Outside Options and Optimal Bargaining Dynamics

Andrew McClellan*

February 12, 2021

Abstract

We study how to design optimal bargaining strategies in a bargaining model with two players, P and A, when A's outside option changes over time. We solve for P's optimal strategy and find a new, but intuitive, set of bargaining dynamics. When A's outside option increases, A is tempted to cease bargaining, leading P to increase A's continuation value by gradually promising A a larger share of the surplus (decreasing demands) and giving A more time to explore his outside option before being forced to make a decision (decreasing pressure). We explore comparative statics and show that although P's value of bargaining is decreasing in A's outside option, it increases when the expected value of A's outside option tomorrow rises. We show P's optimal strategy can be implemented without commitment.

1 Introduction

Outside options are an important determinant of bargaining outcomes. When a firm and worker negotiate over wages, the worker's outside option limits the set

*University of Chicago Booth School of Business Email: andrew.mcclellan@chicagobooth.edu. We would like to thank Dhruva Bhaskar, Simon Board, Eric Budish, Isa Chaves, Jack Fanning, Emir Kamenica, Rohit Lamba, Elliot Lipnowski, Pietro Ortoleva, Doron Ravid and Jakub Steiner for helpful comments. This work is supported by the Liew Family Junior Faculty Fellowship at the University of Chicago Booth School of Business.

of wage offers he will find acceptable. As negotiations go on, the worker's outside option may change. For example, while negotiating with Firm 1, the worker's outside option may go up if he receives a competing offer from competing firms or go down if no competing offer arrives and he becomes pessimistic about the likelihood he will receive an offer from competing firms in the future. If the worker rejects Firm 1's offer today and his outside option becomes better, previous wages offered by Firm 1 may no longer be acceptable to the worker. How should the offers Firm 1 makes depend on changes in the worker's outside option? Should Firm 1 make a take-it-or-leave-it (TIOLI) offer to the worker or give him time to explore his outside options? These questions point to two fundamental decisions while bargaining: how much to *demand* (i.e., how high or low to make the wage) and how much *pressure* to apply (i.e., how long to let the worker consider the offer).

We bring a dynamic contracting approach to a classic "split-the-pie" bargaining problem between two players, P and A, to which we add a dynamic outside option for player A. We solve for P's optimal bargaining strategy and find that delay and bargaining breakdowns are both prevalent *and* efficient. The bargaining process features a history dependence that resembles haggling: when A's outside option is high, A threatens to walk away and take his outside option, leading P to gradually and permanently lower his demands. The pressure exerted by P decreases as well, giving A more time to explore his outside options before P makes a TIOLI offer to A. Our results show a *complementarity* between the choice of how much to demand and how much pressure to apply.

Unsurprisingly, we find that an increase in A's outside option lowers P's expected value from bargaining as P is forced to make lower demands to prevent A from taking his outside option. However, P's expected value of bargaining is *increases* when the expectation of A's outside option tomorrow increases. This may be surprising at first glance: A's outside option is more likely to increase tomorrow and such an increase lowers P's utility. However, an increase in the expected change of A's outside option increases the value for A of exploring his outside options, which allows P to increase his demands while still incentivizing A to continue bargaining. This result rationalizes why we see firms take actions during negotiations that may increase a worker's future outside options (e.g., a firm negotiating with a current employee may write a positive recommendation letter for the worker or give the worker time to interview with competing firms).

Unlike much of the bargaining literature, our results generate efficient delay and gradual concessions in bargaining demands without asymmetric information or behavioral types. The intuition for why delay is efficient can be seen by viewing delay in reaching an agreement as "experimentation" for A. Consider a worker deciding whether to take a firm's wage offer today. If he expects to receive a reasonable offer from the firm tomorrow, he may prefer to delay, knowing he can take the offer tomorrow if his outside option goes down. In this way he can enjoy the benefits when his outside option increases, but still be protected against the risk that it decreases. The firm may also benefit from allowing the worker to explore his outside option if it can appropriate a larger part of the surplus by decreasing its wage offer when the worker's outside option is low.

The dynamics of P's bargaining strategy are driven by the fact that A has the option of walking away and taking his outside option at any time. P's mechanism must therefore ensure A's continuation value from bargaining is sufficiently high. Treating A's choice to take his outside option early as a deviation, we must consider deviations in an infinite-dimensional space, making analysis of the problem difficult. Nevertheless, we identify a binding class of constraints on deviations for A and a tractable relaxed problem incorporating only these constraints that yields a solution to our full problem.

The optimal offer process, although it features non-stationary dynamics, is still simple and intuitive. It can be characterized by three objects: a demand function, a split threshold and a breakdown threshold. A split is made when A's outside option goes below the split threshold. The placement of the split threshold tells us how much pressure is placed on A: the higher the threshold, the less time A has to explore his outside options before being given a TIOLI offer. We find both P's demand and the location of the split threshold change over the course of the game, decreasing in the maximum of A's past outside options. When A's outside option reaches a new high, P gradually lowers his demand and the pressure on A, keeping it fixed until A's outside option again reaches a new high or an agreement is reached. The bargaining process does not always end with players reaching an agreement: if A's outside option goes above the breakdown threshold, A walks away and takes his outside option. Our model shows that delay and breakdowns, natural outcomes in many real-world bargaining settings, can arise in a complete information environment and are in fact necessary for an efficient outcome.

We also study how the structure of changes in A's outside option affects P's value from bargaining. Would P rather A's outside option change frequently in small amounts or infrequently in large jumps? Imposing a constraint on the expected speed at which A's outside option may change, we show that whenever A's outside option is a martingale and changes frequently, we can construct another martingale outside option with only infrequent large jumps jumps whose optimal mechanism yields strictly higher utility for P. This result highlights how different types of outside option processes add distortions to the optimal mechanism through A's individual rationality constraint and shows that jump processes introduce less distortions than continuously changing processes.

Although the commitment assumption is reasonable in many of our examples, it is natural to ask how a large a role this assumption plays. We show that our results are robust to relaxing the assumption that P can commit to his bargaining strategy and looking at a classic discrete time alternating offers version of our model in which we construct an equilibrium that converges to our optimal mechanism in the frequent offer limit. This result shows that, if we allow P to select his preferred equilibrium, the loss from dropping commitment is negligible.

Related Literature

Our paper brings a dynamic contracting approach to a bargaining environment and so lies in the intersection of these two literatures. The two important features of our environment, the outside option and the stochasticity of the bargaining environment, have both been the focus of attention in the bargaining literature. The importance of the outside option in bargaining is well known and has been studied in axiomatic bargaining (Nash (1950)), strategic bargaining (Binmore et al. (1989)), in conjuction with reputation (Compte and Jehiel (2002), Lee and Liu (2013)) and in relation to the Coase conjecture (Board and Pycia (2014)). These papers assume players' outside options stay fixed throughout the game. The literature on changing bargaining envrioments has recieved growing attention in recent years and has looked the impact of newly arriving players (Fuchs and Skrzypacz (2010), Chaves (2019)), the impact of transparency of outside options (Hwang and Li (2017)), the arrival of information about a seller's type (Daley and Green (2018)) and changing costs of supplying a good (Ortner (2017)). These papers have typically focused on studying stationary equilibria where players' strategies are stationary in their beliefs about their opponent's type.

Rubinstein (1982) established the uniqueness of equilibrium outcomes in an infinite-horizon alternating-offers bargaining model and found that an agreement is reached immediately. The finding of no delay in reaching an agreement is at odds with some real-world phenomena (e.g., haggling, labor strikes, etc.), which spurred the search for bargaining models that generate delay. The majority of the literature has looked to incomplete information to generate delay. Papers in the Coasian bargaining literaure generate a cream-skimming style of delay, finding equilibrium with a gradual, but deterministic, downward movement in demands. Papers in the reputational bargaining literature, such as Abreu and Gul (2000), generate a war-of-attrition style of delay, finding equilibrium in which any concession in bargaining demands leads to immediate agreement. Our paper by contrast generates delay in a complete information environment and features a new set of dynamics to players' demands with sporadic and gradual concessions, along with periods of intransigence where players hold firm to their demands, and breakdowns in bargaining after significant delay. Backus et al. (2020) empirically document delay in bargaining and observe that there is often delayed disagreement and gradually changing offers, features that are outside standard bargaining models but are found in our model.

Efficient delay arises in Merlo and Wilson (1995) and Cripps (1998), who study models where the size of the surplus to be split is stochastic. In these models, players may benefit from delay only if the expected discounted total surplus tomorrow is greater than the surplus today. The efficiency of delay in our model is driven instead by the option value of choosing between the outside option and a split. An important difference is that the outside option is changing rather than the size of the pie changing introduces very different strategic forces. The dynamic outside option leads to a rich set incentive constraints that P must satisfy to ensure A continues bargaining; its these incentive constraints, which do not arise when only the size of the pie changes, that lead to the history dependence in P's bargaining strategies.

The dynamics in the efficient offer processes feature a backloading of incentives as in Ray (2002) and a downward rigidity to P's demands. Similar types of rigidity are also found in Harris and Holmstrom (1982), who find such rigidity in wages that arise from a competitive market for workers, Thomas and Worrall (1988), who study the design of self-enforcing contracts, and in McClellan (2020), who studies the design of approval rules to incentivize experimentation. Our paper differs from these both in the set of tools available to P (he decides how *and* when to split the surplus, rather than choose transfers in each period as in Thomas and Worrall (1988)) and that our solution can be implemented without commitment (this feature was found also in McClellan (2020)).

2 Model

Two players, P and A, bargain over how to split a surplus of size one. Each player $i \in \{P, A\}$ has a utility function $u_i : [0, 1] \to \mathbb{R}$ over the share of the surplus they receive if an agreement is reached. Time runs continuously between 0 and ∞ and both players discount time at a rate of $r > 0$.

The game ends when either an agreement to split the surplus is reached or one player takes their outside option. Each player can take their outside option at any time, in which case both players will receive their outside options. P has a constant outside option with value $\nu \geq 0$. A's outside option is given by a stochastic process Y with a value at time t of $Y_t \in [\underline{Y}, \overline{Y}]$, where $\underline{Y} \geq 0$. At each time t, the entire history $\{Y_s : \ 0 \leq s \leq t\}$ is publicly observable. Let $\mathcal{F} = \{\mathcal{F}_t\}_{t\in[0,\infty)}$ be the filtration generated by Y. We consider two types of stochastic processes:[1]

Diffusion Process: The evolution of Y_t is given by

$$dY_t = \mu(Y_t)dt + \sigma(Y_t)dB_t,$$

where B_t is a standard Brownian motion on the canonical probability space subject to standard conditions. Both $\mu(\cdot)$ and $\sigma(\cdot)$ are Lipschitz continuous on $(\underline{Y}, \overline{Y})$ with $\sigma(y) > 0 \ \forall y \in (\underline{Y}, \overline{Y})$. The boundaries $\underline{Y}, \overline{Y}$ may be absorbing or reflecting.

Search Process: A receives iid draws $Z_t \sim F$ with finite support at a Poisson rate $\zeta(Y_t) > 0$ which is Lipschitz continuous and strictly increasing in Y_t. When a new Z_t arrives, A's outside option is then the max of Z_t and his outside option right before t: $Y_t = max\{Z_t, Y_{t-}\}$. There is a reflecting barrier at $\widetilde{Y}$, so that if $Y_t > \widetilde{Y}$, it jumps down to $\widetilde{Y}$ in the next instant: $Y_{t+} = \widetilde{Y}$. Otherwise, Y_t evolves according to $dY_t = \eta(Y_t)dt$ for some Lipschitz continuous $\eta(\cdot) < 0$ on $(\underline{Y}, \overline{Y})$. We assume $\underline{Y}$

[1] We discuss in the Online Appendix the properties of these two stochastic processes that are used in our proof. The intuition for our results does not rely Y_t taking one of these two forms, but these two stochastic processes are useful for simplifying the analysis while still capturing a wide range of settings.

cannot be reached in finite time from Y_0.[2] We provide an explicit construction of Y in the Online Appendix.

To simplify the exposition, we take $\widetilde{Y} = \overline{Y}$ when discussing a diffusion process.

Our main result focuses on the case in which P is allowed to commit to how he makes demands, which we call a mechanism. We can heuristically think of P as making a demand at each moment in time that A accepts or rejects. We define a mechanism by the outcome it induces.

Definition 1. *A* ***mechanism*** *consists of $\mathcal{F}$-measurable functions $(\tau, d_\tau, \alpha_\tau)$ where*

i. τ is a stopping time that gives the time when the game ends; that is, a split is made or one player takes his outside option.

ii. $d_\tau \in \{0, 1\}$ is a decision rule that equals 1 if and only if a split is made at time τ.

iii. $\alpha_\tau \in [0, 1]$ gives P's share of the surplus if a split is made at time τ.

P's expected payoff from a mechanism $(\tau, \alpha_\tau, d_\tau)$ is

$$J(\tau, d_\tau, \alpha_\tau) = \mathbb{E}_{Y_0}[e^{-r\tau}\big(d_\tau(u_P(\alpha_\tau) - \nu) + \nu\big)],$$

and A's expected payoff is

$$V(\tau, d_\tau, \alpha_\tau) = \mathbb{E}_{Y_0}[e^{-r\tau}\big(d_\tau(u_A(1 - \alpha_\tau) - Y_\tau) + Y_\tau\big)].$$

For notational convenience, we will drop dependence on Y_0 in $\mathbb{E}_{Y_0}$ where it causes no confusion. We can easily incorporate flow costs c_i from bargaining for player i. Because the expected flow costs to i from a mechanism using τ are $\mathbb{E}[\int_0^\tau e^{-rt} c_i dt] = \mathbb{E}[\frac{1-e^{-r\tau}}{r} c_i]$, adding flow costs to the model is equivalent to subtracting $\frac{c_i}{r}$ from i's utility adding $\frac{c_i}{r}$ to both u_i and i's outside option.

Without loss, we focus on mechanisms in which A never takes his outside option.[3] To ensure that A does not take his outside option early, we impose a dynamic individual rationality constraint on the set of mechanisms P can use.

[2] This assumption is purely made for simplifying the statement of our results; the structure of the optimal mechanism when relaxing this is discussed in the Appendix.

[3] Replacing any instance of A taking his outside option with P doing so does not change players' payoffs.

Definition 2. $(\tau, d_\tau, \alpha_\tau)$ *is* ***dynamically individually rational*** *if for every* $t \geq 0$ *and history up to* t*,* A*'s continuation value is weakly greater than* Y_t*.*

We place two relatively weak assumptions on the primitives of the model. Our first assumption imposes that players' utilities are concave, which ensures players cannot benefit from randomization over split amounts.

Assumption 1. *The utility functions satisfy* $u_i''(\cdot) \leq 0 < u_i'(\cdot)$ *on* $[0, 1]$, $i \in \{P, A\}$ *with strict concavity for some* i *and bounded derivatives.*

Our next assumption ensures that the expected future discounted value of A's outside option is lower than his current outside option. This assumption will be used to show that if A knows there is no possibility of reaching an agreement in the future, then A's best option will be to take his outside option immediately.

Assumption 2. $e^{-rt}Y_t$ *is a supermartingale.*

We note that Assumption 2 holds when Y_t is a supermartingale. Assumption 2 is a natural property to impose on A's outside option.[4] Consider a firm-worker wage negotiation where Y_t represents the value of searching for new job offers. Because A always has the option to ignore incoming job offers, reentering the search market immediately cannot be worse for A than than rejecting all P's offers for some length of time before reentering the search market.

Discussion

Outside Option: While a firm-worker negotiation is our main example, our model fits many other settings where outside options may change. For example:

i. Buyer-seller negotiations where the evolution of buyer's outside option represents the entry and exit of competing sellers.

ii. Debt negotiations between a bond holder and a politician considering defaulting on sovereign debt where the evolution of the politician's outside option

[4] Let $\bar{Y}_t := \sup_\tau \mathbb{E}_{Y_t}[e^{-r\tau}Y_\tau]$ be A's optimized value choosing when to take his outside option Y_t. After walking away from bargaining, A should still be able to explore his outside option and decide when to take it. Thus, A's continuation upon walking way from bargaining will be $\bar{Y}_t$. We could rewrite the model, only now replacing Y_t with $\bar{Y}_t$. Standard optimal stopping results tell us that $\bar{Y}_t$ is a supermartingale.

represents a change in the public approval for the politician if she exits negotiations and defaults.

iii. Peace negotiations where the evolution of the outside option represents a change in the costs of restarting the conflict or likelihood of winning the subsequent conflict.

The two stochastic processes we consider give us flexibility in how A's outside option changes over time and show that our results hold under a wide range of stochastic processes for Y. A diffusion process fits situations in which A's outside option changes continuously over time, such as a housing market in which the value of searching for a new house changes in small amounts each day. The diffusion model also fits learning models in which Y is A's expected utility of a fixed, but unknown, outside option ψ and A learns about the value of ψ over time via a Brownian news process.

Our search process is a generalization of standard search processes and allows for the value of search to change over time. It has a natural interpretation in the job search example. A searches for new outside offers that have value Z_t when they arrive, which he can then decide to accept or reject.[5] In the absence of the arrival of a new offer, A becomes pessimistic about the likelihood additional offers will arrive. Letting Y_t be the optimized value of searching for new offers, this pessimism corresponds to a decrease in Y_t and ζ. The reflecting barrier $\widetilde{Y}$ accounts for the fact that upon rejecting a Z_t offer, A's outside option may fall; however, it need not fall back to its level prior to the arrival of Z_t if its arrival makes A more optimistic about receiving future offers.[6] The class of search processes includes as limiting cases standard search models such as stationary search with recall (take $\zeta(\cdot)$ constant, $\eta(\cdot) = 0$ and $\widetilde{Y} = \overline{Y}$) and stationary search without recall (take $\widetilde{Y} = \underline{Y}$ and $\zeta(\underline{Y}) > 0$). We discuss the optimal mechanism for these limiting cases,

[5] Taking Z_t can be interpreted as the value of accepting an outside offer immediately or the value of leaving the negotiations with P to bargain with a new firm.

[6] We can microfound this more formally by supposing that there are K possible offers for A and the next offer arrives at a constant Poisson rate. K is unknown to both players and has a geometric distribution. Every time A receives an offer, his belief that there are additional offers to be found jumps up to some $\bar{p}$, after which it drifts down over time as long as a new offer is not received. Let $s(t)$ be the length of time at t since the arrival of a new offer. Then $Y_t = \sup_{\tau} \mathbb{E}[e^{-r\tau} Z_\tau | s(t)]$ and $\widetilde{Y} = \sup_{\tau} \mathbb{E}[e^{-r\tau} Z_\tau | s(t) = 0]$. By standard optimal stopping arguments (e.g., Peskir and Shiryaev (2006)), Y_t is a super-martingale.

which will be similar to the optimal mechanism in our main results, in the Online Appendix.

Observable Y_t: The assumption of common knowledge of Y_t is similar to other papers in the literature on changing bargaining environments and is economically reasonable in many situations. In the firm-worker example, the offers that a worker has from other firms often can be verifiably disclosed by the worker. In the buyer-seller example, the presence of competing sellers is likely common-knowledge.

Commitment: In contrast to much of the bargaining literature, we allow P to commit to his demands. In many situations, such an assumption is reasonable: in firm/worker or seller/buyer negotiations, P changing his bargaining demands today may affect bargaining outcomes with the future workers or buyers. If a firm reneges on a offer, it will lose credibility in future negotiations. When P is a long-run player, such repeated game punishments will enforce the commitment solution. Solving the problem with commitment will prove useful when we study equilibria in a discrete-time version of our model. The commitment solution gives an upper bound on P's payoffs in any equilibrium, which would otherwise be difficult to solve for. Having identified the upper bound, it will be much easier to construct an equilibrium that achieves this upper bound.

3 Delay

Waiting to agree to a split is inefficient: a split that is enacted in the future would be better for both players if it were enacted immediately. By Assumption 2, we know that delay in taking A's outside option is also inefficient. It seems natural to conjecture that an efficient outcome features no delay. The economic intuition for why this conjecture is wrong can be seen by viewing delay as A experimenting with his outside option and the option to accept a split of the surplus as insurance against a decrease in his outside option. Bargaining creates option value for A.

For an illustrative example, suppose $dY_t = dB_t$ in $(\underline{Y}, \overline{Y})$ and $\nu = u_P(0) = 0$. If $Y_0 = u_A(1)$, the only possible bargaining split that achieves no delay and A would accept is to give the entire surplus to A. Consider an alternative offer by P in which he asks A to wait for Δ length of time and commits to give $1 - \Delta^2$ to A. If A waits and his outside option goes up, he can take his new higher outside option, but if his outside option goes down, he can take the split if $Y_\Delta < u_A(1 - \Delta^2)$. This option value protects him against a decrease in his outside option. A's expected

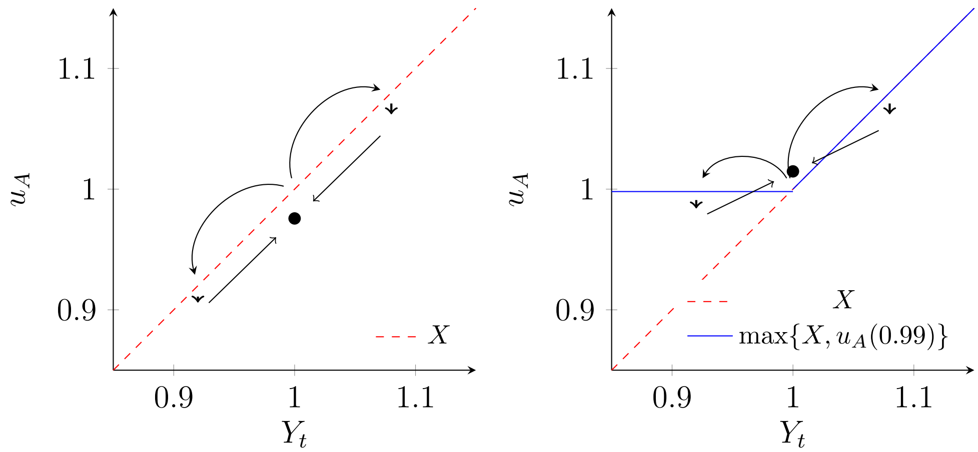

Figure 1: For illustrative purposes, we treat the movement of Y_t as a random walk. The upper curved arrows indicate the movement of Y_0 to Y_Δ, and the downward arrows indicate discounting costs to $e^{-r\Delta}Y_\Delta$. The black dot indicates the expected value of waiting until $t = \Delta$ to make a decision.

utility of waiting is equal to $u_A(1) + \frac{\sqrt{\Delta}}{\sqrt{2\pi}} + O(\Delta)$. For small Δ, this policy yields a higher value than $u_A(1)$. The driving force for this result is the fact that allowing A to choose the max of $u_A(1-\Delta^2), Y_\Delta$ creates a kink in the underlying payoff for A. The convexity this kink creates is enough to make it beneficial for A to delay and take a lottery over payoffs tomorrow. P is also better off because with positive probability he receives Δ^2 share of the surplus compared to 0 before.

