

MARKOV BARGAINING GAMES

Martin W. Cripps*
University of Warwick,
Coventry CV4 7AL, UK.

First version September 1993

This version January 1997

ABSTRACT: I consider an alternating offer bargaining game which is played by a risk neutral buyer and seller, where the value of the good to be traded follows a Markov process. For these games the existence of a perfect equilibrium is proved and the set of equilibrium payoffs and strategies are characterised. The main results are: (a) If the buyer is less patient than the seller, then there will be delays in the players reaching an agreement, the buyer is forced into a suboptimal consumption policy and the equilibrium is ex-ante inefficient. (b) If the buyer is more patient than the seller, then there is a unique and efficient equilibrium where agreement is immediate.

JEL Classification : 026

Keywords : Bargaining, Uncertainty

* This paper has benefited from three very careful referees who greatly improved the presentation and eliminated a number of major errors in the proof of the first proposition. I am grateful to Klaus Schmidt, Avner Shaked and to the University of Bonn for their hospitality and support whilst I completed this work. And I must thank Saul Jacka and seminar participants at Birmingham, Erasmus, Exeter and Warwick for their comments.

1. Introduction

How does one player bargain with another over the price of a good if the value of the good follows a stochastic process? In the model below the players have an indivisible good which they bargain over by making alternating offers. The value of this good evolves through time, it is determined by a Markov chain and the players observe the value of the good at the start of each period's bargaining. The fact that the future value of the good is random gives a possible benefit to waiting, because the value may grow in the future, but makes an agent's decisions more complex.

This problem arises if, for example, the owner of oil reserves is bargaining over their price with an oil company and the future price of oil follows a stochastic process. Here there is a cost to delay, because of the players' rates of time preference, but at the same time the oil company will not choose to extract the reserves immediately but will generally wait until the oil price has reached a threshold level. I will show that, if the oil company is more patient than the seller, then the reserves will be sold immediately and the oil company chooses when to extract the reserves. In this case there is a unique equilibrium. Moreover, the bargain between the oil company and the seller is determined by each player's ability to delay agreement long enough to upset the oil company's optimal extraction strategy. The bargain reached is a natural generalisation of the deterministic solution of Rubinstein (1982) and the equilibrium is efficient. If the seller is more patient than the buyer then there can be delays in reaching agreement, because the seller cannot force the impatient oil company to follow a good extraction policy. As a consequence, at the equilibrium of the bargaining game the seller will delay the agreement, to force the company to follow an extraction policy that suits the seller. The oil company then extracts the oil immediately after a bargain is agreed with the seller. The equilibrium in this case bears no simple relationship to the Rubinstein (1982) solution and in this case the equilibrium is not generally efficient.

There are two classes of conclusions we can draw from this model: those relating to bargaining games, and those relating to the optimal waiting literature. First, there is a class of complete information bargaining games where delays in reaching agreement do occur, these delays are sometimes efficient but can impose costs on players. This delay

does not arise because of any signalling (as in Admati and Perry (1987)), nor does it occur because of any complex strategic effects (as in Perry and Reny (1993) or in Muthoo (1990)), but is simply because of the benefit that both players perceive in waiting. Second there is a class of models where it is optimal to wait before exploiting a resource, but where the observed delay is not in general equivalent to the amount of delay observed in a one person model.

The paper proceeds in the following way. In Section 2 there is a description of the stochastic process generating the uncertainty in our model and an outline of the bargaining game. Section 3 characterises the general solution to the bargaining problem and establishes the existence of an equilibrium.

2. The Model

The first element of the model is a stochastic process, Π , which determines the evolution of a random sum of money and I will call this variable the “cake”. The process is a homogeneous Markov chain with countable states $N := \{1, 2, \dots\}$. The function $f(i)$ gives the size of the cake in state $i \in N$, the state at time t is denoted i_t . States are ordered so the largest cake occurs when $i = 1$ and the cake size monotonically converges to zero as i increases, that is, $1 = f(1) > f(i) > f(i + 1) \geq 0$ for $i \in N$ and $\lim_{i \rightarrow \infty} f(i) = 0$. The transition probabilities for Π are denoted $\{\pi_{ij}\}_{i \in N, j \in N}$, where $\sum_{j \in N} \pi_{ij} = 1$ for all $i \in N$ and $\pi_{ij} \geq 0$ for all ij . For each $i \in N$, the probability distribution $\{\pi_{ij}\}_{j \in N}$ gives the probability of a transition from state i today to the state $j \in N$ in the next period. For example, if the cake is size $f(i)$ in period t , then the expected size next period is $\sum_{j \in N} \pi_{ij} f(j)$. The matrix which has $\{\pi_{ij}\}_{j \in N}$ as its i th row is denoted P . To simplify notation, I denote the i th row of P as the vector P_i . Therefore, $P_i \phi := \sum_j \pi_{ij} \phi(j)$ (respectively $P_i P \phi := \sum_j \sum_k \pi_{ij} \pi_{jk} \phi(j)$) gives the expected value in one period (respectively two periods) of the random variable $\phi : N \rightarrow \mathfrak{R}$ when the current state is i .

