_l(0, z_{\mathcal{I}}^*) \geq 0$ and (44) is maximum for $\varphi_l^t \equiv 1$, resulting in $E_l(\sigma^t) \geq u_l(0)$. Thus $\widetilde{\psi}_l^t \equiv \psi_l^t$ is best.

Step 3 (Multi-leg deviation)**:** Consider any (unrestricted) deviation $\widetilde{\psi}_l^t$ by $l \in \mathcal{I}$ from his assigned strategy ψ_l^t after history h^t. And let $\widetilde{\Psi}^t$ be the strategy profile where all types $k \neq l$ play their assigned strategy in Ψ^t, except l who plays $\widetilde{\psi}_l^t$. For any path σ^t, $\widetilde{\Psi}^t$ determines an increasing sequence $(\theta_n)_{n\in\mathbb{N}}$ of event times. Now, let $\widetilde{\psi}_l^0 = \psi_l^t$ and for any n

$$\widetilde{\psi}_l^n(h^{t+s}) = \begin{cases} \widetilde{\psi}_l^t(h^{t+s}) \text{ if } s < \theta_n \\ \psi_l^t(h^{t+s}) \text{ if } s \geq \theta_n \end{cases} \tag{47}$$

Since $\widetilde{\psi}_l^n$ and $\widetilde{\psi}_l^{n+1}$ are identical on $[t, t+\theta_n)$ so are the corresponding $\sum_{k\in\mathcal{J}} b_k(t)\int_{[0,\theta_n)} e^{-r_l s} dG_{lk}^t(s)$. And since $\widetilde{\psi}_l^{n+1}$ is the first-leg deviation from $\widetilde{\psi}_l^n \equiv \psi_l^t$ on $[t+\theta_n, \infty)$, by Step 2 we can write for any l and n

$$E_l(\sigma^t|\widetilde{\psi}_l^{n+1}, \Psi_{-l}^t) \leq E_l(\sigma^t|\widetilde{\psi}_l^n, \Psi_{-l}^t) \tag{48}$$

where the notation $E_l(\sigma^t|\Psi)$ stresses what profile Ψ is in play on σ^t. Since the probability that $\lim_{n\to\infty} \theta_n = \Theta < \infty$ is nil either $\lim_{n\to\infty} \theta_n = \infty$ or $\theta_n = \infty$ for some finite n. In either case

$$\begin{aligned} E_l(\sigma^t|\widetilde{\psi}_l^t, \Psi_{-l}^t) &= \lim_{n\to\infty} E_l(\sigma^t|\widetilde{\psi}_l^n, \Psi_{-l}^t) \\ &\leq E_l(\sigma^t|\widetilde{\psi}_l^0, \Psi_{-l}^t) = E_l(\sigma^t|\psi_l^t, \Psi_{-l}^t) \end{aligned} \tag{49}$$

Thus, no deviation $\widetilde{\psi}_l^t$ from ψ_l^t can benefit the arbitrary type l. □

***Corollary 3*:** In the PBE of Theorem 3, players learn their opponent's type in finite time.

***Proof*:** Assume that types $j \in \mathcal{J} = \mathcal{J}'(0)$ are indexed according to a *decreasing* order of the corresponding $\lambda_j = \Lambda_j\big(0_{\mathcal{I}}, 0_{\mathcal{J}}, z_{\mathcal{I}}^*, z_{\mathcal{J}}^*\big)$. If $j < |\mathcal{J}|$ is the active type starting at time t_j, then it stops being active when its belief falls to zero at time $t_{j+1} = t_j + \theta_j$ with $\theta_j = \frac{-ln\left(1-b_j(t_j)\right)}{\mu_j}$ and $\mu_j = \lambda_i$, where $i \in \mathcal{I}$ is the targeted type. Meanwhile all types $k > j$ are waiting and have beliefs that *increase* to $b_k(t_{j+1}) = \frac{b_k(t_j)}{1-b_j(t_j)} < 1$, as long as $k < |\mathcal{J}|$. It follows that all θ_j for $j < |\mathcal{J}|$ are finite and so is their finite sum. □