This intuition is economically relevant in many bargaining situations. A firm might be able to make a TIOLI offer that a worker would choose to accept. However, in order to get the worker to forego his outside option this offer may require such a high wage that the firm may prefer to give the worker a lower offer but grant the worker time to explore his other options before deciding whether to accept the firm's offer. Such non-TIOLI offers are often used by firms.

A natural benchmark we might consider is that of a social planner who, for some $\rho \in [0,1]$, places ρ and $1-\rho$ weight on P's and A's utility, respectively. Ignoring the dynamic individual rationality requirement, the social planner will then choose a mechanism that solves

$$\sup_{(\tau,\alpha_\tau,d_\tau)} \mathbb{E}[e^{-r\tau}\big(d_\tau\big(\rho[u_P(\alpha_\tau)-\nu)+(1-\rho)[u_A(1-\alpha_\tau)-Y_\tau]\big)+\rho\nu+(1-\rho)Y_\tau\big)]$$

The solution takes the familiar class of *stationary policies*: namely, we stop whenever Y_t first crosses one of two stationary thresholds and the split amount when an agreement is reached is constant.

Proposition 1. *There are* $(b, B, \alpha^b) \in \mathbb{R}^3_+$ *with* $b < B$ *such that the social planner's mechanism is* $\tau = inf\{t : \; Y_t \notin (b, B)\}, \; d_\tau = \mathbb{1}(Y_\tau \leq b), \; \alpha_\tau = \alpha^b$.

This stationary structure is familiar from standard optimal stopping problems.[7] However, when $\rho > 0$, the social planner's problem does not take into account the incentive constraint that A must find it optimal to delay taking his outside option until the prescribed time. In general, a stationary mechanism which respects A's incentive constraints *will not* be efficient.

This is easiest to see when $\nu = u_P(0) = 0$. Consider a stationary mechanism that respects A's incentive constraints and calls for P to demand $\alpha^b > 0$. Let's go to the moment when Y_t has reached B and P is about to take his outside option. P clearly has no incentive to take his outside option early, so A must be indifferent between continuing and walking away at B. Suppose P were to come to A and promise to always demands $\frac{\alpha^b}{2}$ and let A choose when to accept this demand or take his outside option. This continuation strategy increases the value of bargaining for A, and so, if under the stationary mechanism A was indifferent between continuing and taking his outside option at B, he will now strictly prefer to continue bargaining. Moreover, this would also increases P's utility because a split will now be reached with positive probability. This argument implies that any stationary policy with $\alpha^b > 0$ can be improved upon. We will need to look at a larger class of mechanisms to find the P-optimal mechanism.

4 Mechanism Design Problem

We now turn to the design of the optimal mechanism. Even though we allow for arbitrarily complex mechanisms, the solution is quite simple and intuitive. The optimal mechanism is measurable with respect to only two state variables, Y_t and the running maximum $M_t = \max_{s \in [0,t]} Y_s$, and can be described by three objects: a demand function $\alpha^*(M_t)$, a split threshold $S^*(M_t)$, and a breakdown threshold $\overline{R}^*$. An agreement is reached whenever $Y_t \leq S^*(M_t)$, with P demanding $\alpha^*(M_t)$. P takes his outside option if and only if $\overline{R}^*$ is crossed before $S^*(M_t)$. The location of

[7]The proof follows almost immediately from Lemma A.1 in the Online Appendix and is hence omitted.

the split threshold $S^*(M_t)$ corresponds to the amount of *pressure* being placed on A: the higher S^* is, the less time A has to explore his outside option before being forced to make a decision. Theorem 1 shows that α^* and S^* are both decreasing over the course of the game.

Theorem 1. *There is an optimal mechanism $(\tau^*, d^*_\tau, \alpha^*_\tau)$ given by, for some decreasing continuous functions $S^*(\cdot), \alpha^*(\cdot)$ and threshold $\overline{R}^*$,*

$$\tau^* = inf\{t:\ Y_t \notin (S^*(M_t), \overline{R}^*)\},\ d^*_\tau = \mathbb{1}(Y_\tau \leq S^*(M_\tau)),\ \alpha^*_\tau = \alpha^*(M_\tau).$$

All proofs are relegated to the Appendix or Online Appendix.

A can not recall past outside options, so why should M_t play a role in the optimal mechanism? Although M_t is payoff irrelevant, it captures the additional continuation value P promises A to ensure A continues bargaining. A is most tempted to take his outside option whenever Y_t is highest, namely when $Y_t = M_t$. P must then increase A's continuation value to prevent A from walking away, which is done by decreasing both α^* and S^*. Both of these changes are rigid: α^* and S^* never rise after being lowered, implying a persistent effect from A having a higher past outside option. Our results show that P provides incentives is a "smooth" way by keeping both α^* and S^* constant until M_t increases. The stochasticity of increases in M_t generates demands by P that gradually decrease, but not without long periods of P holding firm to his demands.

P takes his outside option with positive probability when there is delay in the optimal mechanism, an outcome we call a *bargaining breakdown*. Unlike much of the bargaining literature, a breakdown in our model may happen after significant delay. Intuitively, delay creates benefits through experimentation, which is only useful if the result of the experimentation, a higher Y_t, is sometimes taken. Otherwise, delay would be inefficient and P would be better off making an immediate TIOLI offer.

Gradual concession and significant delay in reaching a split or breakdown are often observed in real life bargaining. Using data on negotiations taking place on eBay, Backus et al. (2020) find frequent gradual concessions in demands and delayed bargaining breakdowns, which they note cannot be explained by most bargaining models. Our results both generate and show the efficiency of such dynamics.

Fixing Y_t, the fact that α^* and S^* are decreasing means that the higher A's outside options have been in the past, the longer it will be until an agreement is

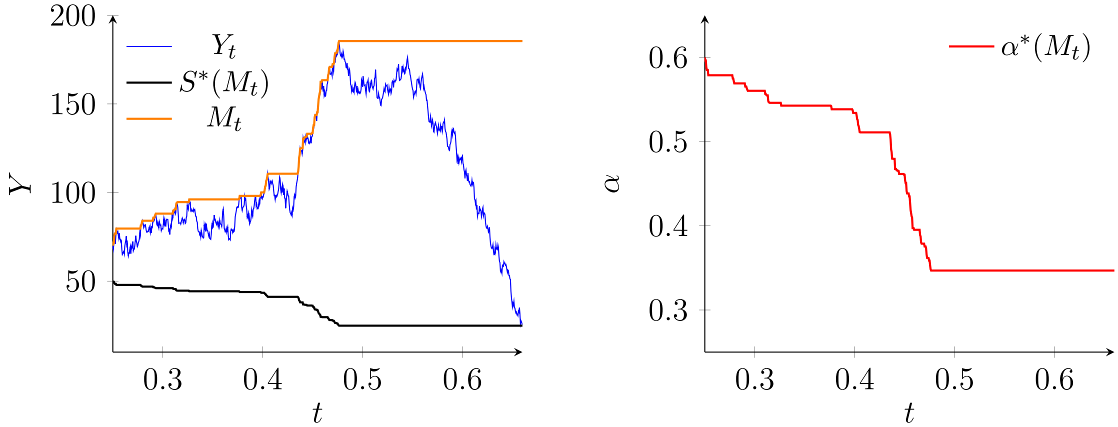

Figure 2: The demand and split threshold are monotonic, decreasing rapidly in spurts and then remaining fixed as Y_t goes down.

reached. For an outside observer, this might appear to be something like anchoring effects or loss aversion. Our results show how such dynamics arise with standard preferences.

Proof Sketch

It is expositionally useful to begin by slightly weakening the requirement that the mechanism be dynamically individually rational. Suppose A were to deviate from P's mechanism by taking his outside option early at some stopping time τ'. A's expected utility would then be $V(\tau \wedge \tau', d_\tau \mathbb{1}(\tau < \tau'), \alpha_\tau)$. P will need to ensure that A prefers not to quit at τ' for every τ'. We call this a DIR constraint:[8]

$$DIR : \sup_{\tau'} \; V(\tau \wedge \tau', d_\tau \mathbb{1}(\tau < \tau'), \alpha_\tau) \leq V(\tau, d_\tau, \alpha_\tau),$$

Quitting early is the only deviation by A we need to consider: by committing to take his outside option upon A rejecting P's offer at τ, P can ensure A does not

[8] DIR is slightly weaker than the dynamic individual rationality constraint in that it can violate dynamic individual rationality on a probability zero event. Lemma O.A.1 in the Online Appendix shows that any dynamically individually rational mechanism satisfies DIR. After solving the problem with only DIR, we will verify that the solution is dynamically individually rationality.

delay in accepting an individually rational offer. P's problem is then

$$J^* = \sup_{(\tau, d_\tau, \alpha_\tau)} J(\tau, d_\tau, \alpha_\tau) \tag{1}$$
$$\text{subject to } DIR.$$

For those familiar with the dynamic contracting literature, the most natural approach to this problem would be to treat the A's continuation value as a state variable and use a dynamic programming approach to solve for P's optimal strategy as in Sannikov (2008). However, because we also need to keep track of Y_t as a state variable, using this approach in our model would involve solving a PDE, which is not feasible. Looking at 1, the main difficulty we face is that DIR involves finding A's best response for an arbitrary $(\tau, d_\tau, \alpha_\tau)$ among an infinite dimensional set of τ', which is intractable. We will need to find a way to relax the DIR constraint in order to make the problem more tractable.

Before introducing the relaxation, we discretize the model, using a discrete-time framework with period length Δ. We replace Y_t with a discretized outside option X_t which moves along a countable grid of points $\mathcal{G}^\Delta$ and define $M_t^X = \max_{s \in \{0, \Delta, ..t\}} X_s$. We will drop the X from M_t^X where it causes no confusion. If Y is a diffusion process, then we take X to be a generalized random-walk, with $X_t = X_{t-\Delta} + w_t$ where $w_t \in \{-\epsilon, 0, \epsilon\}$ for some $\epsilon > 0$. If Y is a search process, then we construct X in the following way. In each period, $Z_t \sim F^\Delta$ arrives with probability $\zeta^\Delta(X_{t-\Delta})$ and $supp(F^\Delta) \subseteq \mathcal{G}^\Delta$. If a Z arrives, then $X_t = max\{Z_t, X_{t-\Delta}\}$. If no Z arrives, X moves down to the next highest grid point: $X_t = max\{x \in \mathcal{G}^\Delta : \ x < X_{t-\Delta}\}$.[9] As we take $\Delta \to 0$, we will assume that the discrete time process X converges to its continuous time counterpart Y.[10] We now define J and V as before but with expectation over X rather than Y.

We now relax DIR by limiting the set of deviations A may take. Suppose A takes his outside option whenever X goes above some threshold B. A's utility from such a strategy is $V(\tau \wedge \tau_+(B), d_\tau(B), \alpha_\tau)$ where

$$\tau_+(B) := min\{t : \ X_t \geq B\}, \ d_\tau(B) := d_\tau \mathbb{1}(\tau < \tau_+(B)).$$

[9] We are dropping the reflecting barrier $\widetilde{Y}$ here. This is without loss in the limit as $\Delta \to 0$, where our X_t process, when above $\widetilde{Y}$, will move below $\widetilde{Y}$ instantaneously in the limit.

[10] Details on how to construct $\mathcal{G}^\Delta$ and law of motion for X to ensure convergence to Y are provided in the Appendix.

Our relaxed problem will limit A to choose deviations of this form. For a threshold B, we define the constraint $RDIR(B)$ to be

$$RDIR(B) : V(\tau \wedge \tau_+(B), d_\tau(B), \alpha_\tau) \leq V(\tau, d_\tau, \alpha_\tau).$$

Consider A's first-best mechanism (i.e., letting $\rho = 0$ in the social planner's problem) and let R^A be the threshold at which A takes his outside option. It is easy to see that P must take his outside option immediately at $X_t \geq R^A$ in any dynamically individually rational mechanism; otherwise, A can get his first-best payoff by walking away. We therefore restrict attention to mechanisms with $\tau \leq min\{t : X_t \geq R^A\}$.

Let $\mathcal{X}_N = \{X^0, X^1, ..., X^N\}$ be the points of $\mathcal{G}^\Delta$ in $[X_0, R^A)$ in ascending order. Our relaxed mechanism-design problem is given by

$$\begin{aligned} &\sup_{(\tau, d_\tau, \alpha_\tau)} \; J(\tau, d_\tau, \alpha_\tau) \\ &\text{subject to } RDIR(X^n) \; \forall X^n \in \mathcal{X}_N. \end{aligned} \tag{2}$$

We employ a Lagrangian approach to solve this relaxed problem. There exists Lagrange multipliers $(\lambda(X^0), ..., \lambda(X^N)) \in \mathbb{R}_-^{N+1}$ associated with the $RDIR(X^n)$ constraints such that the solution to 2 solves

$$\begin{aligned} \sup_{(\tau, d_\tau, \alpha_\tau)} \; \mathbb{E}\Bigg[& e^{-r\tau}\big(d_\tau\{u_P(\alpha_\tau) - \nu - \lambda(X^0)(u_A(1-\alpha_\tau) - X_\tau)\} + \nu - \lambda(X^0)X_\tau\big) \\ & + \sum_{n=1}^{N} \lambda(X^n)\big\{e^{-r(\tau\wedge\tau_+(X^n)}\big(d_\tau(X^n)(u_A(1-\alpha_\tau) - X_{\tau\wedge\tau_+(X^n)}) + X_{\tau\wedge\tau_+(X^n)}\big) \\ & - e^{-r\tau}\big(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau\big)\big\}\Bigg] + \lambda(X^0)X_0. \end{aligned} \tag{3}$$

Although the Lagrangian in 3 may appear complicated, we can use optimal stopping arguments to pin down the structure of the solution.

We show the optimal $(\tau, d_\tau, \alpha_\tau)$ in 3 possesses a kind of "local stationarity." Let us focus on the optimal rule before $\tau_+(X^1)$ (that is, before X_t goes above X^1). Then α_τ is equal to

$$\alpha^0 := argmax_{\alpha\in[0,1]} \; u_P(\alpha) - \lambda(X^0)u_A(1-\alpha),$$

which does not depend on X_τ. We also show there exists a split threshold S^0 such that P stops and implements a split whenever $X_t \leq S^0$. Because of the discretization of X, we may require P to randomly stop with probability $\gamma^0 \in [0,1)$ whenever X_t is at the grid point right above S^0 (this randomization plays a negligible role for small Δ). The split threshold, randomization and demand are constant as long as $t < \tau_+(X^1)$.

At $\tau_+(X^1)$, the structure of mechanism will change. We show that the mechanism is again locally stationary until the next X^n above $X_{\tau_+(X^1)}$ is crossed. Repeating these arguments, the optimal mechanism after any history depends only on the current X_t and the set of thresholds in $\mathcal{X}_N$ that have been crossed, for which M_t is a sufficient statistic. The optimal mechanism after any history takes a similar form as before $\tau_+(X^1)$: P's demands are independent of X_τ and given by a function $\alpha^\Delta(M_\tau)$. There exists a threshold $S^\Delta(M_t)$ such that a split is made immediately at any $X_t \leq S^\Delta(M_t)$ and with probability $\gamma^\Delta(M_t) \in [0,1)$ whenever X_t is at the grid point right above $S^\Delta(M_t)$. We show that the outside option is taken whenever $X_t \geq \overline{R}^\Delta$ for some $\overline{R}^\Delta \geq X_0$. We then define $S^*, \alpha^*, \overline{R}^*$ by taking the limit of $S^\Delta, \alpha^\Delta, \overline{R}^\Delta$ as $\Delta \to 0$; we can ignore the role of γ^Δ in the limit. We verify that $(\tau^*, d^*_\tau, \alpha^*_\tau)$ satisfies dynamic individual rationality and is optimal in continuous time.

The main thing left to understand is *how* the mechanism changes with M_t. Whenever M_t increases, P needs to increase A's continuation value. Decreasing α^* provides a clear way of doing so. Decreasing S^* provides another means, as A benefits from the additional time to explore his outside option. A lower S^* is costly for P as it both lengthens the time until a split is reached and increases the probability that M_t increases before an agreement in reached, forcing P to increase A's continuation value. Given these two means of increasing A's continuation value, it is not immediately clear which P will use. P could decrease α^* and increase S^*, decrease S^* and increase α^*, or decrease both. The fact that α^* and S^* are both decreasing in M_t comes from a complementarity between these two aspects of the bargaining strategy.

To build some intuition for this complementarity, consider the choice of α^* and S^* at some (Y_t, M_t). As M_t increases, P must deliver A a larger continuation value at y. For each choice of S^*, an increase in A's continuation value leads to a lower corresponding choice of α^*, which directly reduces P's marginal utility of S^* in two ways. First, discounting is not as costly for P when α^* is lower, thereby reducing the benefit to P of a higher S^*. Second, when P considers an increase in S^*, he would need to compensate A with a further decrease in α^* in order to maintain

A's continuation value. Due to the concavity of u_P and u_A, a decrease in α^* when starting at a lower α^* is, in utility terms, more costly for P and less beneficial to A. Together, the necessary compensating decrease in α^* in response to an increase in S^* is now more costly for P.

A lower α^* also changes the responsiveness of A's continuation value to an increase in S^*. When α^* is lower, the value of stopping and taking a split is higher for A, thereby reducing the decrease in A's utility caused by an increase in S^*. This force reduces the necessary utility compensation to A for an increase in S^*, thereby increasing the marginal utility to P of increasing S^* when α^* is lower. Nevertheless, we find that this effect is always smaller than the previous two effects and P's marginal benefit of increasing S^* is lower when α^* is lower. When P needs to increase A's continuation value, P will find it profitable to use a decrease in S^* as way to mitigate the decrease in α^* that would otherwise be necessary.

Our proof shows several other notable features of the optimal mechanism. When delay is optimal in the continuous time limit and we make Assumption 2 strict,[11] the solution to our relaxed problem satisfies dynamic individual rationality in the discrete-time model for sufficiently small Δ. Thus, our proof gives the discrete-time optimal mechanism and shows that the qualitative features of Theorem 1 are not artifacts of the continuous time structure.

We find that A's continuation value is exactly equal to his outside option when $Y_t = M_t$. This allows us to show that optimal continuation mechanism when $Y_t = M_t$ will be the same as the optimal mechanism from starting at $Y_0 = M_t$. Therefore, the form of the optimal continuation mechanism at any (Y_t, M_t) is *independent* of the starting Y_0. This independence from the starting value Y_0 is a standard feature in Markovian individual decision-maker problems, but does not always arise when we include strategic interactions between players. This independence relies on the flexibility of P's mechanism. For example, if we were to restrict P to only choose among stationary policies (i.e., a constant demand and split threshold), the choice of an optimal policy *would*, in general, depend on Y_0.

Corollary 1. $(\tau^*, d^*_\tau, \alpha^*_\tau)$ *is optimal for every* Y_0.

Given the history dependence of the optimal mechanism, we may be concerned that if continuation play at some history was extremely inefficient, players would

[11] That is, we assume there is an $r' < r$ such that $e^{-r't}X_t$ is a supermartingale for all sufficiently small Δ.

have an incentive to renegotiate the mechanism. Proposition 2 below shows the optimal mechanism is resistant to such concerns. We say a mechanism is *constrained Pareto efficient* if there is no dynamically individually rational mechanism that weakly increases both players' utilities and strictly increases at least one player's utility. Proposition 2 shows we can add off-path continuation mechanisms to $(\tau^*, d^*_\tau, \alpha^*)$ so that it is constrained Pareto-efficient after all on- *and* off-path histories.

Proposition 2. *There exists an optimal mechanism that is constrained Pareto-efficient after all histories.*

5 Comparative Statics

We can use our characterization of P's optimal mechanism to study how changes in the evolution of the Y affect J^*, the value of the optimal mechanism. It is easy to show that P's continuation value at t is decreasing in Y_t: an increase Y_t when $Y_t < M_t$ increases the probability P will have to lower his future demands and when $Y_t = M_t$ forces P to immediately lower his demands. A more interesting question is what the impact of an increase in the distribution tomorrow's outside option is, namely an increase in μ or ζ. Because an increase in Y_t hurts P, a natural guess is that an increase in the distribution of Y tomorrow will also hurt P.

Proposition 3 below shows this conjecture is incorrect and the comparative static in fact goes in the *opposite* direction. This result illustrates an important distinction between the *current level of* Y and the *expected level of* Y *tomorrow*: increasing the current level hurts P while increasing the expected level tomorrow benefits P.