A *one person optimal stopping problem* arises if one individual can eat all of the cake themselves, but must decide the right moment to eat it. Formally, one can state the problem as: what is the best time to stop the process $\{\delta^t f(i_t)\}_{t=1}^{\infty}$ (where $\delta \leq 1$ is the player’s discount factor)? A strategy $\tau(\cdot)$ which determines when to stop the process as a

function only of past and current states is called a stopping time. The expected payoff from adopting the optimal stopping time in state i is the *value function*: $v_\delta(i) := \sup_\tau E_i \delta^\tau f(i_\tau)$, (here the supremum is taken over all stopping times, with the convention that $f(i_\infty) = 0$ and E_i denotes expectations taken from an initial state i). The optimal strategy (if it exists) is denoted τ^* . The value function for this problem is uniquely characterised as the smallest solution to the equation

$$(1) \quad v_\delta(i) = \max\{ f(i) , \delta P_i v_\delta \}, \quad \forall i \in N.$$

That is, if the function $h(i)$ also satisfies $h(i) = \max\{f(i), \delta P_i h\}$ then $v_\delta(i) \leq h(i)$ for all $i \in N$.¹ The Hamilton-Jacobi-Bellman equation (1) says that if the optimal strategy is being played in state i , then the maximum expected payoff is obtained either by eating today's cake, or by waiting until next period and then playing the optimal strategy. The optimal strategy appears to be easy to compute (if $v_\delta(i) = f(i)$ stop and if $v_\delta(i) > f(i)$ continue), however, such a strategy may not always achieve the payoff $v_\delta(i)$.² Below, I prove a result that ensures the existence of an optimal strategy for the Markov chain Π , which is described above. The result follows from the special structure of the function $f(\cdot)$: (a) $f(\cdot)$ is bounded above, (b) the only accumulation point of $f(\cdot)$ is at zero.³

LEMMA 1 *If Π is an irreducible aperiodic Markov chain and $\delta \leq 1$, then the stopping time τ^* , defined by the first time i_t enters the set $G := \{i | f(i) = v_\delta(i)\}$, satisfies $v_\delta(i) = E_i \delta^{\tau^*} f(i_{\tau^*})$ on the Markov chain Π .*

Proof: See the Appendix.

I now describe a bargaining game where the players bargain over the random cake described above. The game is played by a seller, “she”, and a buyer, “he”, of a randomly

¹In general there are many solutions to (1), for example when $\delta = 1$ the function $h(i) = 1$ for all $i \in N$ is a solution, so it is essential to choose the smallest solution to (1).

²In general there only exists a strategy yielding a payoff within $\epsilon > 0$ of $v_\delta(i)$: stop if $v_\delta(i) - \epsilon < f(i)$ and if not continue.

³The proposition makes some assumptions on the form of the Markov process Π . A Markov process is irreducible if it is possible to move from any one state to any other in a finite number of steps, that is all states *communicate* with all other states. An irreducible Markov chain cannot have absorbing states. Relaxing the assumption of irreducibility will complicate the results, but all the results below will apply *mutatis mutandis* to any irreducible component of the Markov chain. A sufficient condition for the Markov chain to be aperiodic is for $\pi_{ii} > 0$ for all $i \in N$.

valued asset, once the buyer has obtained the rights to the asset he can consume it whenever he wishes. Both players have a reservation level of zero. Play proceeds as follows: In the first period both players observe the size of the cake, $f(i_1)$, and then the seller suggests a deal. She proposes a sum of money x_1 for which she is willing to transfer her rights over the process to the buyer. The buyer then accepts or rejects her proposal. If he accepts, then the bargaining terminates. If he rejects, then the game moves into the second period where the random variable $f(i_2)$ is observed and it is the buyer's turn to propose a deal. If this division is accepted then the bargaining ends, if not play continues to a third period. Play continues with alternating offers until an agreement is reached. The new feature is that, at the start of period t the players observe a random variable which is the value of the good in that period, however, if they do not agree today, they do not know the size of the cake they will bargain over in the next period.

Let $0 < \gamma \leq 1$ be the seller's discount factor and let $0 < \delta \leq 1$ be the buyer's discount factor.⁴ If there is an agreement at time t in state i where the buyer pays x_t , then the seller and buyer's payoffs are respectively $(\gamma^t x_t, \delta^t (v_\delta(i) - x_t))$. If the players bargain forever and never reach an agreement then they both get a payoff of zero. The buyer's payoff is, of course, determined by his ability to consume the cake in an optimal state once agreement is reached. I have assumed that the players are risk neutral so that the surplus they divide is linear. A pure strategy for each player in the game is a function mapping the history of the state, offers and rejections to an action today. Let $I := [0, 1]$ be the set of possible offers a player can make and let $R := \{Y, N\}$ be the set of responses to these offers, then a complete description of the events in a given period is an element of $H := N \times I \times R$. A history to period t is an element of the set H^{t-1} . Thus, a strategy for the seller is a sequence of functions $\sigma := \{\sigma_t\}_{t=1}^\infty$ such that: (a) for t odd $\sigma_t : H^{t-1} \times N \rightarrow I$, (b) for t even $\sigma_t : H^{t-1} \times N \times I \rightarrow R$. Similarly, a strategy for the buyer is a sequence of functions $\rho := \{\rho_t\}_{t=1}^\infty$ such that: (a) for t even $\rho_t : H^{t-1} \times N \rightarrow I$, (b) for t odd $\rho_t : H^{t-1} \times N \times I \rightarrow R$. The equilibrium concept used throughout is that of subgame perfect equilibrium.

⁴It is not necessary for the discount factors to be strictly less than unity, because the stochastic process Π can make the cake shrink even when $\delta = \gamma = 1$.