After complete information is reached according to Corollary 3 the two remaining active types play a perfect information game and could deviate from the extreme choices $x_{\mathcal{J}} = 0$ and $z_{\mathcal{J}} = z_{\mathcal{J}}^*$. They could do so as suggested in sections 3.3. So one can construct a PBE that allows more efficient reciprocating behavior once full information is reached. However, any trembling in the beliefs would again require the use of extreme choices. So, a *sequential* equilibrium requires that extreme choices be used even once full information has been reached. Formally

***Corollary 4*:** The PBE of Theorem 3 with extreme choices $x_{\mathcal{J}} = 0$ and $z_{\mathcal{J}} = z_{\mathcal{J}}^*$ at all times forms a sequential equilibrium.

***Proof*:** Introduce trembling ϵ_i so that $\beta_i(t) = b_i(t) + \epsilon_i > 0$ and $\sum_{i \in \mathcal{I}} \epsilon_i = 0$. Replacing $b_i(t)$ by $\beta_i(t)$ everywhere above yields a PBE. As all $\epsilon_i \to 0$ one easily shows that it converges to the PBE of Theorem 3 with the given constant extreme choices. □

5. Conclusion

This paper extends the players' action space from simple offer, cost control, and acceptance moves of strategic bargaining theory to logical instructions such as "take-it-or-leave-it" or "accept at least this much". Most importantly, it allows the timing of all such instructions to be the players' decision variables within the time continuum rather than the fiat of the theorist. The natural generalization of the discounted payoff objective of discrete-time games is defined by a Lebesgue-Stieltjes integral that reflects the expected payoff of acceptance by either side minus the expected costs of waiting.

Within the perfect information framework, I construct and characterize subgame perfect equilibria. I illustrate with various examples that could not possibly arise in the standard Rubinstein model. In particular, disagreement over the sharing of the pie can persist for a while and is the generic choice. Under a mild cost incentive compatibility assumption and with continuous Bayesian updating of beliefs, I obtain perfect Bayesian and sequential equilibria. Their structure suggests that uncertainty favors ungenerous and hurtful bargaining choices. But complete information is reached in finite time allowing more efficient choices.

These findings add support to the growing suspicion that the uniqueness and efficiency results of the standard Rubinstein model are the product of its temporal monopoly assumption and do not survive when it is lifted. When the timing of decisions is returned to the players, the cost of waiting a *fixed* additional period, upon which the Rubinstein model depends, evaporates thus allowing a wider range of optimal bargaining choices. These are generically inefficient regardless of whether information is complete or not.

Appendix

This appendix contains five technical propositions supporting some statements that were made in the text without proof.

***Proposition 1*:** If (7) is included in strategy ψ_i then Assumption A1 is satisfied for i.

***Proof*:** Assume that h^t contains infinitely many i-events. Consider two distinct such events α_i^s and α_i^τ with $s < \tau \leq t$. Clearly $\mathbb{P}(h_t) \leq \mathbb{P}(\alpha_i^\tau | \alpha_i^s)$. But if there are infinitely many such event times in a finite time interval: $\forall \epsilon > 0\ \exists s, \tau : \tau - s < \epsilon$. So $\mathbb{P}(h^t) < \mathcal{O}(\epsilon)$, and therefore $\mathbb{P}(h^t) = 0$. $\square$

***Proposition 2*:** Under Assumption 1, at any time t there is probability one that, for some $\theta > 0$, $(t, t+\theta)$ is a leg.