Let Y', Y be two outside option processes that satisfy Assumption 2. We say that Y' is *more likely to increase* than Y if either both Y', Y are diffusion proccesses with respective drifts μ', μ such that $\mu'(y) \geq \mu(y)$ $\forall y \in [\underline{Y}, \overline{Y}]$ or both Y', Y are search processes with respective jump rates ζ', ζ such that $\zeta'(y) \geq \zeta(y)$ $\forall y \in [\underline{Y}, \overline{Y}]$.

Proposition 3. *If Y' is more likely to increase than Y, then P's value of bargaining is higher under Y' than Y.*

The intuition for these Proposition 2 is simple. Increasing μ or ζ makes the total "size of the pie" larger by increasing the value of experimentation for A. Keeping the mechanism fixed, as the value of experimentation increases, A's individual rationality constraint no longer binds. P can then increase his demands while still satisfying dynamic individual rationality and thereby extract the increase in the size of the pie.

However, an increase in the Y_0 increases the social planner's value from bargaining, so we can also view an increase in Y_0 as also increasing the size of the pie. Why does a increase in the size of the pie through μ or ζ change the sign of the comparative static on J^* when compared to an increase in Y_0? The difference comes from how the increase in value is allocated to players. The additional value created by an increase in the increase of Y_0 is already given entirely to A: due to the DIR constraint, P is forced to give A a larger continuation value and is not able to extract the additional surplus created. In contrast, when μ or ζ increase, A's participation constraint at $t = 0$ continues to bind and P alone benefits from the increase in surplus. The key difference is that an increase in μ or ζ does not improve A's bargaining position today.

Proposition 3 implies that P would benefit from taking actions, or encouraging A to take actions, that make A's expected outside option tomorrow better. Consider a firm P and worker A bargaining over a new contract upon the expiration of an old contract. P benefits from taking actions such as writing a positive letter of recommendation or permitting A to take time to interview with other firms even though these actions may increase the probability P is forced to increase his wage offer or A accepts an outside offer. Such actions are often observed in firm-worker relationships and are rationalized by our model.

Our next two results fix the expected change in Y_t by focusing martingale processes. Martingale processes arise naturally when changes in Y_t are driven by the arrival of information.

We start by showing that increasing the "speed" of the evolution of Y_t increases J^*. For a diffusion process, this corresponds to an increase in $\sigma(y)$. For a search process, this corresponds to an decrease in $\eta(y)$ when we decease $\zeta(y)$ an appropriate amount to ensure Y remains a martingale. Holding the mechanism fixed, these changes increase both P's and A's expected utility by reducing delay time.

Proposition 4. *If Y is a martingale below $\widetilde{Y}$, then J^* increases if $\sigma(y)$ or $-\eta(y)$ increase for all $y \in [\underline{Y}, \overline{Y}]$.*

The different structures of diffusion and search processes brings another natural question: would P rather changes in A's outside option come in small increments, as in a diffusion process, or in large jumps, as in a search process?

To answer this question, we look at a limiting case of a search process. A *pure jump process* for Y_t, when at Y_0, jumps down to some $y_- < Y_0$ at Poisson rate ζ_- and up to some $y_+ > Y_0$ at Poisson rate ζ_+; otherwise it is stationary. The optimal mechanism for a pure jump process can be shown to take the same form as in Theorem 1 (this is discussed in more detail in the Online Appendix).

In order to stay away from the limit as σ or $-\eta \to \infty$, in which delay becomes insignificant, we provide some discipline on the set of proccesses we compare by placing a bound on the expected speed of changes in Y_t. For some continuous concave function H, we impose a *capacity constraint* $-\mathbb{E}[\frac{d}{dt}H(Y_t)|Y_t] \leq dt$. Such constraints have been used in experimentation models (e.g., Zhong (2019)). When the evolution of Y_t is driven by the arrival of new information, the capacity constraint represents a cap on the amount of information that can be acquired in any instant.

Proposition 5. *For every diffusion process and search process that satisfies the capacity constraint, there is a pure jump process that satisfies the capacity constraint and generates higher expected utility for P.*

For an arbitrary diffusion or search process we find an upper bound on P's expected utility by solving the design problem for that process when only imposing an ex-ante participation constraint for A. This upper-bound mechanism uses a constant split threshold and demand function. We then find a pure jump process that increases P's utility when we keep the upper-bound mechanism fixed.

This result is driven by two forces. First, as shown by Zhong (2017), moving to a pure jump process increases P's payoff when we keep the upper bound mechanism fixed. Second, the pure jump process we construct is less susceptible to the distortions introduced by the dynamic individual rationality constraint. This is most easily seen when comparing it to a diffusion process. The upper bound mechanism is dynamically individually rational with a pure jump process but not with a diffusion process. Y_t will go above Y_0 with probability one with a diffusion process, at which point A's continuation value will be too low and P will then need to lower his demands. With our pure jump process, a jump up to y_+ is the first time Y_t goes above Y_0 and leads to P immediately taking his outside option.

When changes in Y are driven by the arrival of new information, Proposition

5 tells us about P's preference on the information structure. P benefits when information arrives in large, "decisive" amounts that trigger an immediate end to bargaining and prefers that beliefs stay the same in the absence of a decisive signal. This result illustrates the optimality of backloading information as a way to reducing the distortions caused by the arrival of non-decisive information through the dynamic individual rationality constraint.

6 Alternating Offers Bargaining

The assumption that P can commit to his bargaining strategy sets us apart from much of the bargaining literature. While we believe such an assumption is reasonable in many settings such as the firm-worker example, it is natural to ask how robust our results are to relaxing commitment. Our optimal mechanism relies on P promising to lower future demands. How credible are these promises without commitment? Using a canonical discrete time alternating offers protocol, Theorem 2 below shows that the outcome of P's optimal mechanism can be implemented in the frequent offer limit even without commitment power.

The alternating offers game is as follows. Players make demands in a pre-specified alternating order at $t = 0, \Delta, 2\Delta, ...$ A's outside option follows the discretization used in the proof of Theorem 1.[12] Within each period, both players first observe the realization of X_t, after which the proposing player i makes his demand. Player $k \neq i$ can then either accept, reject or take his outside option. If k accepts i's demand, the game ends and the agreed upon split is made. If k rejects i's demand, then i is given a chance to take his outside option or move to the next period. Both players discount by $e^{-r\Delta}$. We make Assumption 2 strict by assuming that there exists an $r' < r$ such that $e^{-r't}X_t$ is a supermartingale as $\Delta \to 0$; we note that this is satisfied when X_t is a supermartingale. Finally, to simplify the proof, we assume there is a public randomization device observed by both players at the beginning of each period.

Theorem 1 gives us strategies that yield an upper-bound on P's equilibrium value of bargaining. We need to consider if these strategies can be supported in equilibrium. It is not clear ex-ante whether this is possible. For example, when $X_t = M_t$, P promises to lower his demands at the split threshold. But when A's

[12] An earlier version of the paper showed Theorem 2 holds for more general discrete-time processes.

outside option reaches the split threshold and P makes his demand, he may be tempted to renege and increase his demand. Because past outside options cannot be taken, now that X_t is low, it may be rational for A to accept P's higher demand. Foreseeing this, back when $X_t = M_t$ A may not view P's promise to decrease his demand in the future as credible and may instead choose to take his outside option.

We construct off-path punishment equilibria to prevent this type of unraveling. When player i increases his demand at t higher than he is called to, we move to a punishment equilibrium in which k rejects i's demand and makes a high demand of his own demand at $t+\Delta$. If k can credibly threaten to take his outside option upon i rejecting at $t+\Delta$, i will find it optimal to accept k's high demand. This allows us to threaten i with a harsh punishment for increasing his demand at t. In the proof, we build subsequent off-path equilibria that make k's threat to take his outside option credible.

Theorem 2. *There exists a sequence of subgame-perfect equilibria with equilibrium payoff J^Δ for P when the period length is Δ such that $\lim_{\Delta \to 0} J^\Delta = J^*$.*

Although this is not quite an "anything goes" type of result (for example, equilibria must satisfy dynamic individual rationality constraints for both A and P), there exist a multiplicity of equilibria in this setting. We view Theorem 2 as a possibility result: If, as is standard in mechanism design, we allow P to choose his preferred equilibrium, the loss to P from relaxing commitment is negligible in the frequent offer limit

Our proof also allows us to, under some conditions, find an equilibrium that *exactly* achieves the commitment upper-bound at positive Δ. The equilibrium we construct in Theorem 2 uses the mechanism we derived in the discrete-time approximation of X_t in the proof of Theorem 1. When there is delay in the optimal mechanism (i.e., $S^*(X_0) < min\{\widetilde{Y}, X_0\}$), the discrete time strategies we use are able to achieve the commitment upper-bound for positive Δ.

Corollary 2. *If $S^*(Y_0) < min\{\widetilde{Y}, Y_0\}$, then J^Δ is equal to the value of P's optimal discrete time mechanism for sufficiently small Δ.*

7 Conclusion

In this paper, we study a bargaining game in which one player's outside option may change over time. A changing outside option leads to a rich set of dynamics in the optimal bargaining outcome when one side can commit to their demands. The committed party gradually decreases the demands he makes and the pressure being placed on the other party over the course of the game, with periods of intransigence followed by quick spurts of concession reminiscent of haggling. Our model shows a new interplay between demands and pressure and finds they are complementary in providing incentives to continue bargaining. Our results show how natural features of bargaining outcomes such as delay, gradual concessions and breakdowns can arise in a complete information bargaining model. We study how changes in the evolution to one parties' outside option change the value of the optimal mechanism, finding, for example, an increase in the drift of A's outside option may increase P's value of bargaining. We also explore how to relax the assumption that one party can commit to his demand process by studying a classic alternating-offers bargaining game, finding subgame perfect equilibrium that implements our optimal mechanism when the period length becomes small.

References

Dilip Abreu and Faruk Gul. Bargaining and reputation. *Econometrica*, 68:85–117, 2000.

Eitan Altman. *Constrained Markov decision processes.* CRC Press, 1999.

Matthew Backus, Thomas Blake, Brad Larsen, and Steven Tadelis. Sequential bargaining in the field: Evidence from millions of online bargaining interactions. *Quarterly Journal of Economics*, 135:1319–1361, 2020.

Rabi N. Bhattacharya and Edward C. Waymire. *Stochastic processes with applications.* Society for Industrial and Applied Mathematics, 2009.

Ken Binmore, Avner Shaked, and John Sutton. An outside option experiment. *The Quarterly Journal of Economics*, 104:753–770, 1989.

Simon Board and Marek Pycia. Outside options and the failure of the coase conjecture. *American Economic Review*, 104:656–671, 2014.

Isaıas N Chaves. Privacy in bargaining: The case of endogenous entry. Working Paper, 2019.

Olivier Compte and Philippe Jehiel. On the role of outside options in bargaining with obstinate parties.. *Econometrica*, 70:1477–1517, 2002.

Martin W Cripps. Markov bargaining games. *Journal of Economic Dynamics and Control*, 22:341–355, 1998.

Brendan Daley and Brett S. Green. Bargaining and news. Working Paper, 2018.

William Fuchs and Andrzej Skrzypacz. Bargaining with arrival of new traders. *American Economic Review*, 100:802–836, 2010.

Milton Harris and Bengt Holmstrom. A theory of wage dynamics. *Review of Economic Studies*, 49:315–333, 1982.

Ilwoo Hwang and Fei Li. Transparency of outside options in bargaining. *Journal of Economic Theory*, 167:116–147, 2017.

Jihong Lee and Qingmin Liu. Gambling reputation: Repeated bargaining with outside options. *Econometrica*, 81:1601–1672, 2013.

Andrew McClellan. Experimenation and approval mechanisms. Working Paper, 2020.

Antonio Merlo and Charles Wilson. A stochastic model of sequential bargaining with complete information. *Econometrica*, pages 371–399, 1995.

John Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.

Juan Ortner. Durable goods monopoly with stochastic costs. *Theoretical Economics*, 12:817–861, 2017.

Goran Peskir and Albert Shiryaev. *Optimal stopping and free-boundary problems.* Birkhauser Basel, 2006.

Debraj Ray. The time structure of self-enforcing agreements. *Econometrica*, pages 547–582, 2002.

Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, pages 97–109, 1982.

Yuliy Sannikov. A continuous-time version of the principal-agent problem. *Review of Economic Studies*, pages 957–984, 2008.

Jonathan Thomas and Tim Worrall. Self-enforcing wage contracts. *The Review of Economic Studies*, pages 541–554, 1988.

Weijie Zhong. Time preference and information acquisition. Working Paper, 2017.

Weijie Zhong. Optimal dynamic information acquisition. Working Paper, 2019.

Appendix A

In this Appendix we provide proofs of Theorems 1 and 2. All remaining proofs can be found in the Online Appendix.

Discrete Time Model

Probability Space: We replace Y with a Markovian stochastic process $X = \{X_t : t \in \{0, \Delta, ...\}\}$ on a filtered probability space $(\Omega^\Delta, \mathcal{F}^\Delta, (\mathcal{F}_t^\Delta)_{t\geq 0}, \mathbb{P}^\Delta)$. For a fixed sample point $\omega \in \Omega^\Delta$, the function $(t, X_0) \mapsto X_t(\omega, X_0)$ is the sample path associated with ω when starting at X_0; we assume this function is increasing in X_0.[13] We let $\tau(\omega, X_0)$ be the stopping time induced by (ω, X_0), with similar definitions for $d_\tau(\omega, X_0)$ and $\alpha_\tau(\omega, X_0)$.

Discretization: In each period $t \in \{0, \Delta, ...\}$, X_t and a public randomization device are observed by both players, P then makes a demand, after which A decides to accept, reject or take his outside option. If A rejects, we move to the next period which players discount by $e^{-r\Delta}$. The discrete time X_t moves on a countable grid $\mathcal{G}^\Delta = \{..., X^{-1}, X^0, X^1, ...\}$ where X^i is increasing in i and $X^0 = X_0$. The distance between a grid point $x = X^n$ and the next highest grid point is $\epsilon_x = X^{n+1} - X^n$.

Notation: We say $(\tau', d'_\tau, \alpha'_\tau)$ is the *continuation mechanism* of $(\tau, d_\tau, \alpha_\tau)$ after history $h_t = (x_0, ..., x_t)$ if $\tau'(\omega, X_0) = \tau(\omega, X_0) - t$, $d'_\tau(\omega, X_0) = d_\tau(\omega, X_0)$, and $\alpha'_\tau(\omega, X_0) = \alpha_\tau(\omega, X_0)$ for all $\omega \in \Omega^\Delta$ such that $(X_0, X_1(\omega, X_0), ..., X_t(\omega, X_0)) =$

[13]Because our discrete time $X_{t+\Delta}$ will be increasing in X_t in a *FOSD* sense, assuming the mapping from (t, X_0) into $X_t(\omega, X_0)$ is increasing in X_0 is without loss.

$(x_0, x_1, ..., x_t)$. A mechanism $(\tau, d_\tau, \alpha_\tau)$ is *stationary in* (X, M) if, for each (x, m), the distribution of outcomes is the same for all continuation mechanisms of $(\tau, d_\tau, \alpha_\tau)$ after every history h_t with $(X_t, M_t) = (x, m)$.

Diffusion Approximation: We follow the discretization in the text. Our grid $\mathcal{G}^\Delta$ is defined by $X^0 = X_0 = Y_0$ and $X^j = X^{j-1} + \epsilon$. Let $q_+(X_{t-\Delta}) = \mathbb{P}(w_t = \epsilon | X_{t-\Delta})$ and $q_-(X_{t-\Delta}) = \mathbb{P}(w_t = -\epsilon | X_{t-\Delta})$. To ensure for convergence of X to Y as $\Delta \to 0$, we let $\sigma_0 = \max\limits_{y \in [\underline{Y}, \overline{Y}]} \sigma(y)$ and set

$$\epsilon = \sigma_0 \sqrt{\Delta},\ q_+(x) = \frac{1}{2}(\frac{\sigma^2(x)}{\sigma_0^2} + \frac{\mu(x)}{\sigma_0}\sqrt{\Delta}),\ q_-(x) = \frac{1}{2}(\frac{\sigma^2(x)}{\sigma_0^2} - \frac{\mu(x)}{\sigma_0}\sqrt{\Delta}).$$

Convergence of X to Y as $\Delta \to 0$ follows from Theorem 4.1 of Bhattacharya and Waymire (2009). With some abuse of language, we will continue to call the random walk a diffusion process.

Search Approximation: We again follow the discretization in the text. Let $W_t(w) = \int_0^t \eta(W_s(w))ds$ with $W_0(w) = min\{\widetilde{Y}, w\}$; if no Z arrive in $[t, t+\Delta]$, then $Y_{t+\Delta} = W_\Delta(Y_t)$. We define our grid $\mathcal{G}^\Delta$ in the following way. Let $\widetilde{Y} = X^k \in \mathcal{G}^\Delta$ for some k. For $j \geq k$, let $X^{j+1} = min\{z \in supp(F) : \ z > X^j\}$. For $j < k$, let $X^{j-1} = W_\Delta(X^j)$. If $Y_0 > \widetilde{Y}$ and $Y_0 \notin supp(F)$, then we insert Y_0 into $\mathcal{G}^\Delta$. We take $X_0 = \min\limits_{x \in \mathcal{G}^\Delta} |x - Y_0|$ and rename indices so that $X^0 = X_0$. If $x \leq \widetilde{Y}$, we set $\zeta^\Delta(x) = 1 - exp(-\int_0^\Delta W_s(x)ds)$; if $x > \widetilde{Y}$, we set $\zeta^\Delta(x) = \zeta(x)\Delta$. We determine the value of Z in discrete time by drawing $Z' \sim F$ and rounding up to the nearest point in $\mathcal{G}^\Delta$. For an arbitrary function f, $\mathbb{E}_Z[f(Z)] = \sum_{z \in supp(F^\Delta)} f(z)\mathbb{P}^\Delta(Z = z)$. A formal construction of the discrete time search process is provided in the Online Appendix alongside the construction of the continuous time search process.

The solution to our discrete time problem will hold for general transition probabilities for X, not just those needed for convergence of X to Y, as long as $e^{-rt}X_t$ is a supermartingale, which we will assume throughout. By the convergence of Y, it is easily shown that $e^{-rt}X_t$ is a supermartingale for sufficiently small Δ if, for some $r' < r$, $e^{-r't}Y_t$ is a supermartingale. We can potentially run into problems with ensuring $e^{-rt}X_t$ is a supermartingale if $e^{-rt}Y_t$ is a martingale over some range of $[\underline{Y}, \overline{Y}]$. In such a case, the proof of Theorem 1 will proceed by deriving the optimal continuous time mechanism when the discount rate is $r'' > r$ and taking

$r'' \to r$.

Derivation of the Optimal Mechanism

Our first Lemma characterizes A's first-best mechanism and will allow us to partially characterize the optimal mechanism. If A takes his outside option at some x in his first-best mechanism, then, because A can achieve his first-best payoff by taking his outside option immediately, P cannot incentivize A to continue bargaining at x while satisfying dynamic individual rationality. We show that the set of such x takes a simple form. The proof is standard and deferred to the Online Appendix.

Lemma A.1. *For some S^A, R^A, A's first-best mechanism is*

$$\tau^A = min\{t : \; X_t \notin (S^A, R^A)\}, \; d^A_\tau = \mathbb{1}(X_\tau < R^A), \; \alpha^A_\tau = 1.$$

We restrict attention to mechanisms with $\tau \leq \tau_+(R^A)$ and $d_\tau = 0$ at $X_\tau \geq R^A$.

We now turn to characterizing the optimal mechanism in our relaxed problem 2. Our next Lemma shows that we can restrict attention to mechanisms which are stationary in (X, M). The proof is a relatively straightforward application of results in Altman (1999) and is therefore deferred to the Online Appendix.

Lemma A.2. *There exists an optimal mechanism that is stationary in (X, M).*

We let $(\tau^\Delta, d^\Delta_\tau, \alpha^\Delta_\tau)$ be an optimal mechanism which is stationary in (X, M). Players' continuation values after every history h_t with $(X_t, M_t) = (x, m)$ will be the same as at the first time s such that $(X_s, M_s) = (x, m)$. Define $\tau_1(x, m) = min\{t : \; (X_t, M_t) = (x, m)\}$. For expositional ease, we will henceforth say "$(\tau, d_\tau, \alpha_\tau)$ is the continuation mechanism at (x, m)" as short-hand for "$(\tau, d_\tau, \alpha_\tau)$ is the continuation mechanism for $(\tau^\Delta, d^\Delta_\tau, \alpha^\Delta_\tau)$ after a history $h_{\tau_1(x,m)}$."

Let $(\tau^{(x,m)}, d^{(x,m)}_\tau, \alpha^{(x,m)}_\tau)$ be the continuation mechanism of $(\tau^\Delta, d^\Delta_\tau, \alpha^\Delta)$ at (x, m) and $\tau^{(x,m)}_+(x')$ be the length of time after $\tau_1(x, m)$ until X goes above x': $\tau^{(x,m)}_+(x') = min\{s : \; s \geq \tau_1(x, m), \; X_s \geq x'\} - \tau_1(x, m)$. Similar to τ_+ and $\tau^{(x,m)}_+$, we define a lower threshold stopping rule $\tau_-(S) = min\{t : \; X_t \leq S\}$ and let $\tau^{(x,m)}_-(x') = min\{s : \; s \geq \tau_1(x, m), \; X_s \leq x'\} - \tau_1(x, m)$. When conditioning on (x, m) in an expectation, we mean conditional on the history $h_{\tau_1(x,m)}$.