3. Equilibrium in the Bargaining Game

This section presents the main results of the paper. First some general results are given. The first result in this section is a characterisation of the equilibrium payoffs of the bargaining game which owes a great deal to the method pioneered in Shaked and Sutton (1984), then an existence result is established in Proposition 1. The last general result shows that the players will always agree a bargain in a state i , if they both believe that the expected future value of the good is less than its current value in state i . Then the two cases, (a) where the buyer is more patient than the seller, and (b) where the seller is more patient than the buyer, are treated separately. Proposition 3 proves that if the buyer is more patient than the seller then the game has a unique equilibrium payoff in each state i . At this equilibrium an agreement is reached in the first period. An example then shows that there can be delay in reaching an agreement if the seller is more patient than the buyer. The final proposition shows that if the seller is more patient, then the buyer will choose to consume the good immediately agreement is reached, that is, the timing of agreement in the bargaining determines the buyer's consumption strategy.

First some additional notation is needed. Define $a(i)$ (respectively $A(i)$) to be the infimum (respectively supremum) of the set of all subgame perfect equilibrium expected payoffs for the seller if she is the proposer and the subgame begins in state i . Similarly, define $b(i)$ and $B(i)$ to be the bounds on the buyer's equilibrium payoffs when he is the proposer in a subgame beginning in state i . The structure of the game is stationary, so these bounds also apply after any history leading to state i .

LEMMA 2 *If Π is an irreducible aperiodic Markov chain, then*

$$(2) \quad a(i) = \max\{v_\delta(i) - \delta P_i B, \gamma^2 P_i P a\},$$

$$(3) \quad A(i) = \max\{v_\delta(i) - \delta P_i b, \gamma^2 P_i P A\},$$

$$(4) \quad b(i) = \max\{v_\delta(i) - \gamma P_i A, \delta^2 P_i P b\},$$

$$(5) \quad B(i) = \max\{v_\delta(i) - \gamma P_i a, \delta^2 P_i P B\}.$$

Where $a(i)$ is the smallest solution to (2), $A(i)$ is the smallest solution to (3), $b(i)$ is the smallest solution to (4), and $B(i)$ is the smallest solution to (5). That is: if $h(i) = \max\{v_\delta(i) - \delta P_i B, \gamma^2 P_i P h\}$ then $h(i) \geq a(i)$ for all

$i \in N$, if $h(i) = \max\{v_\delta(i) - \delta P_i b, \gamma^2 P_i P h\}$ then $h(i) \geq A(i)$ for all $i \in N$,
if $h(i) = \max\{v_\delta(i) - \gamma P_i A, \delta^2 P_i P h\}$ then $h(i) \geq b(i)$ for all $i \in N$, and if
 $h(i) = \max\{v_\delta(i) - \gamma P_i a, \delta^2 P_i P B\}$ then $h(i) \geq B(i)$ for all $i \in N$.

Proof: See the Appendix.

The lemma above characterises the worst equilibrium payoff and the best equilibrium payoffs of the players. The next proposition proves the existence an equilibrium. Existence is proved by constructing a Markovian, or stationary, equilibrium where the players use strategies which depend only on the current state and not on previous events.

PROPOSITION 1 *If Π is an irreducible aperiodic Markov chain, then there exists a subgame perfect equilibrium of the bargaining game.*

Proof: See the Appendix.

I now give a general result on the states in which the players are certain to come to an agreement. The condition used in the proposition below, $v_\delta(i) = v_\gamma(i) = f(i)$, says that in state i both of the players agree that it is not worth waiting for the cake to grow in the future and it is better to eat it now. If in state i even the most patient of the two players agrees that it is not worth waiting for the cake to grow, then the proposition shows that the players reach an agreement in this state. This result is therefore the stochastic equivalent of the result that says agreement is immediate when the cake shrinks.

PROPOSITION 2 *If $v_\delta(i) = v_\gamma(i) = f(i)$ for some state i , then the players reach an agreement in state i .*

Proof: See the Appendix.

This completes the general results. I will now study the case where the buyer is more patient than the seller. I will show that in this case in each state i the players have unique equilibrium payoffs that are efficient. I will also show that these payoffs can

be obtained at a Markov perfect equilibrium where the players agree immediately and the buyer then follows his optimal stopping policy. I will argue that this equilibrium is the natural extension of the Rubinstein solution to a stochastic environment, where the players' bargaining strengths are completely determined by their ability to delay agreement. Moreover, the form of the solution is identical to Rubinstein (1982), provided that one interprets the terms γP_i and δP_i as the seller and buyer's *random* discount factors. These results contrast with what is true when the buyer is less patient than the seller ($\gamma > \delta$), in this case the equilibria of the game generally exhibit delays before the players come to an agreement. Also, I show that the equilibrium when the buyer is less patient does not implement the buyer's optimal stopping policy, and that it is not a simple generalisation of the Rubinstein solution.