***Proof*:** If there fails to be such a θ, then there must exist, with probability $p > 0$, a decreasing sequence $(\theta_n)_{n\in\mathbb{N}}$ such that $(t + \theta_n)$ are event times and $\lim\limits_{n\to\infty} \theta_n = 0$. But then $h^{t+\theta_1}$ would contain infinitely many events with positive probability p. $\square$

***Proposition 3*:** If $\Delta F_i^t(0) = F_i^t(0^+) - F_i^t(0) > 0$ and $(t, t+\theta)$ is a leg then

$$\int_{[0,\theta)} e^{-r_i s} u_i(\xi_j^t(s)) dF_i^t(s) = \Delta F_i^t(0) u_i(\xi_j^t(0)) + \int_{(0,\theta)} e^{-r_i s} u_i(\xi_j^t(s)) dF_i^t(s)$$

***Proof*:** For any $\epsilon \in (0, \theta)$ define the integrable function

$$\nu_\epsilon(s) = \begin{cases} u_i(\xi_j^t(0)) & \text{if } s \in [0, \epsilon] \\ e^{-r_i s} u_i(\xi_j^t(s)) & \text{if } s \in (\epsilon, \theta) \end{cases}$$

One then has (for instance by the Dominated Convergence Theorem)

$$\int_{[0,\theta)} e^{-r_i s} u_i(\xi_j^t(s)) dF_i^t(s) = \lim_{\epsilon\to 0} \int_{[0,\theta)} \nu_\epsilon(s) dF_i^t(s)$$
$$= \lim_{\epsilon\to 0} \Big(u_i(\xi_j^t(0))(F_i^t(\epsilon) - F_i^t(0)) + \int_{(\epsilon,\theta)} e^{-r_i s} u_i(\xi_j^t(s)) dF_i^t(s) \Big) \qquad \square$$

***Proposition 4*:** If $(\theta_n)_{n\in\mathbb{N}}$ be the (increasing) sequence of event times on path σ^t then $\lim\limits_{n\to\infty} e^{-r_i \theta_n} \Phi^t(\theta_n) = 0$.

***Proof*:** If the sequence is finite or empty let us add to it an arbitrary sequence increasing to infinity. If $\lim\limits_{n\to\infty} \theta_n = \infty$ the result is obvious. Otherwise $\lim\limits_{n\to\infty} \theta_n = \Theta < \infty$. But then $\Phi^t(\Theta) = 0$ or the finite history $h^{t+\Theta}$ would contain infinitely many events with positive probability. Thus, by left-continuity of Φ^t, $\lim_{n \to \infty} \Phi^t(\theta_n) = 0$. ☐

***Proposition 5*:** Assume that $g : [0, \theta] \to \mathbb{R}$ is strictly positive and continuous. Let $\mathcal{F} = \{f : [0, \theta] \to [0, 1];$ with f strictly monotonous$\}$. And let $H[f] = \int_0^\theta g(s) df(s)$. Then an extremum of H is reached in $\mathcal{F}$ if and only if g is constant on $[0, \theta]$.

***Proof*:** Assume, for instance, that f is strictly increasing and yields a maximum of H on $\mathcal{F}$. If g is not constant one can find a non-empty interval $(a, b) \subset [0, \theta]$ such that

$$\sup_{s \in (a,b)} g(s) = \gamma < g(\tau) = \max_{s \in [0,\theta]} g(s)$$

Assume for instance that $\tau > b$ (if $\tau < b$ or $\tau < a$ the construction of f^* is similar) and let

$$f^*(s) = \begin{cases} f(s) & \text{for } s \in [0, a] \\ \frac{1}{2}(f(a) + f(s)) & \text{for } s \in (a, b) \\ f(s) + \frac{1}{2}(f(a) - f(b)) & \text{for } s \in [b, \tau] \\ f(s) & \text{for } s \in (\tau, \theta] \end{cases}$$

One easily verifies that $f^* \in \mathcal{F}$ and that

$$H[f^*] \geq H[f] + \tfrac{1}{2}(f(b) - f(a))(g(\tau) - \gamma) > H[f]$$

So f can't yield a maximum of H. If g is constant the result is obvious. ☐

References

Busch, L., and Q. Wen (1995): "Perfect Equilibria in a Negotiation Model," *Econometrica*, 63:545-565.