With some abuse of notation, we let $J(x, m)$ and $V(x, m)$ be P's and A's respective continuation values from $(\tau^\Delta, d^\Delta_\tau, \alpha^\Delta_\tau)$ at the beginning of a period t

with $(X_t, M_t) = (x, m)$. We define $\mathcal{R}$ to be the set of (x, m) which are reached with positive probability under τ^Δ:

$$\mathcal{R} = \{(x,m) : \ \mathbb{P}\big(\exists t < \tau^\Delta \ s.t.\ (X_t, M_t) = (x,m)\big) > 0\}.$$

We use a Lagrangian approach to solve the relaxed problem in 2. By Theorem 9.10 of Altman (1999),[14] there exist multipliers $(\lambda(X^0), ..., \lambda(X^N)) \in \mathbb{R}_-^{N+1}$ such that 2 has a solution that solves:

$$\begin{aligned}
\mathcal{L}(X_0) = \max_{(\tau, d_\tau, \alpha_\tau)} \ & \mathbb{E}\Big[e^{-r\tau}\big(d_\tau\{u_P(\alpha_\tau) - \nu - \lambda(X^0)(u_A(1-\alpha_\tau) - X_\tau)\} + \nu - \lambda(X^0)X_\tau\big) \\
& + \sum_{n=1}^{N} \lambda(X^n)\big\{e^{-r(\tau\wedge\tau_+(X^n))}\big(d_\tau(X^n)(u_A(1-\alpha_\tau) - X_{\tau\wedge\tau_+(X^n)}) + X_{\tau\wedge\tau_+(X^n)}\big) \\
& - e^{-r\tau}\big(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau\big)\big\}\Big] + \lambda(X^0)X_0. \qquad (4)
\end{aligned}$$

Our next Lemma uses 4 to characterize when, prior to $\tau_+(X^1)$, a split is reached and what the split amount is. A standard solution to optimal stopping problems is to stop the first time some chosen threshold is crossed. We show that the same type of solution holds prior to $\tau_+(X^1)$ when we generalize the notion of the stopping threshold. We say a mechanism uses a (S, γ)-*split threshold* if it stops and implements a split with probability $\gamma \in [0, 1)$ in each period t with $X_t = S + \epsilon_S$ and with probability one when $X_t \le S$.

Lemma A.3. *If a split is made with positive probability prior to $\tau_+(X^1)$, then there are unique α^0 and S' such that every stationary solution to 4 has $\alpha_\tau = \alpha^0$ if $\tau < \tau_+(X^1)$ and uses an (S^0, γ^0)-split threshold prior to $\tau_+(X^1)$ for some (S^0, γ^0) such that S^0 is either S' or $\max\{x \in \mathcal{G}^\Delta : x < S'\}$*

Proof. From 4, if a split is made before $\tau_+(X^1)$, then α_τ is equal to α^0 as defined in the text, which is independent of X_τ and unique by our concavity assumption.

If it is strictly optimal to stop immediately at $t = 0$, then we are done. Suppose it is not. Let $K(x')$ be the continuation value in 4 at $\tau_+(X^1)$ when $X_{\tau_+(X^1)} = x'$. Let $L(x, \tau, d_\tau)$ be the continuation value in 4 at $(X_t, M_t) = (x, X_0)$ if we

[14] Theorem 9.10 assumes a Slater-type condition which is satisfied in our problem by the mechanism $(\tau^A, d_\tau^A, \alpha_\tau^A)$. Although Theorem 9.10 is stated in terms of a "total cost" formulation of the constrained optimization problem, Chapter 10 of Altman (1999) shows how to translate the the discounted cost structure of our problem into the total cost formulation.

use a continuation mechanism (τ, d_τ, α^0) until $\tau_+(X^1)$ and revert to the optimal mechanism at $\tau_+(X^1)$:

$$\begin{aligned}L(x,\tau,d_\tau) =&\mathbb{E}_x\big[e^{-r\tau}\big\{d_\tau[u_P(\alpha^0) - \nu - \lambda(X^0)(u_A(1-\alpha^0) - X_\tau)] \\ &+ \nu - \lambda(X^0)X_\tau\big\}\mathbb{1}(\tau < \tau_+(X^1)) + e^{-r\tau_+(X^1)}K(X_{\tau_+(X^1)})\mathbb{1}(\tau \geq \tau_+(X^1))\big].\end{aligned}$$

Define $L^*(x) = \max\limits_{(\tau,d_\tau)} L(x,\tau,d_\tau)$. Then $\mathcal{L}(X_0) = L(X_0, \tau^\Delta, d_\tau^\Delta) = L^*(X_0)$. Let $D(x)$ be the value of stopping at $X_t = x$ when $t < \tau_+(X^1)$ and $\mathcal{C}$ be the continuation region:

$$\begin{aligned}&D(x) = max\{u_P(\alpha^0) - \lambda(X^0)u_A(1-\alpha^0), \nu - \lambda(X^0)x\}, \\ &\mathcal{C} = \{x:\ x < X^1,\ \exists(\tau, d_\tau)\ s.t.\ L(x,\tau,d_\tau) \geq D(x)\ and\ \mathbb{P}(\tau > 0) > 0.\}.\end{aligned}$$

By standard arguments it is weakly optimal to continue at x if $x \in \mathcal{C}$ and strictly optimal to stop immediately at x if $x \notin \mathcal{C}$.

A useful property of X is that the distribution of $X_{\tau_+(X^1)}$ is independent of the current X_t. Therefore, the expected continuation value at $\tau_+(X^1)$ when $X_0 = x$ (i.e., $\mathbb{E}_x[K(X_{\tau_+(X^1)})]$) is the same for all x. Because it is not strictly optimal to stop immediately at X_0 and $D(X_0) \geq D(x)\ \forall x \leq X_0$, it must be that $\mathbb{E}_{X_0}[K(X_{\tau_+(X^1)})] > D(X_0)$.

Take any $x' \in \mathcal{C}$ such that $(x', X_0) \in \mathcal{R}$. We will show $L^*(x'') > L^*(x')\ \forall x'' > x'$. Let $(\tau_{x'}, d_{\tau,x'}) \in argmax_{(\tau,d_\tau)}\ L(x',\tau,d_\tau)$ such that $\mathbb{P}(\tau_{x'} > 0) > 0$ and $x'' > x'$. Define $\tilde{\tau}(\omega, x'') = \tau_{x'}(\omega, x')$ and $\tilde{d_\tau}(\omega, x'') = d_{\tau,x'}(\omega, x')$. For each ω, we have $\tau_+(X^1)(\omega, x'') \leq \tau_+(X^1)(\omega, x')$ and so

$$\tilde{\tau}(\omega, x'') \wedge \tau_+(X^1)(\omega, x'') \leq \tau_{x'}(\omega, x') \wedge \tau_+(X^1)(\omega, x').$$

Because $X_t(\omega, x'') \geq X_t(\omega, x')$, the value of stopping before $\tau_+(X^1)$ is weakly higher when starting at x''. Additionally, $\mathbb{P}(\tilde{\tau} \geq \tau_+(X^1)|x'') > \mathbb{P}(\tau_{x'} \geq \tau_+(X^1)|x')$, which increases the value of L because $\mathbb{E}_x[K(X_{\tau_+(X^1)})] > D(x)\ \forall x < X^1$. Therefore, $L^*(x'') \geq L(x'', \tilde{\tau}, \tilde{d_\tau}) > L^*(x')$.

Let $S' = max\{x:\ L^*(x) = u_P(\alpha^0) - \lambda(X^0)u_A(1-\alpha^0), x \leq X_0\}$. If there was an $x \in \mathcal{C}$ such that $x < S'$ and $(x, X_0) \in \mathcal{R}$, then $L^*(x) \geq D(x) = D(S') = L^*(S')$, a contradiction. Therefore, if it is weakly optimal to immediately implement a split at S', then it is strictly optimal to immediately implement a split at all $x < S'$ which are reached with positive probability. Every solution to 4 must either stop at S' with probability one or, if $S' \in \mathcal{C}$, mix at S' and stop with probability one at $max\{x \in \mathcal{G}^\Delta : x < S'\}$. $\square$

The continuation mechanism at $\tau_+(X^1)$ when $X_{\tau_+(X^1)} = X^\ell$ solves $K(X^\ell)$. Defining $\underline{\lambda}(X^\ell) = \sum_{k\le\ell}\lambda(X^k)$, $K(X^\ell)$ is equal to

$$\max_{(\tau,d_\tau,\alpha_\tau)} \mathbb{E}_{X^\ell}\Big[e^{-r\tau}\Big(d_\tau\big\{u_P(\alpha_\tau)-\nu-\underline{\lambda}(X^\ell)(u_A(1-\alpha_\tau)-X_\tau)\big\}+\nu-\underline{\lambda}(X^\ell)X_\tau\Big) \\ +\sum_{j>\ell}\lambda(X^j)\big\{e^{-r(\tau\wedge\tau_+(X^j))}\Big(d_\tau(X^j)\big(u_A(1-\alpha_\tau)-X_{\tau\wedge\tau_+(X^j)}\big)+X_{\tau\wedge\tau_+(X^j)}\Big)\Big) \\ -e^{-r\tau}(d_\tau(u_A(1-\alpha_\tau)-X_\tau)+X_\tau)\big\}\Big]+\sum_{k\le\ell}\lambda(X^k)X^k.$$

The similarity to 4 is clear and we can apply the same arguments as in Lemma A.3 when solving $K(X^\ell)$ to conclude that a (S,γ)-split threshold and constant split amount are used until $\tau_+(X^\ell+\epsilon_{X^\ell})$. Repeating these arguments, we conclude that, for some functions $S(\cdot)$, $\gamma(\cdot)$ and $\alpha(\cdot)$, when $M_t = m$, the optimal continuation mechanism from $\tau_+(m)$ to $\tau_+(m+\epsilon_m)$ uses a $(S(m),\gamma(m))$-split threshold and a split gives P $\alpha(m)$.

Complementary slackness conditions imply $\lambda(X^n) = 0$ whenever $RDIR(X^n)$ is slack. It is easily verified that α, S and γ only change when we reach a new X^ℓ for which $RDIR(X^\ell)$ binds. Otherwise, they remain constant. For any $X^k < X^\ell$ such that $RDIR(X^n)$ is slack for all $X^n \in (X^k, X^\ell)$, $S(m), \alpha(m), \gamma(m), J(x,m)$ and $V(x,m)$ are constant in m for all $m \in [X^k, X^\ell)$.

Before providing more structure to $\alpha(\cdot)$ and $S(\cdot)$, it is useful to look at P's and A's continuation values. The next Lemma pins down properties of $V(x,m)$ and is important for establishing that our mechanism is dynamically individually rational. In order to not continually specify whether we are at an (x,m) which can be reached in the optimal mechanism, it is useful to specify "off-path" continuation mechanisms. For any $(x,m) \notin \mathcal{R}$ with $x < min\{R^A, m\}$, we specify that the continuation mechanism at date t with $(X_t, M_t) = (x,m)$ is A's first-best mechanism; for all other $(x,m) \notin \mathcal{R}$, the continuation mechanism takes the outside option immediately.[15]

Lemma A.4. *$V(x,m) \ge x$, with equality if $x = m$.*

Proof. The desired properties hold at (m,m) with $m \ge R^A$ and all $(x,m) \notin \mathcal{R}$. We proceed by induction. Suppose $m \ge X_0$ is such that $V(m',m') = m'$ $\forall m' > m$.

[15] Even though A's continuation value is higher in these off-path continuation mechanisms, the decision of when to stop is entirely in the hands of P, so there is no deviation by A which can lead to these off-path mechanisms.

Let $I_-(m) = \mathbb{1}(\tau^\Delta < \tau_+(m))$, and $I_+(m', m) = \mathbb{1}(\tau^\Delta \geq \tau_+(m),\ X_{\tau_+(m)} = m')$. $RDIR(m)$ implies

$$
\begin{aligned}
&\mathbb{E}[e^{-r\tau^\Delta}(d_\tau^\Delta(u_A(1-\alpha_\tau^\Delta) - X_{\tau^\Delta}) + X_{\tau^\Delta})I_-(m)] + \sum_{m' \geq m} \mathbb{E}[e^{-r\tau_+(m)} m' I_+(m', m)] \\
&\leq \mathbb{E}[e^{-r\tau^\Delta}\big(d_\tau^\Delta(u_A(1-\alpha_\tau^\Delta) - X_{\tau^\Delta}) + X_{\tau^\Delta}\big)I_-(m)] + \sum_{m' \geq m} \mathbb{E}[e^{-r\tau_+(m)} V(m', m')\ I_+(m', m)], \\
&\Rightarrow \sum_{m' \geq m} \mathbb{E}[e^{-r\tau_+(m)} m' I_+(m', m)] \leq \sum_{m' \geq m} \mathbb{E}[e^{-r\tau_+(m)} V(m', m') I_+(m', m)]. \qquad (5)
\end{aligned}
$$

Because $V(m', m') = m'\ \forall m' > m$, we have $V(m, m) \geq m$, with equality if $RDIR(m)$ binds.

Suppose there is an $(x, m) \in \mathcal{R}$ such that $V(x, m) \leq x$. Define $\check{\tau} = \tau^{(m,m)} \wedge \tau_-^{(m,m)}(x) \wedge \tau_+^{(m,m)}(m + \epsilon_m)$-namely, the length of time after $\tau_1(m, m)$ until the mechanism stops or X_t exits the interval $(x, m + \epsilon_m)$. If $x > S(m)$, then $\tau^{(m,m)} < \tau_-^{(m,m)}(x) \wedge \tau_+^{(m,m)}(m + \epsilon_m)$ implies $d_\tau^{(m,m)} = 0$. If $x = S(m)$, then $V(x, m) \leq x$ implies $u_A(1 - \alpha(m)) \leq x$, so A's utility at $\tau^{(m,m)}$ when $\tau^{(m,m)} < \tau_-^{(m,m)}(x) \wedge \tau_+^{(m,m)}(m + \epsilon_m)$ is always weakly less than $X_{\tau^{(m,m)}}$. Together, $V(x, m) \leq x$ and $V(m', m') = m'\ \forall m' > m$ imply

$$
\begin{aligned}
V(m, m) = \mathbb{E}\big[e^{-r\check{\tau}}\big(&V(x, m)\mathbb{1}(\check{\tau} = \tau_-^{(m,m)}(x)) \\
&+ \big(d_\tau^{(m,m)}(u_A(1 - \alpha(m)) - X_{\tau^{(m,m)}}) + X_{\tau^{(m,m)}}\big) \\
&\cdot \mathbb{1}(\check{\tau} = \tau^{(m,m)} < \tau_-^{(m,m)}(x) \wedge \tau_+^{(m,m)}(m + \epsilon_m)) \\
&+ V(X_{\tau_+^{(m,m)}(m+\epsilon_m)}, X_{\tau_+^{(m,m)}(m+\epsilon_m)})\mathbb{1}(\check{\tau} = \tau_+^{(m,m)}(m + \epsilon_m))\big)|(m, m)\big] \\
&\leq \mathbb{E}[e^{-r\check{\tau}} X_{\check{\tau}}|(m, m)] \\
&\leq m, \qquad (6)
\end{aligned}
$$

where the final inequality follows from Doob's Optional Stopping Theorem. The first inequality is strict if $V(x, m) < x$, which contradicts $V(m, m) \geq m$. Therefore, when $V(m', m') = m'\ \forall m' > m$, we have $V(x, m) \geq x\ \forall x \leq m$.

We conclude the inductive step by showing $V(m, m) = m$. Suppose $V(m, m) > m$. Let $\hat{m} = max\{x :\ x < m,\ RDIR(x)\ binds,\ (x, x) \in \mathcal{R}\}$.[16] For all $x \in [\hat{m}, m]$, $V(x, m')$ is constant in $m' \in [x, m]$ because all $RDIR(m')$ constraints are

[16] $\hat{m}$ exists because $RDIR(X_0)$ binds; if it were slack P could increase his demand before $\tau_+(X^1)$ without violating any constraints.

slack. Because $V(\hat{m},\hat{m}) = V(\hat{m},m)$, taking $x = \hat{m}$ in 6, we get $V(m,m) \leq m$ if $V(\hat{m},\hat{m}) \leq \hat{m}$. Therefore $V(m,m) > m$ implies $V(\hat{m},\hat{m}) > \hat{m}$.

Because $RDIR(\hat{m})$ is binding, the same steps as in 5 imply

$$\sum_{m' \geq \hat{m}} \mathbb{E}[e^{-r\tau_+(\hat{m})} m' I_+(m',\hat{m})] = \sum_{m' \geq \hat{m}} \mathbb{E}[e^{-r\tau_+(\hat{m})} V(m',m') I_+(m',\hat{m})]. \tag{7}$$

For a diffusion process, because $X_{\tau_+(\hat{m})} = \hat{m}$, 7 implies $V(\hat{m},\hat{m}) = \hat{m}$, a contradiction. For a search process, 7 and $V(\hat{m},\hat{m}) > \hat{m}$ imply $V(m'',m'') < m''$ for some $m'' \in (\hat{m},m)$ with $\mathbb{P}(X_{\tau_+(\hat{m})} = m'') > 0$. By our previous arguments, $V(m'',m'') = V(m'',m) \geq m''$, a contradiction. Therefore, $V(m,m) = m$. $\square$

The next Lemma shows that P's continuation value is always higher than his outside option.

Lemma A.5. *$J(x,m) \geq \nu$ $\forall (x,m) \in \mathcal{R}$ with strict inequality if $x < m$. If P doesn't take his outside option with probability one when $(X_t, M_t) = (m+\epsilon_m, m+\epsilon_m)$, then $J(m,m) > J(m+\epsilon_m, m+\epsilon_m)$.*

Proof. If $\nu > J(m,m)$, P is better off taking his outside option immediately at (m,m). Doing so does not change A's continuation value because $V(m,m) = m$ by Lemma A.4. Therefore, $J(m,m) \geq \nu$ $\forall m$ such that $(m,m) \in \mathcal{R}$.

We show $J(x,m) \geq \nu$ by induction. For the base case, let $\bar{m} = max\{m' : (m',m') \in \mathcal{R}\}$. Then $J(\bar{m},\bar{m}) \geq \nu$ and, vacuously, $J(\bar{m},\bar{m}) \geq J(m',m')$ for all $m' > \bar{m}$ such that $(m',m') \in \mathcal{R}$. We show the inductive step in two steps. In the first step, for an arbitrary m such that $(m,m) \in \mathcal{R}$, we show that if $J(m,m) \geq J(m',m') \geq \nu$ for all $m' > m$ such that $(m',m') \in \mathcal{R}$, then $J(x,m) \geq \nu$ for all $(x,m) \in \mathcal{R}$. In the second step, we let $m_- = max\{m' : m' < m \text{ and } (m',m') \in \mathcal{R}\}$ and show that $J(m_-,m_-) \geq J(m,m)$.

Take an m such that $J(m,m) \geq J(m',m') \geq \nu$ for all $(m',m') \in \mathcal{R}$ with $m' > m$. If $(x,m) \in \mathcal{R}$, then $(x',m) \in \mathcal{R}$ for all $x' \in [x,m]$ because the only way to reach (x,m) from (m,m) is to go through (x',m). We show, for any $(x,m) \in \mathcal{R}$ and $x' \in [x,m]$, that $J(x,m) \geq J(x',m) \geq \nu$. This is done by induction as well. For the base case, take $x = m$. For the inductive step, suppose $J(x,m) \geq J(x',m) \geq \nu$ $\forall x' \in [x,m]$. Let $x_- = max\{x'' \in \mathcal{G}^\Delta : x'' < x\}$. If $(x_-,m) \notin \mathcal{R}$, we are done.

Suppose $(x_-, m) \in \mathcal{R}$. Note that

$$\begin{aligned}
J(x,m) = {} & \mathbb{E}[e^{-r\tau_-^{(x,m)}(x_-)} J(x_-,m)\mathbb{1}(\tau_-^{(x,m)}(x_-) \leq \tau_+^{(x,m)}(x+\epsilon_x) \wedge \tau^{(x,m)})) \\
& + e^{-r\tau^{(x,m)}}(d_\tau^{(x,m)}(u_P(\alpha(m)) - \nu) + \nu)\mathbb{1}(\tau^{(x,m)} < \tau_-^{(x,m)}(x_-) \wedge \tau_+^{(x,m)}(x+\epsilon_x)) \\
& + e^{-r\tau_+^{(x,m)}(x+\epsilon_x)} J(X_{\tau_+^{(x,m)}(x+\epsilon_x)}, max\{m, X_{\tau_+^{(x,m)}(x+\epsilon_x)}\}) \\
& \cdot \mathbb{1}(\tau_+^{(x,m)}(x+\epsilon_x) \leq \tau_-^{(x,m)}(x_-) \wedge \tau^{(x,m)})|(x,m)], \\
& \leq \mathbb{E}[e^{-r\tau_-^{(x,m)}(x_-)} J(x_-,m)\mathbb{1}(\tau_-^{(x,m)}(x_-) \leq \tau_+^{(x,m)}(x+\epsilon_x) \wedge \tau^{(x,m)})) \\
& + e^{-r\tau^{(x,m)}}(d_\tau^{(x,m)}(u_P(\alpha(m)) - \nu) + \nu)\mathbb{1}(\tau^{(x,m)} < \tau_-^{(x,m)}(x_-) \wedge \tau_+^{(x,m)}(x_+)) \\
& + e^{-r\tau_+^{(x,m)}(x+\epsilon_x)} J(x,m)\mathbb{1}(\tau_+^{(x,m)}(x+\epsilon_x) \leq \tau_-^{(x,m)}(x_-) \wedge \tau^{(x,m)})|(x,m)].
\end{aligned}$$

Thus, due to discounting, $J(x,m)$ is strictly less than either $J(x_-,m)$ or, when $\tau^{(x,m)} < \tau_-^{(x,m)}(x_-) \wedge \tau_+^{(x,m)}(x_+)$, $d_\tau^{(x,m)}(u_P(\alpha(m)) - \nu) + \nu$. If $S(m) = x_-$, then $J(x_-,m) = u_P(\alpha(m))$, which, because $J(x,m) \geq \nu$ and $d_\tau^{(x,m)}(u_P(\alpha(m)) - \nu) + \nu \in \{u_P(\alpha(m)), \nu\}$, implies $J(x_-,m) > J(x,m)$. If $S(m) < x_-$, then $d_\tau^{(x,m)} = 0$ whenever $\tau^{(x,m)} < \tau_-^{(x,m)}(x_-) \wedge \tau_+^{(x,m)}(x+\epsilon_x)$, which implies $J(x_-,m) > J(x,m)$. This concludes the inductive step and shows $J(x,m) \geq \nu$ for all x such that $(x,m) \in \mathcal{R}$, with strict inequality if $x < m$.