The proposition below provides a complete characterisation of the equilibria for a large class of bargaining games. It shows that, if the seller is less patient than the buyer, then she will transfer the good immediately to the buyer. This allows him to receive the full value from the optimal stopping policy. This will also maximise both of the players' potential surplus, because if the seller is relatively myopic any delay in the transfer of control will impose costs on both parties (the seller bears a cost because she is forced to wait for the receipt of the cash and the buyer bears a cost because he fails to follow the optimal exploitative strategy). The equilibrium is, therefore, efficient.⁵

PROPOSITION 3 *If $\delta \geq \gamma$ and either: (a) $\gamma < 1$, or (b) $\gamma = \delta = 1$ and Π is an irreducible process with all states transient, then all states $i \in N$ have unique equilibrium payoffs to the two players. This unique equilibrium payoff can be achieved at a Markov perfect equilibrium where there is no delay in reaching agreement. If the unique equilibrium payoffs to the seller are denoted $\alpha(i)$ and the buyer's unique equilibrium payoffs are denoted $\beta(i)$, then these functions are determined by the equations*

$$(6) \quad \alpha(i) - \gamma \delta P_i P \alpha = v_\delta(i) - \delta P_i v_\delta,$$

$$(7) \quad \beta(i) - \gamma \delta P_i P \beta = v_\delta(i) - \gamma P_i v_\delta.$$

⁵In part (b) of Proposition 3 I allow $\gamma = \delta = 1$ (so both players are infinitely patient), but the process Π is transient, which implies that $f(i_t) \rightarrow 0$ almost surely. This case, therefore, arises when the value of the cake shrinks because of the nature of the random process rather than the players' impatience.

Proof: See the Appendix.

The solutions (6) and (7) for the unique equilibrium payoffs in the game, derived in the proof above, have a simple interpretation. If the terms $(\delta_2 =)\delta P_i$ and $(\delta_1 =)\gamma P_i$ are interpreted as random levels of discounting and $v_\delta(i)$ is interpreted as the notional size of the cake in state i , then (6) and (7) are analogous to the solution of the deterministic bargaining problem. Since, with an abuse of notation, we can write the seller's payoff in state i as $v_\delta(i)(1-\delta_2)/(1-\delta_1\delta_2)$ and the buyer's payoff in state i as $v_\delta(i)(1-\delta_1)/(1-\delta_1\delta_2)$. Therefore, if the buyer is more patient than the seller, the outcome of the bargaining game is *qualitatively* equivalent to that in the deterministic case — there is a unique solution with generalised bargaining strengths. The only difference is the replacing of the size of the cake with the optimal stopping policy of the buyer v_δ .

COROLLARY 1 *If $\delta \geq \gamma$ and either (a) $\gamma < 1$, or (b) $\gamma = \delta = 1$ and Π is an irreducible process with all states transient, then the bargaining game has unique equilibrium payoffs and the equilibrium is efficient.*

Proof: See the Appendix.

I cannot explicitly solve for the equilibria for when $\gamma > \delta$, however, it is certain that the equilibria differ from the solution given above. This can be easily verified, since the functions $\alpha(i)$, $\beta(i)$ calculated in (6) and (7) in general will not satisfy $\alpha(i) \geq \delta^2 P_i P \alpha$ and $\beta(i) \geq \gamma^2 P_i P \beta$ when $\gamma > \delta$. (Of course, if γ is sufficiently close to δ and Π is chosen suitably, then an identical equilibrium can be found.) A second difference is that there will be delays in reaching agreements when $\gamma > \delta$. This is shown in the following example.

Example : Delay in agreements when $\gamma > \delta$

In this example: (a) the seller is more patient than the buyer ($\gamma > \delta$), (b) in period $t = 1$ the cake is size $1/2$ ($f(i_1) = 1/2$), (c) in all future periods the cake is unity with probability one, ($f(i_t) = 1$ for $t = 2, 3, \dots$). Now consider the proposal made by the seller in period 1. If she sells the cake today, then the buyer must choose between having a cake size $1/2$ today or waiting and getting δ in a period's time. (If the buyer is impatient ($\delta < 1/2$) he will choose to eat the cake today and not wait.) From period 2 onwards

the cake is size one, so the standard Rubinstein solution applies and, as the buyer is the proposer in period 2, the players' payoffs at the perfect equilibrium beginning in period 2 are $(\gamma(1 - \delta)/(1 - \delta\gamma), (1 - \gamma)/(1 - \delta\gamma))$ for the seller and the buyer respectively. If the buyer rejects the seller's proposal in the first period, then his expected payoff from the next period is $\delta(1 - \gamma)/(1 - \delta\gamma)$, so the highest payoff the seller can get from an agreement in the first period of bargaining is $\delta - \delta(1 - \gamma)/(1 - \delta\gamma)$ when $\delta \geq 1/2$ and $1/2 - \delta(1 - \gamma)/(1 - \delta\gamma)$ if $\delta < 1/2$. However, if the seller waits until period 2 she can get a payoff of $\gamma^2(1 - \delta)/(1 - \delta\gamma)$ by accepting the buyer's proposal in the second period. It is easy to see that if $\delta < \gamma$, then the seller prefers to wait until period 2 to trade. (Notice that, although the cake has a deterministic path in this example, the result will also apply to stochastic model.)

This example shows that the players do not agree immediately when the buyer is less patient than the seller. The proposition below extends the example by showing that, if $\gamma > \delta$, the bargaining only reaches agreement in states i where $f(i) = v_\delta(i)$, or where $f(i) < v_\delta(i)$ and the seller receives a price of zero. The condition $f(i) = v_\delta(i)$ implies that the buyer will choose to consume the cake immediately after the bargaining is finished, so there is no delay between trade and consumption. One can interpret this outcome as the seller delaying agreement to force the buyer to follow a consumption strategy which suits her objectives. However, the seller will not force the buyer to mimic her consumption policy completely, because although agreement is reached in states where the buyer expects the cake to shrink, it is not necessary for the (more patient) seller to expect the cake to shrink at an agreement ($v_\gamma(i) = f(i)$). The reason this does not happen is because the buyer is willing to compensate the seller for an early agreement in states where she would rather delay, $v_\delta(i) = f(i) < v_\gamma(i)$.