Coles, M., and A. Muthoo (2003): "Bargaining in a Non-Stationary Environment," Journal of Economic Theory 109:70-89.

Cramton, P.C. (1992): "Strategic Delay in Bargaining with Two-Sided Uncertainty," *Review of Economic Studies* 59:205-225.

De Vos, M., and D. Vermeir (2002): "Dynamic Decision Making in Logic Programming and Game Theory," Proceedings AI2002: Advances in Artificial Intelligence: 36-57.

Laraki, R., Solan E., and N. Vieille (2005): "Continuous-Time Games of Timing," *Journal of Economic Theory* 120:206-238.

Perry, M., and P. Reny (1993): "A Non-Cooperative Bargaining Model with Strategically Timed Offers," *Journal of Economic Theory* 59:50-77.

Poole, D., (1997): "The Independent Choice Logic for Modelling Multiple Agents under Uncertainty," *Artificial Intelligence* 94:5-56.

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

Sakovics, J. (1993): "Delay in Bargaining Games with Complete Information," *Journal of Economic Theory* 59:78-95.

Sandholm, T. (2002): "Bargaining with Deadlines," manuscript.

Smith, L., and E. Stacchetti (2003) "Aspirational Bargaining," manuscript.

Stahl, D.O. (1972): "Bargaining Theory," Stockholm Research Institute.

Wang, R., (2000): "Separating Equilibria in a Continuous-Time Bargaining Model with Two-Sided Uncertainty," *International Journal of Game Theory* 29:229-240.

Footnotes

[*]Mathematics Department, San Francisco State University.

Email: langlois@math.sfsu.edu, web: http://userwww.sfsu.edu/~langlois.

[†]This work has been partially supported by NSF grant IIS-0439034

[1]The logical instructions can be viewed as two computer programs interacting with each other as in "Agent Oriented Programming." There is a growing literature on Logic and Game Theory that investigates such issues. See De Vos and Vermeir (2002) and Poole (1997).

[2]Formally, the history h_t is generated by the strategy profile (ψ_i, ψ_j) if for any $s \geq 0$ and $\alpha^s = \left(\alpha_i^s, \alpha_j^s\right) \in h_t$, $\alpha_i^s = \psi_i(h_s)$ and $\alpha_j^s = \psi_j(h_s)$ with $h_s = \{\alpha^\tau \in h_t : \tau < s\}$.

[3]Any $\mathcal{O}(\epsilon)$ that approaches zero with ϵ can be used. It could be defined so that it becomes significant for only very small values of ϵ, so that the probability is 1 in most cases.

[4]There are reasonable alternatives to this assumption that would not affect the results. For instance, one may rule that the offer $\xi_j(t)$ is accepted or that x_i is changed to $(1 - \xi_j(t))$.

[5]In the standard definition a distribution function is *right*-continuous. In this definition it is *left*-continuous since $\mathcal{A}_i^{\mathcal{L}}([0,t))$ refers to an event *prior* to time t.

[6]It is assumed that probabilistic acceptance always yields a_i or $\neg a_i$ rather than $\emptyset_i$.

[7]$U_i(\sigma^t(\tau)) = \begin{cases} -c_i(z_i, z_j) & \text{the current costs if neither side has accepted} \\ u_i(x) & \text{if either side has accepted and } x \text{ is } i\text{'s share} \end{cases}$

[8]It is standard to denote by $\int_{[0,\theta)}$ a Stieltjes integral in order to make the interval of integration (here closed at 0, open at θ) unambiguous. Integration by parts applies since $\int_0^s -r_i e^{-r_i\tau} c_i\left(\zeta_i^t(\tau), \zeta_j^t(\tau)\right) d\tau$ is continuous in s and $F_i^t(s)$ is a distribution function.