We next prove the second step. If P takes his outside option at (m,m), then $J(m,m) = \nu$ and we are done because $J(m_-,m_-) \geq \nu$. Suppose P doesn't take his outside option with probability one at (m,m). If $S(m) \geq m_-$, then $\alpha(m) \leq 1 - u_A^{-1}(S(m))$ because $V(S(m),m) \geq S(m)$. P could make a TIOLI demand of $1 - u_A^{-1}(m_-)$ at (m_-,m_-) and without changing A's continuation value because $V(m_-,m_-) = m_-$. This implies

$$J(m_-,m_-) \geq u_P(1 - u_A^{-1}(m_-)) > u_P(1 - u_A^{-1}(S(m))) \geq J(m,m).$$

If $S(m) < m_-$, then by our arguments above $J(m_-,m) > J(m,m)$. Because $V(m_-,m) \geq m_- = V(m_-,m_-)$, at (m_-,m_-) P could use the continuation mechanism at (m_-,m) without violating A's $RDIR$ constraints. Therefore, $J(m_-,m_-) \geq J(m_-,m) > J(m,m) \geq \nu$. □

The next two Lemmas add more structure to $(\tau^\Delta, d_\tau^\Delta, \alpha_\tau^\Delta)$.

Lemma A.6. $d_\tau^\Delta = \mathbb{1}(X_\tau < \overline{R})$ *for some* $\overline{R} \geq X_0$.

Proof. Suppose P takes his outside option at $(X_t, M_t) = (x, m) \in \mathcal{R}$ with $x < m$ with probability $p \in (0, 1]$. Because $V(x, m) \geq x$, A's continuation value when P does not take his outside option today is weakly greater than x. Because A weakly prefers to continue bargaining, for $p \in (0, 1]$ to be optimal, it must be that $\nu > J(x, m)$, contradicting Lemma A.5. Therefore, there is an optimal mechanism stationary in (X, M) in which P does not take his outside option at any $x < m$. Let $\overline{R} = min\{m : J(m, m) = \nu\}$. Lemma A.5 shows $J(\overline{R}, \overline{R}) \geq J(m, m)$ $\forall m > \overline{R}$ with $(m, m) \in \mathcal{R}$ and so it is optimal for P to take his outside option at such m. □

Lemma A.7. *$S(m)$ and $\alpha(m)$ are decreasing in m.*

Proof. The monotonicity of $\alpha(\cdot)$ follows from pointwise optimization of α in 4 at $(X_t, M_t) = (x, m)$:

$$\alpha(m) = argmax_{\alpha \in [0,1]}\ u_P(\alpha) - \sum_{X^n \leq m} \lambda(X^n) u_A(1 - \alpha).$$

Because $\lambda(X^n) \leq 0$, $\alpha(m)$ is decreasing in m. It then immediately follows that, at the optimal choice of $S(m)$ and $\gamma(m)$, P's utility is strictly increasing in $S(m)$ and $\gamma(m)$, as increasing either decreases both delay time and the chance that P's future demand will decrease. This means that A's utility must be strictly decreasing in $S(m)$ and $\gamma(m)$; otherwise it would be optimal for P to increase $S(m)$ or $\gamma(m)$.

Next, we argue that $S(\cdot)$ is decreasing. Take some $m < \overline{R}$. By stationarity of the mechanism between $\tau_+(m)$ and $\tau_+(m + \epsilon_m)$, both players' continuation value will be the same at $\tau_+^{(m,m)}(m + \epsilon_m)$ and $\tau_+^{(m,m+\epsilon_m)}(m + \epsilon_m)$. Starting at $X_0 = m$, we consider an alternative problem in which we fix continuation values at $\tau_+(m + \epsilon_m)$ to be $J(X_{\tau_+(m+\epsilon_m)}, X_{\tau_+(m+\epsilon_m)})$ and $X_{\tau_+(m+\epsilon_m)}$ for P and A respectively and allow P to choose S, γ and α. We also impose a promise-keeping constraint to deliver W expected utility to A.

Let $\tau^{S,\gamma}$ be the stopping time induced by an (S, γ)-split threshold. We define $\Phi(S, \gamma) = \mathbb{E}_m[e^{-r\tau_+(m+\epsilon_m)} \mathbb{1}(\tau^{S,\gamma} > \tau_+(m + \epsilon_m))]$ and $\phi(S, \gamma) = \mathbb{E}_m[e^{-r\tau^{S,\gamma}} \mathbb{1}(\tau^{S,\gamma} < \tau_+(m + \epsilon_m))]$. The promise-keeping constraint for W pins down the split amount $\alpha(W, S, \gamma)$:

$$W = \Phi(S, \gamma)\mathbb{E}[X_{\tau_+(m+\epsilon_m)}] + \phi(S, \gamma) u_A(1 - \alpha(W, S, \gamma)),$$
$$\Rightarrow \alpha(W, S, \gamma) = 1 - u_A^{-1}\left(\frac{W - \Phi(S, \gamma)\mathbb{E}[X_{\tau_+(m+\epsilon_m)}]}{\phi(S, \gamma)}\right).$$

P's problem is to choose S, γ subject to $\alpha = \alpha(W, S, \gamma)$:

$$\hat{J}(W) := \max_{S,\gamma} \ \Phi(S,\gamma)\mathbb{E}[J(X_{\tau_+(m+\epsilon_m)}, X_{\tau_+(m+\epsilon_m)})] + \phi(S,\gamma)u_P(\alpha(W,S,\gamma)). \quad (8)$$

$\hat{J}$ is decreasing in W. When $W = V(m,m)$, we have $\hat{J}(W) = J(m,m)$ and the solution to 8 is $S(m), \gamma(m)$. When $W = V(m, m+\epsilon_m)$, $\hat{J}(W) = J(m, m+\epsilon_m)$ and the solution to 8 is $S(m+\epsilon_m), \gamma(m+\epsilon_m)$.

Compare two (S,γ)-split thresholds (S_1,γ_1) and (S_2,γ_2) and let $\alpha_j(W) = \alpha(W, S_j, \gamma_j)$. P prefers (S_1,γ_1) to (S_2,γ_2) if

$$\begin{aligned}&\Phi(S_1,\gamma_1)\mathbb{E}[J(X_{\tau_+(m+\epsilon_m)}, X_{\tau_+(m+\epsilon_m)})] + \phi(S_1,\gamma_1)u_P(\alpha_1(W)) \\ &- \left[\Phi(S_2,\gamma_2)\mathbb{E}[J(X_{\tau_+(m+\epsilon_m)}, X_{\tau_+(m+\epsilon_m)})] + \phi(S_2,\gamma_2)u_P(\alpha_2(W))\right] \geq 0.\end{aligned} \quad (9)$$

When we increase W, 9 changes by

$$\phi(S_1,\gamma_1)u_P'(\alpha_1(W))\alpha_1'(W) - \phi(S_2,\gamma_2)u_P'(\alpha_2(W))\alpha_2'(W). \quad (10)$$

Using $\alpha_i'(W) = \frac{-1}{\phi(S_i,\gamma_i)u_A'(1-\alpha_i(W))}$, 10 becomes

$$\frac{u_P'(\alpha_2(W))}{u_A'(1-\alpha_2(W))} - \frac{u_P'(\alpha_1(W))}{u_A'(1-\alpha_1(W))}.$$

By concavity of u_P and u_A, 10 is strictly positive if and only if $\alpha_1(W) > \alpha_2(W)$.

We argue the optimal S is decreasing in W for $W \in [V(m,m), V(m, m+\epsilon_m)]$. Consider a point W' at which P is indifferent between two optimal choices (S_1,γ_1) and (S_2,γ_2) with $S_1 < S_2$. For the sake of contradiction, suppose $\alpha_1(W') \leq \alpha_2(W')$. Let τ^i be the stopping time induced by (S_i,γ_i)-split threshold in the alternative problem. Because $\tau^2 \leq \tau^1$, P's indifference between (S_1,γ_1) and (S_2,γ_2) implies

$$\begin{aligned}\hat{J}(W') =& \mathbb{E}_m[e^{-r\tau^2}u_P(\alpha_2(W'))\mathbb{1}(\tau^2 < \tau_+(m+\epsilon_m))] \\ &+ \mathbb{E}_m[e^{-r\tau_+(m+\epsilon_m)}J(X_{\tau_+(m+\epsilon_m)}, X_{\tau_+(m+\epsilon_m)})\mathbb{1}(\tau^2 > \tau_+(m+\epsilon_m))] \\ =& \mathbb{E}_m\Big[e^{-r\tau^2}\mathbb{1}(\tau^2 < \tau_+(m+\epsilon_m))\mathbb{E}_m\big[e^{-r(\tau^1-\tau^2)}u_P(\alpha_1(W'))\mathbb{1}(\tau^1 < \tau_+(m+\epsilon_m)) \\ &+ e^{-r(\tau_+(m+\epsilon_m)-\tau^2)}\mathbb{E}[J(X_{\tau_+(m+\epsilon_m)}, X_{\tau_+(m+\epsilon_m)})]\mathbb{1}(\tau^1 > \tau_+(m+\epsilon_m))|h_{\tau^2}\big]\Big] \\ &+ \mathbb{E}_m[e^{-r\tau_+(m+\epsilon_m)}J(X_{\tau_+(m+\epsilon_m)}, X_{\tau_+(m+\epsilon_m)})\mathbb{1}(\tau^2 > \tau_+(m+\epsilon_m))], \\ \Rightarrow& \mathbb{E}[J(X_{\tau_+(m+\epsilon_m)}, X_{\tau_+(m+\epsilon_m)})] > u_P(\alpha_2(W')), \\ \Rightarrow& \hat{J}(W') < \mathbb{E}[J(X_{\tau_+(m+\epsilon_m)}, X_{\tau_+(m+\epsilon_m)})].\end{aligned}$$

By Lemma A.5, $J(m, m + \epsilon_m) \geq J(m', m')\ \forall m' \geq m + \epsilon_m$, so

$$\hat{J}(W') \geq \hat{J}(V(m, m + \epsilon_m)) = J(m, m + \epsilon_m) \geq \mathbb{E}[J(X_{\tau_+(m+\epsilon_m)}, X_{\tau_+(m+\epsilon_m)})],$$

a contradiction. Thus, $\alpha_1(W') > \alpha_2(W')$ and P strictly prefers (S_1, γ_1) at higher W. Thus, the optimal choice of S is decreasing in W and $S(m + \epsilon_m) \leq S(m)$. $\square$

We now look in the limit as $\Delta \to 0$. For each Δ, let $S^\Delta(M_t), \alpha^\Delta, \overline{R}^\Delta$ be the respective optimal approval threshold, split amount and breakdown threshold in the discrete time relaxed problem. Let $S^*, \alpha^*, \overline{R}^* = \lim_{\Delta \to 0} S^\Delta, \alpha^\Delta, \overline{R}^\Delta$.[17] Use $S^*, \alpha^*, \overline{R}^*$ to define the mechanism $(\tau^*, d^*_\tau, \alpha^*_\tau)$ as in Theorem 1. If $S^*(Y_0) \geq Y_0$, we take $\tau^* = 0, d^*_\tau = \mathbb{1}(Y_0 < \overline{R})$ and $\alpha^*_\tau = 1 - u_A^{-1}(Y_0)$.

The limit threshold and demand functions S^* and α^* are continuous. This result follows from the fact that $V(m, m) = m$ and, as noted in Lemma A.7, A's utility is decreasing in the S^Δ and α^Δ. A discrete jump down in S^* or α^* at m would mean that A strictly preferred to continue bargaining at $(X_t, M_t) = (m, m)$ for small Δ; P could then increase S^* and α^* without violating dynamic individual rationality. The proof is deferred to the Online Appendix.

Lemma A.8. *S^* and α^* are continuous.*

Before formally proving Theorem 1, we note an auxiliary result that shows that, under some conditions, the mechanism $(\tau^\Delta, d^\Delta_\tau, \alpha^\Delta_\tau)$ is the optimal discrete time dynamically individually rational mechanism. By Lemma A.4, to show dynamic individual rationality it suffices to show that A finds it optimal to accept a split when offered one. We show that when when there is non-vanishing delay in reaching a split as $\Delta \to 0$, namely $S^*(Y_0) < min\{\widetilde{Y}, Y_0\}$, and A's expected discounted value of $X_{t+\Delta}$ is strictly less than X_t, then $(\tau^\Delta, d^\Delta, \alpha^\Delta_\tau)$ is dynamically individually rational for small Δ. The proof is in the Online Appendix.

Proposition A.1. *Suppose there exists an $r' < r$ such that $x \geq \mathbb{E}_x[e^{-r'\Delta} X_\Delta]\ \forall x$ as $\Delta \to 0$. If $S^*(Y_0) < min\{\widetilde{Y}, Y_0\}$, then $(\tau^\Delta, d^\Delta_\tau, \alpha^\Delta_\tau)$ is dynamically individually rational for sufficiently small Δ.*

The fact that $(\tau^*, d^*_\tau, \alpha^*_\tau)$ is optimal in continuous time follows from the convergence of X to Y, thereby yielding Theorem 1. Let $\mathbb{E}^\Delta$ be the expectation over X_t

[17] Because S^Δ and α^Δ are decreasing and uniformly bounded, by Helly's Selection Theorem the sequence as $\Delta \to 0$ has a subsequence converging to decreasing functions S^* and α^*.

with period length Δ and $\mathbb{E}^c$ be the expectation over Y_t. We provide the proof for a diffusion process below. The proof for a search process is similar, but requires more details on the construction of the continuous time search process, and so is deferred to the Online Appendix.

Proof. Let us extend X_t to continuous time by keeping X constant on $[0, \Delta), [\Delta, 2\Delta), \ldots$. That P's utility from $(\tau^\Delta, d_\tau^\Delta, \alpha_\tau^\Delta)$ converges to the $(\tau^*, d_\tau^*, \alpha_\tau^*)$ as $\Delta \to 0$ and his continuous time utility from $(\tau^*, d_\tau^*, \alpha_\tau^*)$-namely,

$$\begin{aligned}\lim_{\Delta\to 0} \mathbb{E}^\Delta[e^{-r\tau^\Delta}(d_\tau^\Delta(u_P(\alpha_\tau^\Delta) - \nu) + \nu)] &= \lim_{\Delta\to 0} \mathbb{E}^\Delta[e^{-r\tau^*}(d_\tau^*(u_P(\alpha_\tau^*) - \nu) + \nu)] \\ &= \mathbb{E}^c[e^{-r\tau^*}(d_\tau^*(u_P(\alpha_\tau^*) - \nu) + \nu)],\end{aligned}$$

is straightforward from the continuity of S^*, α^* and the well-known fact that there exists a probability space such that for any $T, \delta > 0$, for small enough Δ we have $\mathbb{P}(\max_{t\in[0,T]} |X_t - Y_t| > \delta) < \delta$. The same convergence holds for A's discrete and continuous time utility from $(\tau^*, d_\tau^*, \alpha_\tau^*)$. It is easily seen that A's continuous time continuation value under $(\tau^*, d_\tau^*, \alpha_\tau^*)$ is, for each (y, m) the same for every history with $(Y_t, M_t) = (y, m)$. Let $V^c(y, m)$ and $V^\Delta(x, m)$ be A's continuous and discrete time continuation values respectively. Then $(\tau^*, d_\tau^*, \alpha_\tau^*)$ is dynamically individually rational if $V^c(y, m) \geq y$ for all $m < \overline{R}^*$ and $y \in (S^*(m), m]$. Suppose $V^c(y, m) < y$ for some such (y, m). Let (x^Δ, m^Δ) be a sequence such that $(x^\Delta, m^\Delta) \uparrow (y, m)$ and $(x^\Delta, m^\Delta) \in \mathcal{G}^\Delta \times \mathcal{G}^\Delta$. Then $\lim_{\Delta\to 0} V^\Delta(x^\Delta, m^\Delta) = V^c(y, m) < y$, contradicting the fact that $V^\Delta(x^\Delta, m^\Delta) \geq x^\Delta$ by Lemma A.4. Therefore, $(\tau^*, d_\tau^*, \alpha_\tau^*)$ is dynamically individually rational.

All that is left is to show is that $(\tau^*, d_\tau^*, \alpha_\tau^*)$ is optimal in continuous time. Suppose there was a dynamically individually rational continuous time mechanism $(\tau^{**}, d_\tau^{**}, \alpha_\tau^{**})$ such that $J^{**} := J(\tau^{**}, d_\tau^{**}, \alpha_\tau^{**}) > J(\tau^*, d_\tau^*, \alpha_\tau^*)$.

We start by extending u_P and u_A to be defined on $\mathbb{R}$ rather than $[0, 1]$ in any way that satisfies $u_i'' \leq 0 < u_i'$ for $i = P, A$ with each u_i' uniformly bounded away from 0 and ∞. We also expand the range of α_τ to $(-\infty, 1]$. With this extended environment, all the steps in the derivation of the optimal discrete time mechanism go through verbatim and the optimal mechanism takes the same form as when we restricted $\alpha_\tau \geq 0$. [18] We will use this extended model for the rest of the proof.

[18]Note that in our optimal mechanism, the restriction to $\alpha_\tau \geq 0$ was not binding; even if P could choose $\alpha_\tau < 0$, the monotonicity of $\alpha^\Delta(m)$ and the fact that $V(m, m) = m$ implies that, whenever $\alpha^\Delta(m) < 0$, P would be better off taking his outside option at (m, m). Thus, the value

Take some small $\delta, \delta_\alpha, \delta_T > 0$ and T such that $e^{-rT} u_P(1) = \delta_T$. We use a single probability space so that convergence in probability of X_t to Y_t holds over $[0,T]$. Let $\Delta(\delta, \delta_T)$ be a Δ such that $\mathbb{P}(\max_{t\in[0,T]} |Y_t - X_t| > \delta) < \delta$. We allow the discrete time mechanism to depend on the realization of Y as well as X. Because Y is payoff irrelevant, allowing our discrete-time mechanism to be measurable with respect to Y_t changes neither the optimal mechanism $(\tau^\Delta, d^\Delta_\tau, \alpha^\Delta_\tau)$ nor P's value from $(\tau^\Delta, d^\Delta_\tau, \alpha^\Delta_\tau)$.

Let $\lceil t \rceil_\Delta$ be t rounded up the nearest date in $\{0, \Delta, ...\}$ and $\overline{\alpha} < 1 - u_A^{-1}(\overline{Y}) < 0$. We define a discrete-time mechanism $(\hat{\tau}, \hat{d}_\tau, \hat{\alpha}_\tau)$ that approximates $(\tau^{**}, d^{**}_\tau, \alpha^{**}_\tau)$ as long as $\tau^{**} < T$ and Y_t and X_t are close; otherwise it gives $1 - \overline{\alpha}$ to A. Let $\hat{\tau} = \lceil \tau^{**} \wedge T \wedge inf\{t : |Y_t - X_t| > \delta\} \rceil_\Delta$. If $\hat{\tau} < \lceil T \wedge inf\{t : |Y_t - X_t| > \delta\} \rceil_\Delta$ and $X_{\hat{\tau}} < u_A(1 - \alpha_{\hat{\tau}} + \delta_\alpha)$, then $\hat{d}_\tau = d_{\tau^{**}}$ and $\hat{\alpha}_\tau = \alpha_{\tau^{**}} - \delta_\alpha$; otherwise, $\hat{d}_\tau = 1,\ \hat{\alpha}_\tau = \overline{\alpha} - \delta_\alpha$.

Let $\delta, \delta_T \to 0$ (which implies $\Delta(\delta, \delta_T) \to 0$). Because $\delta_\alpha > 0$, A receives a discrete increase in utility from any split. By the convergence of X to Y, for small Δ, A's continuation value at any t in the discrete-time mechanism is strictly higher than than his continuous time continuation value from $(\tau^{**}, d^{**}_\tau, \alpha^{**}_\tau)$. Because $(\tau^{**}, d^{**}_\tau, \alpha^{**}_\tau)$ satisfies dynamic individual rationality, so too will $(\hat{\tau}, \hat{d}_\tau, \hat{\alpha}_\tau)$. Taking $\delta, \delta_T \to 0$, P's value from $(\hat{\tau}, \hat{d}_\tau, \hat{\alpha}_\tau)$ when $\Delta = \Delta(\delta, \delta_T)$ is bounded below by $J^{**} + u'_P(-\delta_\alpha)\delta_\alpha + O(\delta_\alpha^2)$. Because δ_α is arbitrary, we can take $\delta_\alpha \to 0$ to get a dynamically individually rational mechanism that gets arbitrarily close to J^{**} for small Δ. But this contradicts the fact P's utility from $(\tau^\Delta, d^\Delta_\tau, \alpha^\Delta_\tau)$ is the solution to P's relaxed problem and converges to $J(\tau^*, d^*_\tau, \alpha^*_\tau)$. Therefore, no such $(\tau^{**}, d^{**}, \alpha^{**}_\tau)$ exists. □

Appendix B Alternating Offers Bargaining

Define $\chi = u_A(1 - u_P^{-1}(\nu))$. There is an immediate split at t that is individually rational (IR) for both parties if and only if $X_t \leq \chi$. By constrained Pareto efficiency, we know $\overline{R}^* > \chi$. Because, at any (X_t, M_t), the continuation mechanism in our relaxed problem does not depend on X_0 and $S^*(M)$ is decreasing in M, there exists M^* such that the optimal mechanism has an immediate split if and only if $X_0 \leq M^*$. Note that $M^* < \chi$; otherwise $J(M^*, M^*) < \nu$, a contradiction

of the optimal mechanism with these extended utility functions is the same as when restricting $\alpha_\tau \in [0, 1]$.

of Lemma A.5. If $S^*(Y_0) > \widetilde{Y}$, then an immediate agreement is reached. We will treat such situations as if $S^*(Y_0) = Y_0$.