PROPOSITION 4 *If $\gamma > \delta$, then an agreement in bargaining in state i implies that either: $v_\delta(i) = f(i)$, or $v_\delta(i) > f(i)$ and the seller's expected payoff in state i is zero.*

Proof: See the Appendix.

This proposition does not say that $v_\delta(i) = f(i)$ is a sufficient condition for agreement,

so there may be states where the impatient buyer wants to consume the cake but where the seller does not agree. That is, from the buyer's point of view the bargaining does not follow an optimal extraction policy (see the above example for an instance of this). The example below will show that the equilibria need not result in agreement at the seller's optimal stopping time, because the buyer is willing to compensate the seller for early consumption. Thus, when $\gamma > \delta$, the bargaining game need not result in the cake being consumed in a state that is optimal for either of the players as individuals.

Example : The seller coming to an early agreement when $\gamma > \delta$

Suppose that $\gamma > \delta$ and that the cake is size γ in period $t = 1$, but that in periods $t = 2, 3, \dots$ the cake is size unity. Unlike the previous example, suppose that the buyer makes the first offer, so the unique subgame perfect equilibrium in $t = 2$ gives payoffs $((1 - \delta)/(1 - \delta\gamma), \delta(1 - \gamma)/(1 - \delta\gamma))$ to the seller and the buyer respectively. Now consider the buyer's proposal in the first period. If the seller agrees to his proposal, then he obtains the cake in period 1 and he will not wait until period 2, because the current size of the cake exceeds its future discounted value to him, $\gamma > \delta$. The smallest offer that the seller is willing to accept in period 1 is $\gamma(1 - \delta)/(1 - \delta\gamma)$. The buyer is willing to make such an offer if his payoff from agreement in period $t = 1$ exceeds that from waiting until period $t = 2$, that is, if $\gamma - \gamma(1 - \delta)/(1 - \delta\gamma) > \delta^2(1 - \gamma)/(1 - \delta\gamma)$. This is true. Thus, although the seller's optimal time to eat the cake is in period 2, she is willing to trade before this time because the buyer is impatient and willing to reward the seller for earlier agreement. (The calculations above also go through if the cake in period one is $\gamma - \epsilon$, for ϵ sufficiently small, and in this case the seller would strictly prefer to delay eating the cake until period two.)

4. Conclusion

I have characterised the equilibria of the Markov bargaining games. I have completely solved for the unique equilibria when $\delta \geq \gamma$. I have also described some features of the equilibria in the remaining cases.

References

- ADMATI, A.R., AND M. PERRY, 1987, Strategic delay in bargaining, *Review of Economic Studies*, 54, 345-364.
- DYNKIN, E.B., 1960, *The theory of Markov processes*, (Pergamon Press, Oxford).
- KEMENY, J.G., J.L.SNELL AND A.W.KNAPP, 1976, *Denumerable Markov chains*, (Springer Verlag, Berlin).
- MUTHOO, A., 1990, Bargaining without commitment, *Games and Economic Behavior*, 2, 291-297.
- NEVEU, J., 1975, *Discrete parameter martingales*, (North Holland, Amsterdam).
- PERRY, M., AND P.J. RENY, 1993, A non-cooperative model of bargaining with strategically timed offers, *Journal of Economic Theory*, 59, 78-95.
- RUBINSTEIN, A., 1982, Perfect equilibrium in a bargaining model, *Econometrica*, 50, 97-109.
- SHAKED, A., AND J. SUTTON, 1984, Involuntary unemployment as a perfect equilibrium in a bargaining model, *Econometrica*, 52, 1351-64.

APPENDIX

PROOF OF LEMMA 1

I will begin by considering the case where $\delta = 1$ and the irreducible chain Π is recurrent. Assume that $G = \{1\}$, so that τ^* only stops in the maximal state. Since: (a) $\delta = 1$, (b) for recurrent chains there is a probability one of hitting the state $i = 1$, (c) $v_\delta \leq 1$, we can deduce that $v_\delta(i) = 1 = E_i \delta^{\tau^*} f(i_{\tau^*})$.

Now consider two cases case where Π is transient and $\delta = 1$, or where $\delta < 1$. Since $v_\delta(i)$ is bounded, the process $v_\delta(i_{\tau^* \wedge n})$ is a uniformly integrable martingale. From the definition of a martingale we have

$$(8) \quad v_\delta(i) = E_i \delta^{\tau^* \wedge n} v_\delta(i_{\tau^* \wedge n}) = E_i \delta^{\tau^*} v_\delta(i_{\tau^*}) 1_{(\tau^* \leq n)} + E_i \delta^n v_\delta(i_n) 1_{(\tau^* > n)}.$$

Provided I can show that the term $E_i \delta^n v_\delta(i_n) 1_{(\tau^* > n)}$ converges to zero as $n \rightarrow \infty$ I have established the result, because $E_i \delta^{\tau^*} v_\delta(i_{\tau^*}) 1_{(\tau^* \leq n)}$ converges to $E_i \delta^{\tau^*} f(i_{\tau^*})$ as $n \rightarrow$