[9]Note that i observes $\xi_{\mathcal{J}}^t(s)$ and $\zeta_{\mathcal{J}}^t(s)$ rather than individual $\xi_j^t(s)$ and $\zeta_j^t(s)$ in that case.

[10]Discontinuities of F_i^t and F_j^t may occur at 0^+ or θ^+, but *not* at θ^-, by definition here.

[11]It also simplifies the writing of (10) when path σ^t involves the probabilistic *timing* of offers or cost controls. If θ is the random date of an event other than acceptance, with distribution function $f(\theta)$, then i's expected payoff $\widetilde{E}_i(\sigma^t)$ satisfies

$$\widetilde{E}_i(\sigma^t) = \int_{[0,\infty)} \left(\int_{[0,\theta)} e^{-r_i s} dG_{ij}^t(s) + e^{-r_i\theta} \Phi^t(\theta) \widetilde{E}_i(\sigma^{t+\theta}) \right) df(\theta)$$

[12]This is sufficient but not necessary since it is usually possible to construct equilibria from suboptimal behavior in $[0, \theta)$ by resorting to adequate threats in $[\theta, \infty)$.

[13]Indeed we assumed in A2 that illegal acceptance results in a "rejection" event by both. So, a discontinuity in ϕ_i^t must yield an event whether the acceptance is legal or not. Also note that ϕ_i^t extends continuously to $[t, t+s]$ since it is monotonous and bounded.

[14]$\xi_i^t(0)$ can be different from $\xi_i^t(0^+) \equiv x_i^t$ if i makes a take-it-or-leave-it offer $\xi_i^t(0)$ at time t.

[15]Although this instruction is sent at all times, it can be non-null only once, when j's offer is actually accepted. Thus the strategy is not explosive.

[16]With Λ_i as defined in (13) and the notation $x_i^t = \xi_i^t(0^+)$, etc. Clearly $\mu_j \geq 0$ is constant between event times and $\mu_j = 0$ can only occur if costs are nil and $x_j^t = 0$.

[17]i's *non-null* offer can never exceed $\left(1 - \ x_j^*\right)$ by A2.

[18]If AP does not hold, non-acceptance can yield the same expected payoff as countervailing and the conclusion of (ii) does not holds. If $0 < \nu < \theta$ then $E_i(\sigma^{t+\nu}) = u_i(x_j^t)$ by continuity.

[19]j may have an expectation about the shape of ϕ_i^{t+s} but can only observe a_i or $\neg a_i$.

[20]The only difference between ϕ_i^t and φ_i^t is that $\mathcal{A}_i[t, t+s)$ can include an event at time t. As a result φ_i^t may be discontinuous at $s = 0$.

[21]Under condition CA, and for any $i \in \mathcal{I}$, $S_j^{t+s} = \mathcal{A}_j^{\mathcal{L}}[t, t+s)$, $\mathbb{P}^t(S_j^{t+s}) = \varphi_i^t(s)\varphi_j^t(s)$. But $\varphi_i^t(s)$ drops out from (33).

[22]Bayes' Law is the only belief-updating process consistent with dynamic programming.

[23]Any tie-breaking rule can be used in case there are several maximizands.

[24]With the convention $-ln0 = \infty$.

[25]Because $\xi_{\mathcal{J}} \equiv 0$ and by (36)

$$\Delta G_{i\mathcal{J}}^t(0) + \left(\varphi_i^t(0^+) - e^{-r_i\theta}\varphi_i^t(\theta)\varphi_{\mathcal{J}}^t(\theta)\right)u_i(x_{\mathcal{J}}^t) = \left(1 - e^{-r_l\theta}\varphi_l^t(\theta)\varphi_{\mathcal{J}}^t(\theta)\right)u_l(0)$$

[26]A change in the active type on either side is not considered an event.