Equilibrium Definition

Let α_t and β_t be P's and A's respective demands at time t. Fix some $\kappa \in (0, \chi - M^*)$ and let $M^\kappa = M^* + \kappa$. Suppose $X_0 \geq M^\kappa$. We specify on-path strategies that implement the same outcomes as $(\tau^\Delta, d_\tau^\Delta, \alpha_\tau^\Delta)$. The proposer at τ^Δ is called to make a demand that gives P and A, respectively, $\alpha^\Delta(M_\tau)$ and $1 - \alpha^\Delta(M_\tau)$ if $X_{\tau^\Delta} < \overline{R}^\Delta$ and take his outside option if $X_{\tau^\Delta} \geq \overline{R}^\Delta$. Prior to τ^Δ, the proposer demands the entire surplus and the responder rejects. Whether a split is proposed at $S^\Delta(M_t) + \epsilon_{S^\Delta(M_t)}$ when $\gamma^\Delta(M_t) \in (0,1)$ is determined by the public randomization device.

To support these on-path strategies, we define several off-path punishment equilibria. In an *outside equilibrium*, the proposer makes the maximal possible demand subject to the responder's individual rationality (IR) constraint-namely, $\alpha_t = 1 - u_A^{-1}(X_t)$ and $\beta_t = 1 - u_P^{-1}(\nu)$. The responder accepts any IR demand and rejects otherwise. If the responder rejects the demand, the proposer immediately takes his outside option.

We now define an A-*inside equilibrium* beginning at $t + \Delta$. The A-inside equilibrium begins when P proposes at $t + \Delta$ and will depend on the value of X_t. If $X_t \geq M^\kappa$, P uses the same continuation strategies from $t + \Delta$ on as if we were on-path at time Δ when starting at $(X_0, M_0) = (X_t, X_t)$. By Lemma A.4, A's continuation value is X_t.

If $X_t < M^\kappa$, at $t + \Delta$ P demands $\hat{\alpha}_{t+\Delta}$, which is chosen to leave A indifferent between waiting until $t + \Delta$ or taking his outside option immediately at t. A accepts any demand with $\alpha_{t+\Delta} \leq \hat{\alpha}_{t+\Delta}$ and rejects otherwise. For a diffusion process, $\hat{\alpha}_{t+\Delta} = 1 - u_A^{-1}(X_{t+\Delta}) - \delta_A$, where δ_A is chosen so that

$$X_t = e^{-r\Delta}\mathbb{E}_{X_t}[u_A(u_A^{-1}(X_\Delta) + \delta_A))].$$

For a search process, $\hat{\alpha}_{t+\Delta} = 1 - u_A^{-1}(X_t - \epsilon_{X_t}) - \delta_A$ if there is $X_{t+\Delta} < X_t$, $\hat{\alpha}_{t+\Delta} = 1 - u_A^{-1}(X_{t+\Delta})$ if $X_{t+\Delta} \in [X_t, \chi)$. If $X_{t+\Delta} > \chi$, P makes no offer and takes his outside option. We choose δ_A so that

$$X_t = e^{-r\Delta}\big[\zeta^\Delta(X_t)\mathbb{E}_Z[max\{Z_t, X_t\}] + (1 - \zeta^\Delta(X_t))u_A(u_A^{-1}(X_t - \epsilon_{X_t}) - \delta_A)\big].$$

As $\Delta \to 0$, $\delta_A \to 0$ and $\hat{\alpha}_{t+\Delta} \to 1 - u_A^{-1}(X_t)$. P's continuation value at t from the A-inside equilibrium begining at $t + \Delta$ converges to $u_P(1 - u_A^{-1}(X_t)) > \nu$.

We now define a P-*inside equilibrium* beginning at $t+\Delta$. The P-inside equilibrium begins when A proposes at $t+\Delta$ and will also depend on the value of X_t. At $t+\Delta$, A makes a demand $\hat{\beta}_{t+\Delta}$. P accepts any demand with $\beta_{t+\Delta} \leq \hat{\beta}_{t+\Delta}$ and rejects otherwise. For a diffusion process, $\hat{\beta}_{t+\Delta} = 1 - u_P^{-1}(\nu) - \delta_P$; δ_P is set to leave P indifferent between waiting for A's demand and taking his outside option immediately:

$$\nu = e^{-r\Delta} u_P(u_P^{-1}(\nu) + \delta_P).$$

For a search process, $\hat{\beta}_{t+\Delta} = 1 - u_P^{-1}(\nu) - \delta_P$ if $X_{t+\Delta} < X_t$, $\hat{\beta}_{t+\Delta} = 1 - u_P^{-1}(\nu)$ if $X_{t+\Delta} \in [X_t, \chi)$. If $X_{t+\Delta} > \chi$, A makes no offer and takes his outside option. We choose δ_P so that

$$\nu = e^{-r\Delta}[\zeta^{\Delta}(X_t)\nu + (1 - \zeta^{\Delta}(X_t)) u_P(u_P^{-1}(\nu) + \delta_P)].$$

As $\Delta \to 0$, $\delta_P \to 0$ and A's continuation value at t converges to χ. For any $X_t < \chi$, A will find it strictly optimal to wait at t rather than take his outside option for all sufficiently small Δ.

When $X_0 \geq M^{\kappa}$, we specify the following reaction to a deviation. If i is the proposer and deviates at t by making a larger demand when k is called to accept the equilibrium demand, then k rejects i's demand and i moves to the next period. Continuation play is given by the i-inside equilibrium beginning at $t+\Delta$. If k either rejects i's demand when k is called to accept or fails to take his outside option when called to, then i immediately takes his outside option. If i does not take his outside option, we move to the outside equilibrium for all subsequent periods. We note that, for small Δ, the only equilibrium demands for which IR is not binding and are called to be accepted are made at $X_t \leq M^{\kappa}$.

If $X_0 < M^{\kappa}$ and P is the first proposer, then we start in the outside equilibrium. If A is the first proposer, then we fix a small $\delta > 0$ and specify that A demands $u_A^{-1}(X_0 + \delta)$ and P accepts immediately. If A raises his demand, P rejects and we move to an A-inside equilibrium.

Proof of Theorem 2

Proof. For $X_0 \geq M^{\kappa}$, by Theorem 1, the on-path strategies converge to the optimal continuous time mechanism. Because κ was arbitrary, this holds for any $X_0 > M^*$. For $X_0 \leq M^*$, convergence to the continuous time mechanism is clear as $\kappa, \delta \to 0$. We need only verify that the strategies are a sub-game perfect equilibrium.

We first argue that the outside equilibrium is sub-game perfect. The proposer has no incentive to change his demand because he is getting the maximal amount he can subject to the responder's IR constraint. Taking his outside option following a rejection is optimal for the proposer because he has a continuation value in the next period equal to his outside option. His opponent finds it optimal to accept the proposer's demand because it is IR and rejection will lead to the proposing player taking his outside option immediately.

We now argue that our on-path and inside equilibria strategies are sub-game perfect. In on-path play, due to the constrained Pareto-efficiency of continuation play (Proposition 2), there is no deviating demand before τ^{Δ} that would be profitable for the proposer and which the responder would accept. Continuation play is always dynamically individually rational for small Δ,[19]so neither player has an incentive to take their outside option earlier than called to. At $t + \Delta$, the splits agreed upon are IR for both players. Thus, the responder has no incentive to reject a demand he is called to accept because a rejection leads to the proposer immediately taking his outside option. The proposer has no incentive to raise his demands as they will be rejected and his continuation value will be equal to his outside option.

All that is left to show is that player k, when called to reject a deviating demand by player i, finds it optimal to do so. Suppose P were to increase his demand in a period in which A is called to accept P's equilibrium demand and P's increased demand was IR for A. For small Δ, P only makes an equilibrium IR demand for A if $X_t \leq M^{\kappa}$. If A rejects, A's continuation value in the subsequent P-inside equilibrium is approximately $u_A(1 - u_P^{-1}(\nu)) > M^{\kappa}$. Fixing any demand amount $\alpha_t > u_P^{-1}(\nu)$, for all Δ small enough, A will find it optimal to reject P's deviating demand. When $X_0 \geq M^{\kappa}$ or we are in a A-inside equilibrium that began at $X_t \geq M^{\kappa}$, P's equilibrium demand when A is called to accept is bounded away from $u_P^{-1}(\nu)$; otherwise, due to discounting costs, P would be better off taking his outside option at X_0 or the beginning of the A-inside equilibrium, contradicting Lemma A.5. Thus, an increased demand by P will be rejected by A.

Now suppose A increases his demand in a period in which P is called to accept A's equilibrium demand. P's continuation value in the A-inside equilibrium is equal to $J(X_t, X_t)$ if $X_t \geq M^{\kappa}$ or, if $X_t < M^{\kappa}$, approximately $u_P(1 - u_A^{-1}(X_t))$ as $\Delta \to 0$. In a P-inside equilibrium, A's equilibrium demands leave P with approximately a

[19]For A, this follows from Proposition A.1; for P it follows from Lemma A.5.

utility of ν if $X_t < \chi$ and exactly equal to ν if $X_t \geq \chi$. Because $J(X_t, X_t) > \nu$ for $X_t < \chi$ by Lemma A.5, in a P-inside equilibrium, any increase in A's demand will be rejected by P. In the on-path equilibrium, A is making a demand that is called to be accepted only at τ^Δ and $X_{\tau^\Delta} < M^\kappa$. At τ^Δ, A is making a demand which, for small Δ, gives A utility which is bounded away from his outside option: $\beta_t = 1 - \alpha^\Delta(M_t) > u_A^{-1}(X_t)$, which gives P a utility strictly less than (and bounded away from as $\Delta \to 0$) $u_P(1 - u_A^{-1}(X_t))$. P's utility from the A-inside equilibrium if he rejects is approximately $u_P(1 - u_A^{-1}(X_t))$. Thus, an increased demand by A will be rejected by P. $\square$

Appendix O.A Online Appendix

Proof of Lemma A.1

Proof. That $\alpha_\tau = 1$ is immediate. To find the optimal (τ, d_τ), we partition $\mathcal{G}^\Delta$ into a continuation set $\mathcal{C}$ and a stopping set $\mathcal{D}$ and partition $\mathcal{D}$ into $\mathcal{D}_S$, where it is optimal to take the split, and $\mathcal{D}_O$, where it is optimal to take the outside option:

$$\begin{aligned}
\mathcal{C} &= \{x: \ V^*(x) > max\{u_A(1), x\}\}, \\
\mathcal{D}_S &= \{x: \ V^*(x) = u_A(1) \geq x\}, \\
\mathcal{D}_O &= \{x: \ V^*(x) = x > u_A(1)\}.
\end{aligned}$$

It is clear that $min\{x \in \mathcal{D}_O\} > max\{x \in \mathcal{D}_S\}$. As is standard, the optimal stopping rule is then given by $\tau = min\{t: \ X_t \in \mathcal{D}\}$ with $d_\tau = \mathbb{1}(X_\tau \in \mathcal{D}_S)$.

We want to show $\mathcal{C} = (S^A, R^A) \cap \mathcal{G}^\Delta$ for $S^A = max\{x: \ x \in \mathcal{D}_S\}$ and $R^A = min\{x: \ x \in \mathcal{D}_O\}$. Let $x_1, x_2 \in \mathcal{D}$ such that $x_2 > x_1 + \epsilon_{x_1}$ and $x \in \mathcal{C} \ \forall x \in (x_1, x_2)$. Take $x' \in (x_1, x_2)$. We argue that it must be that $x_1 \in \mathcal{D}_S$ and $x_2 \in \mathcal{D}_O$. This suffices to show $\mathcal{C} = (S^A, R^A) \cap \mathcal{G}^\Delta$ as otherwise there would exist such $x_1, x_2 \in \mathcal{D}_S$ or $x_1, x_2 \in \mathcal{D}_O$.

We first consider a diffusion process. Suppose $x_1, x_2 \in \mathcal{D}_S$. Then

$$\begin{aligned}
V^*(x') &= \mathbb{E}_{x'}[e^{-r(\tau_-(x_1) \wedge \tau_+(x_2))} V^*(X_{\tau_-(x_1) \wedge \tau_+(x_2)})] \\
&= \mathbb{E}_{x'}[e^{-r(\tau_-(x_1) \wedge \tau_+(x_2))} u_A(1)] \\
&< u_A(1),
\end{aligned}$$

contradicting $x' \in \mathcal{C}$. Next, suppose $x_1, x_2 \in \mathcal{D}_O$. Then

$$\begin{aligned}
V^*(x') &= \mathbb{E}_{x'}[e^{-r(\tau_-(x_1) \wedge \tau_+(x_2))} V^*(X_{\tau_-(x_1) \wedge \tau_+(x_2)})] \\
&= \mathbb{E}_{x'}[e^{-r(\tau_-(x_1) \wedge \tau_+(x_2))} X_{\tau_-(x_1) \wedge \tau_+(x_2)}] \\
&\leq x',
\end{aligned}$$

where the inequality follows from Doob's Optional Stopping Theorem, contradicting $x' \in \mathcal{C}$.

We next consider a search process. It is easily seen that $V^*(x)$ is increasing in x. Suppose $x_1, x_2 \in \mathcal{D}_S$. The fact that continuing for one additional period is optimal at $x_1 + \epsilon_{x_1}$ and it is optimal to stop at x_1 implies

$$\begin{aligned}
V^*(x_1 + \epsilon_{x_1}) &= e^{-r\Delta}[\zeta^\Delta(x_1 + \epsilon_{x_1})\mathbb{E}_Z[V^*(max\{Z, x_1 + \epsilon_{x_1}\}))] + (1 - \zeta^\Delta(x_1 + \epsilon_{x_1}))u_A(1)] \\
&> u_A(1). \qquad (11)
\end{aligned}$$

For this to be true, we must have $\mathbb{E}_Z[V^*(max\{Z, x_1 + \epsilon_{x_1}\})] > u_A(1)$. The fact that it is not optimal to continue for one more period at x_2 implies

$$e^{-r\Delta}[\zeta^\Delta(x_2)\mathbb{E}_Z[V^*(max\{Z, x_2\})] + (1 - \zeta^\Delta(x_2))u_A(1)] < u_A(1) = V^*(x_2).$$

a contradiction of 11 because $\zeta^\Delta(x)$ and $\mathbb{E}_Z[V^*(max\{Z, x\})]$ are increasing in x.

Now suppose $x_1, x_2 \in \mathcal{D}_O$. When $X_0 = x'$, X_t must travel through x_1 to go below x_1. Therefore $min\{t : X_t \in \mathcal{D}\} = min\{t : X_t \in \mathcal{D}_O\}$ and

$$V^*(x') = \sup_\tau \ \mathbb{E}_{x'}[e^{-r\tau} max\{u_A(1), X_\tau\}] = \sup_\tau \mathbb{E}_{x'}[e^{-r\tau}X_\tau] \leq x',$$

where the inequality follows from Doob's Optional Stopping Theorem, contradicting $x' \in \mathcal{C}$. □

Proof of Lemma A.2

Proof. Our goal is to write the relaxed problem 2 as a constrained Markov decision problem (as in Altman (1999)) with state space is $\{0,1\} \times \mathcal{G}^\Delta \times \mathcal{G}^\Delta$ and generic state (H_t, X_t, M_t) at time t. X_t and M_t are taken from the discretized outside option process while H_t is a indicator variable for whether P has not stopped prior to t. An action in period t is $(a_t, d_t, \alpha_t) \in \{0,1\} \times \{0,1\} \times [0,1]$ where a_t is the action to stop at time t, d_t is an indicator for a split being made when stopping at t and α_t is the share of the surplus going to P when implementing a split t. A history at $t > 0$ takes the form $h_t = (H_0, X_0, M_0, a_0, d_0, \alpha_0, H_1, ..., \alpha_{t-1}, H_t, X_t, M_t)$ with $h_0 = (H_0, X_0, M_0)$. A strategy $(\tau, d_\tau, \alpha_\tau)$ maps histories h_t into a probability mixture over (a_t, d_t, α_t).[20] By our restriction after Lemma A.1, we set $a_t = 1, d_t = 0$ whenever $X_t \geq R^A$. The transition probability for H is $H_{t+\Delta} = H_t(1 - a_t)$ and $H_0 = 1$.

We now show how to write the objective function and constraint set with state space (H, X, M) and in the form used in Altman (1999). The objective function, P's expected utility, is equal to

$$\mathbb{E}[\sum_{t\in\{0,\Delta,\ldots\}} e^{-rt}H_t a_t\big(d_t(u_P(\alpha_t) - \nu) + \nu\big)].$$

[20] Altman (1999) shows it is without loss to focus on behavioral policies which do not condition on past randomization.

We can rewrite $RDIR(X^n)$ as

$$V(\tau \wedge \tau_+(X^n), d_\tau(X^n), \alpha_\tau) - V(\tau, d_\tau, \alpha_\tau) \leq 0. \tag{12}$$

Similar to the objective function, $V(\tau, d_\tau, \alpha_\tau)$ is equal to

$$\mathbb{E}[\sum_{t \in \{0,\Delta,\ldots\}} e^{-rt} H_t a_t \big(d_t(u_A(1-\alpha_t) - X_t) + X_t)\big)].$$

If $X^n = X^0$, then $V(\tau \wedge \tau_+(X^n), d_\tau(X^n), \alpha_\tau) = X^0$ and we are done. When $X^n > X^0$, $V(\tau \wedge \tau_+(X^n), d_\tau(X^n), \alpha_\tau)$ is equal to

$$\begin{aligned}&\mathbb{E}[\sum_{t \in \{0,\Delta,\ldots\}} e^{-rt} H_t a_t (d_t(u_A(1-\alpha_t) - X_t) + X_t) \mathbb{1}(M_t < X^n)] \\ &+ \mathbb{E}[\sum_{t \in \{\Delta,\ldots\}} e^{-rt} H_t \mathbb{1}(M_{t-\Delta} < X^n) \mathbb{1}(X^n \leq X_t) X_t].\end{aligned}$$

The first line is the payoff of the mechanism stopping prior to $\tau_+(X^n)$ and the second line is the payoff from taking the outside option at $\tau_+(X^n)$.[21] Let $\mathbb{P}_{x,x'} = \mathbb{P}(X_{t+\Delta} = x | X_t = x')$ The second line is equal to

$$\mathbb{E}\big[\sum_{t \in \{\Delta,\ldots\}} e^{-rt} H_t \mathbb{1}(M_{t-\Delta} < X^n) \sum_{x' \geq X^n} \mathbb{P}_{X_{t-\Delta},x'} x'\big]. \tag{13}$$

Because $H_t = H_{t-\Delta}(1 - a_{t-\Delta})$, 13 is equal to

$$\mathbb{E}\big[\sum_{t \in \{0,\Delta,\ldots\}} e^{-rt} H_t (1 - a_t) \mathbb{1}(M_t < X^n) e^{-r\Delta} \sum_{x' \geq X^n} \mathbb{P}_{X_t,x'} x'\big].$$

We can clearly see that $RDIR(X^n)$ depends only on P's strategy and (H, X, M).

By Theorem 8.4 of Altman (1999), there exists a optimal solution to our relaxed problem which is stationary in (H, X, M);[22] that is, conditional on $H_t = 1$, in each period the probability distribution over (a_t, d_t, α_t) is a function only of (X_t, M_t). □

[21] We drop $t = 0$ from the summation because the probability of going above X^n at $t = 0$ is zero.

[22] The application of Theorem 8.4 follows the discussion on pg. 138 of Altman (1999) which, using Altman's terminology, assumes a non-negative immediate cost. We can rewrite $RDIR(X^n)$ to satisfy the non-negative immediate costs by adding $\sum_{t=\{0,\Delta,\ldots\}} e^{-rt}\overline{X}$ to both sides of 12.

Dynamic Individual Rationality

We prove below that DIR is weaker than dynamic individual rationality.

Lemma O.A.1. *Any mechanism that satisfies dynamic individual rationality satisfies DIR.*

Proof. Suppose $(\tau, d_\tau, \alpha_\tau)$ is dynamically individually rational. For any τ' and corresponding history $h_{\tau'}$, we have $V(\tau, d_\tau, \alpha_\tau) - V(\tau \wedge \tau', d_\tau \mathbb{1}(\tau < \tau'), \alpha_\tau)$ is equal to

$$\mathbb{E}\big[e^{-r\tau'}\mathbb{E}[e^{-r(\tau-\tau')}(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau)|h_{\tau'}]\mathbb{1}(\tau' \leq \tau)\big],$$

which is positive by the fact that A's continuation value at $h_{\tau'}$, $\mathbb{E}[e^{-r(\tau-\tau')}(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau)|h_{\tau'}]$, is positive by dynamic individual rationality. Because this holds for all τ', DIR is satisfied. □

Proof of Proposition A.1

We start with a simple Lemma pointing out that the solution to our discrete time relaxed is "almost" dynamically individually rational for any Δ. Lemma A.4 shows A's continuation value is above his outside option at the *beginning* of a period after every history. This is close to, but not quite, dynamic individual rationality. We still need to check that A prefers to accept a split when offered one. The next Lemma shows this will be true except, potentially, when reaching a split at $S^\Delta(X_0) + \epsilon_{S^\Delta(X_0)}$.