∞ . If $\delta < 1$ this is immediate, because $v_\delta \leq 1$. The remaining case has $\delta = 1$ and Π transient. To show that $E_i \delta^n v_\delta(i_n) 1_{(\tau^* > n)}$ converges to zero in this case I will show that $v_1(i_n) \rightarrow 0$ almost surely as $n \rightarrow \infty$. First note that $f(i_t) \rightarrow 0$ almost surely if Π is transient. (By Theorem 4.28 p.102 Kemeny *et al* (1976), for any $\epsilon > 0$ and any $i \in N$ there exists a T such that $\Pr[\exists t > T \text{ s.t. } i_t < i \mid i_0] < \epsilon$, hence for any $\epsilon > 0$ and $\nu > 0$ there exists a T such that $\Pr[\exists t > T \text{ s.t. } f(i_t) > \nu \mid i_0] < \epsilon$.) Define the random variable $z_t = \sup_{m \geq t} |f(i_m)|$, then $\{z_t\}_{t=0}^\infty$ is a non-increasing sequence converging almost surely to zero. But now notice that

$$0 \leq v_1(i_t) \leq \sup_{s \geq t} |f(i_s)| = z_t,$$

and since the right hand side converges almost surely to zero, so too must the left. It follows that $E_i \delta^n v_\delta(i_n) 1_{(\tau^* > n)}$ converges to zero, by the dominated convergence theorem. *Q.E.D.*

PROOF OF LEMMA 2

Assume that there is an equilibrium such that: (1) After any history where the buyer is the proposer in state i , he proposes the payoffs $(v_\delta(i) - b(i), b(i))$ for the seller and the buyer respectively. (2) When the buyer is responder in state i he only accepts offers that give him a payoff of at least $\delta P_i b$. Now consider the seller's optimal response. If she is the proposer after some history, the buyer will accept any offer that gives him no less than his continuation payoff, that is, he will accept the offer $\delta P_i b$. If the seller makes this offer she will receive $v_\delta(i) - \delta P_i b$. The seller's problem is to decide when to make an offer that the buyer is willing to accept, that is, she faces an optimal stopping problem with a reward function $v_\delta(i) - \delta P_i b$ in state i . The seller is the proposer in alternate periods, so she must wait two periods before she can stop the process, so (by (1)) her value function to this optimal stopping problem is described by z the smallest solution to

$$(9) \quad z(i) = \max\{v_\delta(i) - \delta P_i b, \gamma^2 P_i P z\}.$$

By Lemma 1, if the seller offers $v_\delta(j) - z(j)$ to the buyer, when she is the proposer in state j , and accepts offers which give her a payoff of at least $\gamma P_j z$ in state j , then she will achieve the payoff $z(i)$ in state i . The largest possible equilibrium payoff to the seller in state i is $v_\delta(i) - \delta P_i b$, since the buyer would never accept an offer of less than $\delta P_i b$.

Thus, $z(\cdot) = A(\cdot)$, where $A(i)$ is the seller's largest possible equilibrium payoff in state i , provided the initial assumption on $b(\cdot)$ is correct.

If the seller uses the strategy (defined by $z(i)$) above, then what is the buyer's optimal response? The seller accepts proposals from the buyer if they offer her $\gamma P_i A$ in state i , so the buyer also faces an optimal stopping problem with $v_\delta(i) - \gamma P_i A$ as the reward in state i . His value function for this problem is u the smallest solution to

$$(10) \quad u(i) = \max\{v_\delta(i) - \gamma P_i A, \delta^2 P_i P u\}.$$

By Lemma 1, a strategy achieving the payoff $u(i)$ exists. This strategy proposes that the buyer receives $u(i)$ when he is proposer in state i and accepts offers no worse than $\delta P_i u$ in state i . But, $u(i)$ must be his worst possible equilibrium payoff in state i , because the seller can never expect to get more than $\gamma P_i A$ from future play, hence $u(i) = b(i)$. The strategies above give the buyer a payoff of $b(i)$ as proposer in state i and the seller a payoff $A(i)$ as proposer in state i and these strategies do form an equilibrium. The initial assertion is therefore correct. An identical argument will show that there is an equilibrium that supports the payoffs $B(i)$ and $a(i)$. *Q.E.D.*

PROOF OF PROPOSITION 1

At a Markov perfect equilibrium the players will play strategies which only depend on the current state. Let $\alpha(i)$ (respectively $\beta(i)$) give the seller's (respectively buyer's) expected equilibrium payoff if she (respectively he) were the proposer at a Markov perfect equilibrium in state i . Lemma 2 shows that $\alpha(i)$ and $\beta(i)$ satisfy

$$(11) \quad \alpha(i) = \max\{v_\delta(i) - \delta P_i \beta, \gamma^2 P_i P \alpha\},$$

$$(12) \quad \beta(i) = \max\{v_\delta(i) - \gamma P_i \alpha, \delta^2 P_i P \beta\}.$$

I will prove that there exist two functions $\alpha(i)$, $\beta(i)$ satisfying (11) and (12). First, define two sequences of functions $\{\alpha^n(i)\}_{n=1}^\infty$ and $\{\beta^n(i)\}_{n=1}^\infty$ recursively. Let $\beta^1(i) = 0$ and let $\alpha^1(i)$ be the smallest function satisfying $h(i) = \max\{v_\delta(i), \gamma P_i h\}$. Now define $(\alpha^{n+1}, \beta^{n+1})$ to be the smallest functions satisfying

$$\begin{aligned} \alpha^{n+1}(i) &= \max\{v_\delta(i) - \delta P_i \beta^n, \gamma^2 P_i P \alpha^{n+1}\}, \\ \beta^{n+1}(i) &= \max\{v_\delta(i) - \gamma P_i \alpha^n, \delta^2 P_i P \beta^{n+1}, 0\}. \end{aligned}$$