Lemma O.A.2. $u_A(1-\alpha^\Delta(M_\tau)) > X_\tau$ *except, potentially, when* $X_\tau = S^\Delta(X_0) + \epsilon_{S^\Delta(X_0)}$.

Proof. If $u_A(1-\alpha^\Delta(X_0)) \geq S^\Delta(X_0)+\epsilon_{S^\Delta(X_0)}$, then, because $\alpha^\Delta(m)$ and $S^\Delta(m)$ are decreasing, $u_A(1-\alpha^\Delta(M_\tau)) \geq S^\Delta(M_\tau)+\epsilon_{S^\Delta(M_\tau)} \geq X_\tau$. Suppose $u_A(1-\alpha^\Delta(X_0)) < S^\Delta(X_0) + \epsilon_{S^\Delta(X_0)}$. If $u_A(1-\alpha^\Delta(X_0)) < S^\Delta(X_0)$, then, because A's continuation value at $\tau_+(X^1)$ is $X_{\tau_+(X^1)}$

$$\begin{aligned} V(X_0, X_0) &= \mathbb{E}[e^{-r(\tau\wedge\tau_+(X^1))}\big(X_{\tau_+(X^1)}\mathbb{1}(\tau^\Delta > \tau_+(X^1)) + u_A(1-\alpha^\Delta(X_0))\mathbb{1}(\tau^\Delta < \tau_+(X^1))\big)] \\ &< \mathbb{E}[e^{-r(\tau\wedge\tau_+(X^1))}\big(X_{\tau_+(X^1)}\mathbb{1}(\tau^\Delta > \tau_+(X^1)) + X_{\tau^\Delta}\mathbb{1}(\tau^\Delta < \tau_+(X^1))\big)] \\ &\leq X_0, \end{aligned}$$

where the final inequality follows from Doob's Optional Stopping Theorem. Therefore we must have $u_A(1-\alpha^{\Delta}(X_0)) \geq S^{\Delta}(X_0)$. Because $\alpha^{\Delta}(m)$ and $S^{\Delta}(m)$ are decreasing, $u_A(1-\alpha^{\Delta}(M_\tau)) \geq S^{\Delta}(M_\tau) + \epsilon_{S^{\Delta}(M_\tau)} > S^{\Delta}(M_\tau)$ for all M_τ such that $S^{\Delta}(M_\tau) < S^{\Delta}(X_0)$. $\square$

We can now prove Proposition A.1.

Proof. By Lemma A.4, $V(x,m) \geq x$, so to establish dynamic individual rationality we only need to verify that, whenever offered a split, A prefers to accept P's proposed split rather than take his outside option. By Lemma O.A.2, it suffices to show that $u_A(1-\alpha^{\Delta}(X_0)) \geq S^{\Delta}(X_0) + \epsilon_{S^{\Delta}(X_0)}$.

Suppose, for arbitrarily small Δ, $u_A(1-\alpha^{\Delta}(X_0)) < S^{\Delta}(X_0) + \epsilon_{S^{\Delta}(X_0)}$. Let $x = argmin_{x' \in \mathcal{G}^{\Delta}} |x' - \frac{S^*(X_0)+min\{\widetilde{Y}, X_0\}}{2}|$. Let $\tau' = \tau_-^{(x,X_0)}(S^{\Delta}(X_0)) \wedge \tau_+^{(x,X_0)}(X^1)$. Because A strictly prefers to continue than take the split $S^{\Delta}(X_0) + \epsilon_{S^{\Delta}(X_0)}$, we can (weakly) increase A's utility by setting $\gamma^{\Delta}(X_0)$ to 0. This lets us bound A's continuation value at (x, X_0):

$$
\begin{aligned}
V(x,X_0) \leq\; & \mathbb{E}[e^{-r\tau_-^{(x,X_0)}(S^{\Delta}(X_0))} u_A(1-\alpha^{\Delta}(X_0)) \mathbb{1}(\tau_-^{(x,X_0)}(S^{\Delta}(X_0)) < \tau_+^{(x,X_0)}(X^1)) \\
& + e^{-r\tau_+^{(x,X_0)}(X^1)} V(X_{\tau_+(X^1)}, X_{\tau_+(X^1)}) \mathbb{1}(\tau_-^{(x,X_0)}(S^{\Delta}(X_0)) > \tau_+^{(x,X_0)}(X^1)) | (x,X_0)] \\
& < \mathbb{E}[e^{-r\tau_-^{(x,X_0)}(S^{\Delta}(X_0))} (S^{\Delta}(X_0) + \epsilon_{S^{\Delta}(X_0)}) \mathbb{1}(\tau_-^{(x,X_0)}(S^{\Delta}(X_0)) < \tau_+^{(x,X_0)}(X^1)) \\
& + e^{-r\tau_+^{(x,X_0)}(X^1)} X_{\tau_+(X^1)} \mathbb{1}(\tau_-^{(x,X_0)}(S^{\Delta}(X_0)) > \tau'_+(X^1)) | (x,X_0)] \\
& = \mathbb{E}[e^{-r\tau'} X_{\tau'} | (x,X_0)] + \epsilon_{S^{\Delta}(X_0)} \mathbb{E}[e^{-r\tau'} \mathbb{1}(\tau'_-(S^{\Delta}(X_0)) < \tau_+^{(x,X_0)}(X^1)) | (x,X_0)], \\
& \leq \mathbb{E}[e^{-r\tau'} X_{\tau'} | (x,X_0)] + \epsilon_{S^{\Delta}(X_0)}.
\end{aligned} \tag{14}
$$

By our choice of x, the expected length of time until τ' is strictly bounded away from 0 as $\Delta \to 0$. Take $r' < r$ such that $x' \geq e^{-r'\Delta} \mathbb{E}_{x'}[X_\Delta]$ for all $x' \in [S^{\Delta}(X_0), X^1]$. Then $e^{-r't} X_t$ is a supermartingale on $[S^{\Delta}(X_0), X^1]$ and, by Doob's Optional Stopping Theorem, $\mathbb{E}[e^{-r'\tau'} X_{\tau'} | (x, X_0)] \leq x$. Because this holds for all Δ, there is a $k < 1$ such that

$$
\lim_{\Delta \to 0} \mathbb{E}[e^{-r\tau'} X_{\tau'} | (x,X_0)] < \lim_{\Delta \to 0} k \mathbb{E}[e^{-r'\tau'} X_{\tau'} | (x,X_0)] \leq kx.
$$

Substituting this into 14, because $\epsilon_{S^{\Delta}(X_0)} \to 0$, we get $V(x, X_0) < x$ for sufficiently small Δ, a contradiction of Lemma A.4. Therefore, $u_A(1-\alpha^{\Delta}(X_0)) \geq S^{\Delta}(X_0) + \epsilon_{S^{\Delta}(X_0)}$. $\square$

Proof of Lemma A.8

Proof. Take an arbitrary (S, γ) and let $\tau^{S,\gamma}$ be the stopping rule induced by the (S, γ)-split threshold. For arbitrary x, we consider a problem in which we start at $X_0 = x$ and allow A to choose a stopping time $\tau' \leq \tau^{S,\gamma}$ at which point he can choose to take his outside option or take a split giving him $1 - \alpha$. A's value function in this problem is

$$\widetilde{V}^*(x, S, \gamma, \alpha) = \max_{\tau'} \widetilde{V}(x, S, \gamma, \alpha, \tau') \tag{15}$$
$$\text{where } \widetilde{V}(x, S, \gamma, \alpha, \tau') := \mathbb{E}_x[e^{-r(\tau' \wedge \tau^{S,\gamma})} max\{X_{\tau' \wedge \tau^{S,\gamma}}, u_A(1-\alpha)\}].$$

Let $\mathcal{C}(S, \gamma, \alpha)$ be the region for which it is weakly optimal to continue:

$$\mathcal{C}(S, \gamma, \alpha) = \{x : \; \exists \tau' \; s.t. \; \mathbb{P}(\tau' > 0), \; \widetilde{V}(x, S, \gamma, \alpha, \tau') = \widetilde{V}^*(x, S, \gamma, \alpha)\}.$$

Let $B(S, \gamma, \alpha) = max\{x : \; x \in \mathcal{C}(S, \gamma, \alpha)\}$ and $b(S, \gamma, \alpha) = min\{x : \; x \in \mathcal{C}(S, \gamma, \alpha)\}$. By similar arguments as in Lemma A.1, A's optimal stopping time takes the form $\tau' = \min\{t : \; X_t \notin (b(S, \gamma, \alpha), B(S, \gamma, \alpha))\}$ with $b(S, \gamma, \alpha) \geq S$.[23] A only takes his outside option when stopping at $x \geq B(S, \gamma, \alpha)$.

For any (S', γ', α') and $(S'', \gamma'', \alpha'')$ with $\alpha' \geq \alpha''$ and either $S' > S''$ or $S' = S''$ and $\gamma' \geq \gamma''$, it is easily seen that $\widetilde{V}^*(x, S', \gamma', \alpha') \leq \widetilde{V}^*(x, S'', \gamma'', \alpha'')$, which implies $B(S', \gamma', \alpha') \leq B(S'', \gamma'', \alpha'')$.

Let $B^*(m) = B(S^\Delta(m), \gamma^\Delta(m), \alpha^\Delta(m))$ and $b^*(m) = b(S^\Delta(m), \gamma^\Delta(m), \alpha^\Delta(m))$. Because $V(m, m) = m$, we have $\widetilde{V}^*(x, S^\Delta(m), \gamma^\Delta(m), \alpha^\Delta(m)) \geq V(x, m)$ for $x < m$ with equality if $b^*(m) = S^\Delta(m)$ and $B^*(m) = m$. If $B^*(m) > m$, then, for $m' \in [m, B^*(m))$, A would be willing to continue bargaining even if P didn't decrease S^Δ or α^Δ. P could increase his expected utility from $(\tau^\Delta, d^\Delta_\tau, \alpha^\Delta_\tau)$ by increasing $S^\Delta(m')$ to $b^*(m)$ and $\alpha^\Delta(m')$ to $\alpha^\Delta(m)$ while setting $\gamma^\Delta(m') = \gamma^\Delta(m)\mathbb{1}(b^*(m) = S^\Delta(m'))$ without violating $RDIR(m')$. If $B^*(m) < m$, then $\widetilde{V}^*(x, S^\Delta(m), \gamma^\Delta(m), \alpha^\Delta(m)) = x$ for $x > B^*(m)$ and $x \notin \mathcal{C}(S, \gamma, \alpha)$. Because $V(x, m) \geq x$, A would find it weakly optimal to not stop immediately in 15 by taking $B^* = m$ and using $b^* = S^\Delta(m)$, a contradiction of $x \notin \mathcal{C}(S, \gamma, \alpha)$. We must therefore have $B^*(m) = m$. To see that $b^*(m) = S^\Delta(m)$, remember that P would better off if A takes a split sooner. If $b^*(m) > S^\Delta(m)$, P could increase the split threshold from $S^\Delta(m)$ to $b^*(m)$ and make both players strictly better off.

[23]This is without loss because setting $b < S$ is equivalent to $b = S$.

As noted in Lemma A.7, A's utility from increasing the split threshold at $S^{\Delta}(m)$ is strictly negative. A fixed decrease in the split threshold from S to $S' < S$ will, for small Δ, lead to a discrete increase in $\widetilde{V}^*$ and B. A similar conclusion holds for a fixed discrete decrease in α. Suppose there is a discontinuity in S^* or α^* at m_d. Because B is increasing in S and α, there is a $\delta > 0$ such that, for small enough Δ, we have $B^*(m_d + \epsilon_{m_d}) > B^*(m_d) + \delta = m_d + \delta$, a contradiction of $B^*(m_d + \epsilon_{m_d}) = m_d + \epsilon_{m_d}$ when Δ is small enough that $\delta > \epsilon_{m_d}$. Therefore, no such discontinuity exists. $\square$

Search Process Construction

We provide a construction of a search process Y and its discrete time approximation X and prove the convergence of $(\tau^*, d^*_\tau, \alpha^*_\tau)$ under X to $(\tau^*, d^*_\tau, \alpha^*_\tau)$ under Y as $\Delta \to 0$. The proof, while intuitive, is somewhat long and tedious and can be skipped by the reader without much loss.

Let Z^∞ be a sequence $(Z^1, Z^2, ...)$ of iid draws of $Z^k \sim F$ and U^∞ be a sequence $(U^1, U^2, ...)$ of iid draws $U^k \sim U[0,1]$. Let $G(t,w) = exp(-\int_0^t \zeta(W(s,w))ds)$, where $W(s,w)$ is the solution at time s of the differential equation $dW_t = \eta(W_t)dt$ with initial condition $W_0 = min\{\widetilde{Y}, w\}$.

Let $W^0 = min\{\widetilde{Y}, Y_0\}$ and $T_1(W^0) = inf\{t : \ G(t, W^0) \leq U^1\}$ be the arrival time of the first Z draw, namely Z^1. Let $\underline{W}_1 = W(T_1(W^0), W^0)$ be the value of W at $T_1(W^0)$. For $k \geq 1$, let $\underline{W}_k = W(T_{k-1}(W^{k-1}), W^{k-1})$, $W^k = max\{min\{Z^k, \widetilde{Y}\}, \underline{W}_k\}$ and $T_{k+1}(W^k) = inf\{t : \ G(t, W^k) \leq U^{k+1}\}$ be the length of time between the arrival of Z^k and Z^{k+1}. Then Z^k arrives at $\overline{T}^k = \sum_{j=1}^k T_j(W^{j-1})$ with $\overline{T}^0 = 0$. We define Y_t for $t > 0$ by

$$Y_t = \sum_{k=0}^{\infty} W(t - \overline{T}^k, W^k)\mathbb{1}(\overline{T}^k < t < \overline{T}^{k+1}) + \mathbb{1}(t = \overline{T}^k)max\{Z^k, W^k\}.$$

We construct X in a similar way. If $w > \widetilde{Y}$, let $n^F(w) = |\{x \in \mathcal{G}^\Delta : \ x \in (\widetilde{Y}, w]\}|$ be the number of periods X remains above $\widetilde{Y}$ when starting at $X_0 = w$ if no new Z arrives before X crosses $\widetilde{Y}$; if $w \leq \widetilde{Y}$, let $n^F(w) = 0$. Note that $n^F(w) \leq \overline{F} := |supp(F)| + 1$.[24]. Let $\xi(\ell, w)$ be the value of x which is ℓ grid points below w, with $\xi(0,w) = w$ and define $\bar{\xi}(w) = \sum_{x \in \mathcal{G}^\Delta \cap (\widetilde{Y}, w]} \zeta(x)$. We define the

[24]The addition to $|supp(F)|$ accounts for the possible addition of Y_0 to $\mathcal{G}^\Delta$.

discrete time analogue of G by

$$G^{\Delta}(s,w)=\begin{cases}1 & \text{if } s\in[0,\Delta)\\ exp\big(-\sum_{x\in[\xi(\lfloor\frac{s-\Delta}{\Delta}\rfloor),w),w]}\zeta(x)\Delta\big) & \text{if } s\in[\Delta,n^F(w)\Delta),\\ exp\big(-\{\bar{\xi}(w)\Delta+\int_{n^F(w)}^{s}\zeta(W(t-n^F(w)\Delta,w))dt\}\big) & \text{if } s\geq n^F(w)\Delta.\end{cases}$$

where $W(t,w)$ is defined as before. The key difference is that we need to consider the movement downwards of X across points in $\mathcal{G}^{\Delta}$ above $\widetilde{Y}$; this enters G^{Δ} in the sum over $x\in(\widetilde{Y},w]$. Because $\zeta(x)$ is bounded and $\overline{F}$ is finite, it is clear that $G^{\Delta}(s,w)\to G(s,w)$ as $\Delta\to 0$.

We define a discrete time version of W at t when starting at w as

$$W^{\Delta}(t,w)=\mathbb{1}(t<n^F(w)\Delta)\xi(\lfloor\frac{t}{\Delta}\rfloor,w)+\mathbb{1}(t\geq n^F(w)\Delta)W(t-n^F(w)\Delta,w).$$

We define $(Z^{1,\Delta},Z^{2,\Delta},...)$ by rounding $(Z^1,Z^2,...)$ up to the nearest point in $\mathcal{G}^{\Delta}$: $Z^{k,\Delta}=min\{x\in\mathcal{G}^{\Delta}:\ x>Z^k\}$. Let $\mathbb{T}=\{0,\Delta,...\}$. The initial value of X is $W^{0,\Delta}=X_0$ and the length of time until the arrival of $Z^{1,\Delta}$ is $T_1^{\Delta}(W^{0,\Delta})=min\{t\in\mathbb{T}:\ G^{\Delta}(t,W^{0,\Delta})\leq U^0\}$. At $T_1^{\Delta}(W^{0,\Delta})$, the value of W^{Δ} is $\underline{W}_0^{\Delta}=\underline{W}(T_1^{\Delta}(W^{0,\Delta}),W^{0,\Delta})$. For $k\geq 1$, let $\underline{W}_k^{\Delta}=W^{\Delta}(T_{k-1}^{\Delta}(W^{k-1,\Delta}),W^{k-1,\Delta})$, $W^{k,\Delta}=max\{Z^{k,\Delta},\underline{W}_{k-1}^{\Delta}\}$ and $T_{k+1}^{\Delta}(W^{k,\Delta})=min\{t\in\mathbb{T}:\ G^{\Delta}(t,W^{k,\Delta})\leq U^{k+1}\}$. The time of the arrival of $Z^{k,\Delta}$ is then $\overline{T}^{k,\Delta}=\sum_{j=1}^{k}T_j^{\Delta}(W^{j-1,\Delta})$ and $\overline{T}^{0,\Delta}=0$. We define X_t for $t\in\mathbb{T}$ by

$$X_t=\sum_{k=0}^{\infty}W^{\Delta}(t-\overline{T}^{k,\Delta},W^{k,\Delta})\mathbb{1}(\overline{T}^{k,\Delta}<t<\overline{T}^{k+1,\Delta})+\mathbb{1}(t=\overline{T}^{k,\Delta})max\{Z^{k,\Delta},W^{k,\Delta}\}.$$

Let $\overline{\eta}=\max\limits_{y\in[\underline{Y},\overline{Y}]}|\eta(y)|$ and $\underline{\eta}=\min\limits_{y\in[S^*(\overline{R}^*),\widetilde{Y}]}|\eta(y)|$. By our definition of $\mathcal{G}^{\Delta}$, $\overline{\eta}\Delta\geq\max\limits_{x<\widetilde{Y}}\epsilon_x$. In the definition of $\underline{\eta}$ we restrict to $y\geq S^*(\overline{R}^*)$ because we will be concerned with bounds on the speed of the decrease in W when it is above S^*, for which $S^*(\overline{R}^*)$ is a lower bound. We note that $S^*(\overline{R}^*)>\underline{Y}$; otherwise, for small Δ, it would take arbitrarily long to reach $S^{\Delta}(M_t)$ when $M_t\approx\overline{R}^{\Delta}$, which would violate $RDIR$. Let Θ be the length of time it takes W^{Δ}, when starting at $\overline{Y}$, to reach $S^*(\overline{R}^*)$, which is defined by $W^{\Delta}(\Theta,\overline{Y})=S^*(\overline{R}^*)$. This is an upper bound on the maximum amount of time that can occur between arrivals of Z before X or Y stop under $(\tau^*,d_\tau^*,\alpha_\tau^*)$.

Before proving our convergence result, we start with a technical Lemma.

Lemma O.A.3. *$W(t,w)$ and $T_{k+1}(w)$ are Lipschitz continuous for any $k \geq 0$ as long as $W(t,w) \in [\overline{Y}, S^*(\overline{R}^*)]$ and $T_{k+1}(w) \in [0,\Theta]$ with Lipschitz constants K_W, K_T.*

Proof. We start with W. Because we have not reached S^*, we know $|\eta(W(t,w))| \geq \underline{\eta}$. Take $w' < w$. Then $W(t,w') < W(t,w)$. Note that $W(t,w') \geq w' - \overline{\eta}t$ and $W(t,w) \leq w - \underline{\eta}t$. If $W(t,w) \geq w'$, then $\underline{\eta}t < w - w'$ and

$$W(t,w) - W(t,w') \leq (w - \underline{\eta}t - (w' - \overline{\eta}t)) \leq (w - w')\frac{\overline{\eta}}{\underline{\eta}}.$$

Suppose $W(t,w) < w'$ and let $\vartheta(w,w')$ be defined by $W(\vartheta(w,w'),w) = w'$; note that $\vartheta(w,w') \leq (w-w')\underline{\eta}^{-1}$. Then $W(t,w) = W(t-\vartheta(w,w'),w')$ and

$$W(t,w) - W(t,w') = W(t-\vartheta(w,w'),w') - W(t,w') \leq \overline{\eta}\vartheta(w,w') \leq (w-w')\frac{\overline{\eta}}{\underline{\eta}}.$$

To see that $W(t,w)$ is Lipschitz continuous in t, we note that for $t' > t$, $W(t,w) - W(t',w) \leq \overline{\eta}(t'-t)$. Let K_W be the Lipschitz constant for W.