First notice that $\alpha^1 \geq \alpha^2$ and that $\beta^1 \leq \beta^2$. Also, given that $\beta^n \geq \beta^{n-1} \geq \dots \geq \beta^1$ and $\alpha^n \leq \alpha^{n-1} \leq \dots \leq \alpha^1$ the construction ensures that $\alpha^{n+1} \leq \alpha^n$ and $\beta^{n+1} \geq \beta^n$. By induction, therefore, $\{\alpha^n(i)\}$ is a decreasing sequence of functions bounded below by zero and $\{\beta^n(i)\}$ is an increasing sequence of functions bounded above by unity. These sequences converge, by monotone convergence, so there exists limits α and β satisfying (11) and (12).

I must also describe the strategies at a Markov perfect equilibrium to show that the functions $\alpha(i)$, $\beta(i)$, constructed above, are the equilibrium payoffs. By (11), α is the solution to a stopping problem with $v_\delta(i) - \delta P_i \beta$ as the reward in state i , where the seller can only stop the process in even periods. By Lemma 2, a stopping time exists that gives the expected payoff $\alpha(i)$ in state i . This stopping time is the optimal strategy for the seller. A similar strategy can be constructed for the buyer. *Q.E.D.*

PROOF OF PROPOSITION 2

Suppose the game has reached a state where $v_\delta(i) = v_\gamma(i) = f(i)$ and the seller proposes a bargain. Let $x(i)$ be her expected payoff in state i and let $y(j)$ be the function describing the buyer's payoff if he rejects her offer and play moves to state j . Suppose that the proposition is false, that is, agreement is not reached in state i and that the players strictly prefer to delay agreement, $\delta P_i y + x(i) > v_\delta(i)$.

The sum $x(i) + \delta P_i y$ gives the total of the players' expected payoffs from future play. Let the stopping time τ describe the times when players reach agreement after state i . In the states i_t where the players do agree a bargain their total payoff is $v_\delta(i_t)$. Suppose that both players discounted payoffs at the rate δ , then their total expected payoffs $(x(i) + \delta P_i y)$ could be calculated as $E_i[\delta^\tau v_\delta(i_\tau)]$, but instead we can find an upper bound on total expected payoffs, $x(i) + \delta P_i y \leq E_i[\kappa^\tau v_\delta(i_\tau)]$, where $\kappa = \max\{\delta, \gamma\}$. By combining the previous inequalities, and using the fact that $v_\delta(i) = v_\gamma(i) = v_\kappa(i) = f(i)$, we get

$$v_\delta(i) < x(i) + \delta P_i y \leq E_i[\kappa^\tau v_\delta(i_\tau)] \leq v_\kappa(i) = v_\delta(i),$$

where $E_i[\cdot]$ denotes an expectation conditional on the initial state of the stochastic process being i . This is a contradiction, so the players are willing to reach an agreement in state i . *Q.E.D.*

PROOF OF PROPOSITION 3

By Lemma 2 the bounds on the sets of equilibrium payoffs satisfy

$$\begin{aligned} a(i) + \delta P_i B &\geq v_\delta(i), & A(i) + \delta P_i b &\geq v_\delta(i), \\ b(i) + \gamma P_i A &\geq v_\delta(i), & B(i) + \gamma P_i a &\geq v_\delta(i). \end{aligned}$$

Moreover, there is an equilibrium where the seller receives the payoffs $A(i)$ (respectively $a(i)$) when the buyer receives payoffs $b(i)$ (respectively $B(i)$). The left hand side of these inequalities gives the sum of the players' expected payoffs in the states in which they agree. At the equilibrium $(a(i), B(i))$ in state i the total expected payoff to the players is $B(i) + \gamma P_i a$, in states when the buyer makes an offer. This can be estimated by: (1) Calculating the stopping time τ^* determined by the states in which the players expect to reach agreement. (2) Adding the players' payoffs at the states they reach agreement, to get $v_\delta(i_{\tau^*})$. (3) Discounting this back at the rate δ . This gives an estimate $E_i[\delta^{\tau^*} v_\delta(i_{\tau^*})]$. However, the estimate of total payoffs, $E_i[\delta^{\tau^*} v_\delta(i_{\tau^*})]$, overestimates the seller's payoffs because $\delta \geq \gamma$, thus

$$\begin{aligned} E_i[\delta^{\tau^*} v_\delta(i_{\tau^*})] &\geq a(i) + \delta P_i B \geq v_\delta(i), & E_i[\delta^{\tau^*} v_\delta(i_{\tau^*})] &\geq A(i) + \delta P_i b \geq v_\delta(i), \\ E_i[\delta^{\tau^*} v_\delta(i_{\tau^*})] &\geq b(i) + \gamma P_i A \geq v_\delta(i), & E_i[\delta^{\tau^*} v_\delta(i_{\tau^*})] &\geq B(i) + \gamma P_i a \geq v_\delta(i). \end{aligned}$$

But (1) implies that $v_\delta(i) \geq E_i[\delta^{\tau^*} v_\delta(i_{\tau^*})]$ for all $i \in N$, thus we have

$$v_\delta(i) = a(i) + \delta P_i B = A(i) + \delta P_i b = b(i) + \gamma P_i A = B(i) + \gamma P_i a.$$

It is possible to solve these equations for $a(i)$, $A(i)$, $b(i)$, $B(i)$ by substitution. If this is done one gets two equations: the one in $\alpha(i)$ is satisfied by both $A(i)$ and $a(i)$, the other in $\beta(i)$ is satisfied by $b(i)$ and $B(i)$.

$$(13) \quad \alpha(i) - \gamma \delta P_i P \alpha = v_\delta(i) - \delta P_i v_\delta, \quad \forall i,$$

$$(14) \quad \beta(i) - \gamma \delta P_i P \beta = v_\delta(i) - \gamma P_i v_\delta, \quad \forall i.$$

If $\gamma < 1$, then these can be solved by repeated substitution to give a unique positive solution, which implies $\alpha = a = A$ and $\beta = b = B$. If $\gamma = \delta = 1$ and the Markov chain Π is transient, then repeated substitution can also be used because $\Pi^{2n} \rightarrow 0$ as $n \rightarrow \infty$

(Kemeny et al (1976) p.107), the sum increases (as $v_\delta - \delta P_i v_\delta \geq 0$) and is bounded above by v_δ , so (by monotone convergence) the repeated substitution converges.

$$\begin{aligned}\alpha &= \sum_{n=0}^{\infty} \gamma^n \delta^n P^{2n} (v_\delta - \delta P_i v_\delta) \\ \beta &= \sum_{n=0}^{\infty} \gamma^n \delta^n P^{2n} (v_\delta - \gamma P_i v_\delta).\end{aligned}$$

Thus, there is a unique solution to the above equations ($a(i) = A(i)$ and $b(i) = B(i)$), and the game has unique subgame perfect equilibrium payoffs. This equilibrium is also Markov perfect.

Agreement in the bargaining is immediate only if no player benefits from delay, that is if $\alpha(i) \geq \gamma^2 P_i P \alpha$ and $\beta(i) \geq \delta^2 P_i P \beta$.

$$\begin{aligned}\alpha(i) - \gamma^2 P_i P \alpha &\geq \alpha(i) - \gamma \delta P_i P \alpha = v_\delta(i) - \delta P_i v_\delta \geq 0, \\ \beta(i) - \delta^2 P_i P \beta &= (v_\delta(i) - \gamma P_i \alpha) - \delta^2 P_i P (v_\delta - \gamma P \alpha) \\ &= (v_\delta(i) - \gamma P_i \alpha) - \delta^2 P_i P v_\delta + \gamma \delta^2 P_i P^2 \alpha \\ &= (v_\delta(i) - \gamma P_i \alpha) - \delta^2 P_i P v_\delta + \delta P_i [\alpha - v_\delta + \delta P v_\delta], \\ &= (v_\delta(i) - \delta P_i v_\delta) + (\delta - \gamma) P_i \alpha \geq 0.\end{aligned}$$

Q.E.D.

PROOF OF COROLLARY 1

The equilibrium payoffs in this case satisfy: (1) $\alpha(i) + \delta P_i \beta = v_\delta(i) \geq v_\gamma(i)$, (2) $\beta(i) + \gamma P_i \alpha = v_\delta(i) \geq v_\gamma(i)$. This says that the sum of the players' payoffs in state i attains the maximum payoff that either of them could achieve in the one-person stopping problem. So, it is impossible to choose a stopping time and an allocation that makes both players better off. *Q.E.D.*

PROOF OF PROPOSITION 4

Without loss of generality, suppose that the game is in state i of an equilibrium and that the seller is proposing an agreement. If the buyer rejects her proposal then the game moves to state j . Let $y(j)$ be his expected payoff in state j if the bargain in state i is

rejected. Also, let $z(j)$ be the seller's expected payoff if she rejects the buyer's offer in state j tomorrow. Then, in equilibrium $y(j) + z(j) \geq v_\delta(j)$ for all j , because otherwise the buyer could propose a lower price in state j . Thus, if the seller's offer in state i today is rejected, her expected payoff is at least $\gamma P_i(v_\delta - y) \leq \gamma P_i z$. The buyer's expected payoff, if he rejects the offer in state i , is $\delta P_i y$. If the players agree in state i , then the sum of their discounted expected future payoffs is no greater than their current payoff. Rearranging this we get

$$v_\delta(i) \geq \gamma P_i(v_\delta - y) + \delta P_i y = \delta P_i v_\delta + (\gamma - \delta) P_i(v_\delta - y) \geq \delta P_i v_\delta.$$

There are two possibilities, if players agree in state i . Either, $v_\delta(i) > \delta P_i v_\delta$ in state i , which implies $v_\delta(i) = f(i)$ by (1), so the buyer consumes the cake immediately. Or, $v_\delta(i) = \delta P_i v_\delta$, from above, this implies $P_i(v_\delta - y) = 0$. That is, the seller offers the buyer a payoff equal to what he gets in the one person stopping problem commencing in state i and the seller receives nothing. If the seller receives nothing, this equilibrium is equivalent (in payoff terms) to one where the buyer rejects her offer in state i and waits until $f(i_t) = v_\delta(i_t)$, then he receives all of the cake and consumes it immediately. Thus, there is a payoff-equivalent equilibrium where players agree only when $f(i_t) = v_\delta(i_t)$. *Q.E.D.*