Now we turn to T_{k+1}. Remember, $T_{k+1}(w)$ is defined by $\int_0^{T_{k+1}(w)} \zeta(W(s,w))ds = log(U^{k+1})$. Let $w' < w$. Because $W(s,w') < W(s,w)$ and ζ is increasing, $T_{k+1}(w') > T_{k+1}(w)$. Let K_ζ be the Lipschitz constant for ζ. We then have

$$\zeta(W(s,w)) - \zeta(W(s,w')) \leq \zeta(W(s,w') + K_W(w-w')) - \zeta(W(s,w')) \leq K_\zeta K_W(w-w').$$

Therefore, $\int_0^{T_{k+1}(w)} \zeta(W(s,w))ds \leq \int_0^{T_{k+1}(w)} \zeta(W(s,w'))ds + K_W K_\zeta (w-w')T_{k+1}(w)$. Thus we have

$$\begin{aligned}
&\int_{T_{k+1}(w)}^{T_{k+1}(w')} \zeta(W(s,w'))ds + \int_0^{T_{k+1}(w)} \zeta(W(s,w'))ds = \int_0^{T_{k+1}(w)} \zeta(W(s,w))ds \\
\Rightarrow &\int_{T_{k+1}(w)}^{T_{k+1}(w')} \zeta(W(s,w'))ds \leq K_W K_\zeta (w-w')T_{k+1}(w) \leq K_W K_\zeta (w-w')\Theta \\
\Rightarrow &(T_{k+1}(w') - T_{k+1}(w))\zeta(S^*(\overline{R})) \leq K_W K_\zeta (w-w')\Theta
\end{aligned}$$

Thus $T_{k+1}(w)$ is Lipschitz in w as long as neither X or Y have reached S^* between the arrival of the kth and $k+1$st Z. □

Let $(\tau^{*,c}, d^{*,c}_\tau, \alpha^{*,c}_\tau)$ and $(\tau^{*,\Delta}, d^{*,\Delta}_\tau, \alpha^{*,\Delta}_\tau)$ be the continuous time and discrete time outcomes induced by the mechanism defined in Theorem 1. Given our construction, it should be intuitive that $(\tau^{*,\Delta}, d_\tau^{*,\Delta}, \alpha_\tau^{*,\Delta})$ under X_t converges to $(\tau^{*,c}, d_\tau^{*,c}, \alpha_\tau^{*,c})$ under Y_t. The following technical Lemma verifies this fact. It is immediate from this Lemma that P's and A's utility from $(\tau^*, d_\tau^*, \alpha_\tau^*)$ in discrete time converges to their utility from $(\tau^*, d_\tau^*, \alpha_\tau^*)$ in continuous time.

Lemma O.A.4. $(\tau^{*,\Delta}, d_\tau^{*,\Delta}, \alpha_\tau^{*,\Delta}) \xrightarrow{p} (\tau^{*,c}, d_\tau^{*,c}, \alpha_\tau^{*,c})$.

Proof. Let us focus on the event E^L in which $U_k < exp(-\overline{F}\zeta(\overline{Y})\Delta))$ for all $k \in \{1, ..., L\}$ for some large L. This restriction on U_k implies $T_k^\Delta(W^{k,\Delta}) \geq \overline{F}\Delta$ and $W^\Delta(T_k^\Delta(W^{k,\Delta}), W^{k,\Delta}) \leq \widetilde{Y}$. As $\Delta \to 0$, the probability of E^L converges to 1. The probability there are L arrivals of Z before $\tau^{*,c}$ or $\tau^{*,\Delta}$ is less than $(1 - F(\overline{R}^*))^L$, which goes to 0 as $L \to \infty$. To show convergence of $\tau^{*,c}, \tau^{*,\Delta}$, it is therefore sufficient to show that for any $\delta > 0$ and finite L, $\mathbb{P}(|\tau^{*,c} \wedge \overline{T}^L - \tau^{*,\Delta} \wedge \overline{T}^{L,\Delta}| > \delta | E^L) < \delta$ for sufficiently small Δ. Because M, M^X can only increase at the arrival of Z, $|M^X_{\overline{T}^{k,\Delta}} - M_{\overline{T}^k}| < \overline{\eta}\Delta$. If we can show that $\tau^{*,c} \in (\overline{T}^k, \overline{T}^{k+1})$ implies $\tau^{*,\Delta} \in (\overline{T}^{k,\Delta}, \overline{T}^{k+1,\Delta})$ with high probability for all $k \in \{1, ..., L\}$ (and vice versa), then convergence of $d_\tau^{*,\Delta}, \alpha_\tau^{*,\Delta}$ to $d_\tau^{*,c}, \alpha_\tau^{*,c}$ is straightforward.

We argue that there exists $(\Gamma_0, \Gamma_0^T, \Gamma_1, ..., \Gamma_L^T) \in \mathbb{R}_+^{2(L+1)}$ such that, as long as neither Y nor X has reached S^* before $\overline{T}^k$ or $\overline{T}^{k,\Delta}$ respectively, we have $|min\{\widetilde{Y}, W^{k,\Delta}\} - W^k| \leq \Gamma_k\Delta$ and $|T_k^\Delta(W^{k,\Delta}) - T_k(W^k)| \leq \Gamma_k^T\Delta$. We proceed by induction. Taking T_0^Δ, $T_0 = 0$, this holds for $k = 0$ by taking $\Gamma_K = \overline{\eta}$ and $\Gamma_0^T = 0$. For the inductive step, suppose Y and X have both realized k Z arrivals and have not stopped and $|min\{\widetilde{Y}, W^{k,\Delta}\} - W^k| \leq \Gamma_k\Delta$.

We start by putting bounds on the difference between $T_{k+1}^\Delta(w)$ and $T_{k+1}(w)$ in the event that neither X nor Y stop between the kth and $k+1$st arrival of Z. If $w \leq \widetilde{Y}$, then $G(t, w) = G^\Delta(t, w)$, which implies $|T_{k+1}^\Delta(w) - T_{k+1}(w)| \leq \Delta$. When $w > \widetilde{Y}$, for $t \geq n^F(w)\Delta$

$$G^\Delta(t, w) = exp(-\bar{\xi}(w)\Delta)G(t - n^F(w)\Delta, w).$$

Thus $G^\Delta(t, w) \leq G(t, w)$. Take some $\delta_S \in (0, S^*(\overline{R}^*) - \underline{Y})$ and let Δ be small enough that $W(t+\overline{F}\frac{\zeta(\overline{R}^*)}{\zeta(S^*(\overline{R}^*)-\delta_S)}\Delta, w) > S^*(\overline{R}^*)-\delta_S$. Then we have

$$\begin{aligned}
&\int_{n^F(w)}^{t} \zeta(W(s-n^F(w),w))ds + \bar{\xi}(w)\Delta \\
&\leq \int_0^t \zeta(W(s,w))ds + \overline{F}\zeta(\overline{R}^*)\Delta \\
&\leq \int_0^t \zeta(W(s,w))ds + \int_t^{t+\overline{F}\frac{\zeta(\overline{R}^*)}{\zeta(S^*(\overline{R}^*)-\delta_S)}\Delta} \zeta(S^*(\overline{R}^*)-\delta_S)ds \\
&\leq \int_0^{t+\overline{F}\frac{\zeta(\overline{R}^*)}{\zeta(S^*(\overline{R}^*)-\delta_S)}\Delta} \zeta(W(s,w))ds
\end{aligned}$$

We therefore have $G(t+\overline{F}\frac{\zeta(\overline{R}^*)}{\zeta(S^*(\overline{R}^*)-\delta_S)}\Delta, w) \leq G^\Delta(t,w)$. From the definition of T_k^Δ, T_k, we have $|T_{k+1}^\Delta(w)-T_{k+1}(w)| \leq (\overline{F}\frac{\zeta(\overline{R}^*)}{\zeta(S^*(\overline{R}^*)-\delta_S)}+1)\Delta$.

We can then bound the difference between $T_{k+1}^\Delta(W^{k,\Delta}), T_{k+1}(W^k)$ in the event that neither X, Y stop between the kth and $k+1$st arrival of Z. Because $T_{k+1}(w) = T_{k+1}(min\{\widetilde{Y}, w\})$, we have

$$\begin{aligned}
|T_{k+1}^\Delta(W^{k,\Delta}) - T_{k+1}(W^k)| &\leq |T_{k+1}^\Delta(W^{k,\Delta}) - T_{k+1}(W^{k,\Delta})| \\
&\quad + |T_{k+1}(min\{\widetilde{Y}, W^{k,\Delta}\}) - T_{k+1}(W^k)|] \\
&\leq (\overline{F}\frac{\zeta(\overline{R}^*)}{\zeta(S^*(\overline{R}^*)-\delta_S)}+1)\Delta + K_T|min\{\widetilde{Y}, W^{k,\Delta}\}) - W^k| \\
&\leq (\overline{F}\frac{\zeta(\overline{R}^*)}{\zeta(S^*(\overline{R}^*)-\delta_S)} + K_T\Gamma_k + 1)\Delta
\end{aligned}$$

Define $\Gamma_{k+1}^T = \overline{F}\frac{\zeta(\overline{R}^*)}{\zeta(S^*(\overline{R}^*)-\delta_S)} + K_T\Gamma_k + 1$.

We now argue that the values of X and Y just after the $k+1$st arrival of Z will be close. In the event E^L, $W^\Delta(T_{k+1}^\Delta(W^{k,\Delta}), W^{k,\Delta}) = W(T_{k+1}^\Delta(W^{k,\Delta}) - n^F(W^{k,\Delta})\Delta, min\{\widetilde{Y}, W^{k,\Delta}\})$ and so

$$\begin{aligned}
&|W^\Delta(T_{k+1}^\Delta(W^{k,\Delta}), W^{k,\Delta}) - W(T_{k+1}(W^k), W^k)| \\
&= |W(T_{k+1}^\Delta(W^{k,\Delta}) - n^F(W^{k,\Delta})\Delta, min\{\widetilde{Y}, W^{k,\Delta}\}) - W(T_{k+1}(W^k), W^k)| \\
&\leq K_W\big(|min\{\widetilde{Y}, W^{k,\Delta}\} - W^k| + |T_{k+1}(W^k) - T_{k+1}^\Delta(W^{k,\Delta}) + n^F(W^{k,\Delta})\Delta|\big) \\
&\leq K_W\big(\Gamma_k + \Gamma_{k+1}^T + \overline{F}\big)\Delta
\end{aligned}$$

We know $|Z^{k+1,\Delta} - Z^{k+1}| \leq \overline{\eta}\Delta$. Let $\Gamma_{k+1} = K_W\big(\Gamma_k + \Gamma^T_{k+1} + \overline{F}\big) + \overline{\eta}$. From the definition of $W^{k+1,\Delta}$ and W^{k+1}, we have $|min\{\widetilde{Y}, W^{k+1,\Delta}\} - W^{k+1}| \leq \Gamma_{k+1}\Delta$, completing the inductive step.

We next argue that if either Y or X stop before the $k+1$st jump, then the other process stops too with a probability converging to 1 as $\Delta \to 0$. Let $\hat{\theta}(w,b)$ be length of time, in the absence of a Z arrival, until Y crosses some threshold b when starting at $w \geq b$, defined by $W(\hat{\theta}(w,b), w) = b$. Define $\theta(w,m) = \hat{\theta}(w, S^*(m))$. We similarly define the length of time $\hat{\theta}^\Delta(w^\Delta, b)$ it takes X to cross b when starting at w^Δ in the absence of Z arrival, defined by $W^\Delta(\hat{\theta}^\Delta(w^\Delta, b), w^\Delta) = b$, and $\theta^\Delta(w^\Delta, m) = \hat{\theta}^\Delta(w^\Delta, S^*(m))$. Let $D^\Delta_S = \max_{y \in [\underline{Y}, \overline{Y} - \overline{\eta}\Delta]} S^*(y) - S^*(y - \overline{\eta}\Delta)$. Let $m = M_{T_k(W^{k-1})}$ and $m^\Delta = M^X_{T^\Delta_k(W^{k-1,\Delta})}$ be the maximum of Y and X after the kth arrival of Z. Because $|m - m^\Delta| \leq \overline{\eta}\Delta$, we must have $S^*(m^\Delta) \geq S^*(m) - D^\Delta_S$. Because it takes at most $\overline{F}\Delta$ length of time for W^Δ to reach $\widetilde{Y}$ and, from $\widetilde{Y}$ it takes at most $\hat{\theta}(min\{\widetilde{Y}, w\}, S^*(m) - D^\Delta_S)$ length of time to reach $S^*(m^\Delta)$, we have $\theta^\Delta(w,m) \leq \hat{\theta}(min\{\widetilde{Y}, w\}, S^*(m) - D^\Delta_S) + \overline{F}\Delta$.

We now want to show, for small Δ, $\hat{\theta}(w,b)$ is Lipschitz when $w < \widetilde{Y}$ and $b < S^*(\overline{R})$. If $w' < w$, then $W(t,w) < w'$ if $t > \underline{\eta}^{-1}(w - w')$. Having reached w', then length of time to reach b is then $\hat{\theta}(w', b)$; thus $\hat{\theta}(w,b) - \hat{\theta}(w',b) \leq \underline{\eta}^{-1}(w - w')$. Similarly, for $b' \in (S^*(\overline{R}^*), b)$ we have $0 \leq \hat{\theta}(w, b') - \hat{\theta}(w,b) \leq \underline{\eta}^{-1}(b - b')$. Let $K_{\hat{\theta}} = \underline{\eta}^{-1}$ be the Lipschitz constant for $\hat{\theta}$.

Suppose Y stops at $S^*(m)$. This implies that $U^{k+1} \leq G(\theta(W^k, m), W^k)$. Suppose $Z^{k+1,\Delta}$ arrives prior to X reaching $S^*(m^\Delta)$. This implies $U^{k+1} \geq G^\Delta(\theta^\Delta(W^{k,\Delta}, m^\Delta) + \Delta, W^{k,\Delta}) \geq G^\Delta(\hat{\theta}^\Delta(W^{k,\Delta}, S(m) - D^\Delta_S) + \Delta, W^{k,\Delta})$.[25] Because ζ, W are Lipschitz,

[25]We add Δ to θ^Δ to account of the rounding up of θ^Δ to a date in $\{0, \Delta, ...\}$.

so is $G(t, w)$. Let K_G be a Lipschitz constant for G. Then

$$
\begin{aligned}
&|G(\theta(W^k, m), W^k) - G^\Delta(\hat{\theta}^\Delta(W^{k,\Delta}, S(m) - D_S^\Delta) + \Delta, W^{k,\Delta})| \\
&\leq G(\theta(W^k, m), W^k) - G(\hat{\theta}^\Delta(W^{k,\Delta}, S(m) - D_S^\Delta) + (\overline{F}\frac{\zeta(\overline{R}^*)}{\zeta(S^*(\overline{R}^*) - \delta_S)} + 1)\Delta, W^{k,\Delta}) \\
&\leq K_G\big(|W^k - W^{k,\Delta}| + (\overline{F}\frac{\zeta(\overline{R}^*)}{\zeta(S^*(\overline{R}^*) - \delta_S)} + 1)\Delta \\
&+ |\hat{\theta}(W^k, S^*(m)) - \hat{\theta}^\Delta(W^{k,\Delta}, S^*(m) - D_S^\Delta)|\big) \\
&\leq K_G\big(|W^k - W^{k,\Delta}| + (\overline{F}\frac{\zeta(\overline{R}^*)}{\zeta(S^*(\overline{R}^*) - \delta_S)} + 1)\Delta) \\
&+ |\hat{\theta}(W^k, S^*(m)) - \hat{\theta}(min\{\widetilde{Y}, W^{k,\Delta}\}, S^*(m) - D_S^\Delta)| + \overline{F}\Delta\big) \\
&\leq \big(K_G(\Gamma_k(1 + K_{\hat{\theta}}) + (2\overline{F}\frac{\zeta(\overline{R}^*)}{\zeta(S^*(\overline{R}^*) - \delta_S)} + 1)\big)\Delta) + K_{\hat{\theta}} D_S^\Delta
\end{aligned}
$$

Thus, the probability $U^{k+1} \in (G^\Delta(\theta^\Delta(W^{k,\Delta}, m^\Delta) + \Delta, W^{k,\Delta}), G(\theta(W^k, m), W^k))$ is at most $K_G(|W^k - W^{k,\Delta}|(1 + K_{\hat{\theta}}) + K_{\hat{\theta}} D_S^\Delta + \overline{F}\frac{\zeta(\overline{R}^*)}{\zeta(S^*(\overline{R}^*) - \delta_S)}\Delta)$. As $\Delta \to 0$, this goes to 0 and the probability that X stops at $S^*(m^\Delta)$ conditional on Y stopping at $S^*(m)$ goes to one. The case when X at $S^*(m^\Delta)$ first is similar and the probability that Y stops at $S^*(m)$ goes to one as well. □

Proof of Theorem 1 for a Search Process

The proof of Theorem 1 for a search process immediately follows from the following two Lemmas. The first Lemma shows that the upper-bound on our relaxed problem, given by P's utility from $(\tau^\Delta, d_\tau^\Delta, \alpha_\tau^\Delta)$, converges to P's utility from $(\tau^*, d_\tau^*, \alpha_\tau^*)$ as $\Delta \to 0$. The second Lemma shows that $(\tau^*, d_\tau^*, \alpha_\tau^*)$ is dynamically individually rational and there is no dynamically individually rational continuous time mechanism that does strictly better than $(\tau^*, d_\tau^*, \alpha_\tau^*)$.

Lemma O.A.5. *P's utility from $(\tau^*, d_\tau^*, \alpha_\tau^*)$ is optimal in the limit as $\Delta \to 0$:*

$$\lim_{\Delta \to 0} \mathbb{E}^\Delta[e^{-r\tau^\Delta}(d_\tau^\Delta(u_P(\alpha_\tau^\Delta) - \nu) + \nu)] = \lim_{\Delta \to 0} \mathbb{E}^\Delta[e^{-r\tau^*}(d_\tau^*(u_P(\alpha_\tau^*) - \nu) + \nu)].$$

Proof. Convergence is obvious when $S^*(Y_0) > Y_0$. Consider the case in which $S^*(Y_0) = Y_0$. For $\Delta > 0$, $\alpha^\Delta(X_0)$ represents an upper-bound on P's utility from

$(\tau^\Delta, d_\tau^\Delta, \alpha_\tau^\Delta)$. As $|S^\Delta(X_0) - X_0| \to 0$, $\mathbb{P}(\tau^\Delta > \tau_+(X^1)) \to 0$ and

$$\begin{aligned}&\mathbb{E}[e^{-r\tau^\Delta} u_P(\alpha^\Delta(X_0))\mathbb{1}(\tau^\Delta < \tau_+(X^1))\\ &+ e^{-r\tau_+(X^1)} J(X_{\tau_+(X^1)}, X_{\tau_+(X^1)})\mathbb{1}(\tau^\Delta > \tau_+(X^1))] \xrightarrow{\Delta\to 0} u_P(\alpha^*(Y_0))\end{aligned}$$

Optimality in the limit then follows from the fact that $u_P(\alpha^*(Y_0)) = \lim_{\Delta\to 0} u_P(\alpha^\Delta(X_0))$.

Now suppose $S^*(Y_0) < Y_0$. For some small $\delta \in (0, S^*(\overline{R}^*) - \underline{Y})$, let $\overline{\zeta}^\Delta = \max_y \zeta^\Delta(y)$ and $\underline{\eta}^\delta = \min_{y \geq S^*(\overline{R}) - \delta} |\eta(y)|$. Let ℓ^Δ be the maximum number of grid points between $S^\Delta(M)$ and $S^*(M)$, which is at most $\max_{m\in[X_0, \overline{R}^\Delta]} \frac{|S^*(m) - S^\Delta(m)|}{\underline{\eta}^\delta \Delta}$ for small Δ. Because S^* is continuous and S^Δ is monotone, convergence to S^* is uniform; therefore, $\ell^\Delta\Delta \to 0$ as $\Delta \to 0$. At $\tau^* \wedge \tau^\Delta$, we know $\tau^* \vee \tau^\Delta$ arrives in at most $\ell^\Delta + 1$ periods if there is no Z arrival; thus,

$$\mathbb{P}(|\tau^* - \tau^\Delta| \leq (\ell^\Delta + 1)\Delta) \geq 1 - \overline{\zeta}^\Delta(\ell^\Delta + 1)\Delta,$$

and so $\lim_{\Delta\to 0} \mathbb{P}(|\tau^* - \tau^\Delta| \leq (\ell^\Delta + 1)\Delta) = 1$. In the event no Z arrives between $\tau^* \wedge \tau^\Delta$ and $\tau^* \vee \tau^\Delta$, the split amount will be the same at both τ^* and τ^Δ. Convergence of P's expected utility from $(\tau^\Delta, d_\tau^\Delta, \alpha_\tau^\Delta)$ to his expected utility from $(\tau^*, d_\tau^*, \alpha_\tau^*)$ follows. □

Lemma O.A.6. $(\tau^*, d_\tau^*, \alpha_\tau^*)$ *is dynamically individually rational in continuous time. There exists no dynamically individually rational continuous time mechanism which yields strictly higher utility than* $(\tau^*, d_\tau^*, \alpha_\tau^*)$.

Proof. That $(\tau^*, d_\tau^*, \alpha_\tau^*)$ is dynamically individually rational in continuous time follows from the same arguments as in the case of a diffusion process.

Suppose there was a dynamically individually rational $(\tau^{**}, d_\tau^{**}, \alpha_\tau^{**})$ which in continuous time gave P an expected utility of J^{**} with $J^{**} > \lim_{\Delta\to 0} \mathbb{E}^\Delta[e^{-r\tau^\Delta}(d_\tau^\Delta(u_P(\alpha_\tau^\Delta) - \nu) + \nu)]$. We consider below the same extension of u_P and u_A to allow for $\alpha_\tau \in (-\infty, 1]$ as in the proof of Theorem 1 for a diffusion process. For our continuous time process, it is without loss to only consider mechanisms which are measurable to histories of the form $h_t' = (k, Z^1, U^1, ..., Z^k, U^k, \overline{U}_t^{k+1})$ where k is determined by $t \in (\overline{T}^k, \overline{T}^{k+1})$ and $\overline{U}_t^{k+1}$ is the upper-bound on U^{k+1} implied by no arrival of Z^{k+1} by time t. This information maps out then entire history $\{Y_s : 0 \leq s \leq t\}$ and so any mechanism measurable with respect to Y will be measurable with respect to such histories.

We define a discrete time mechanism which is measurable with respect to these same histories; for the same reasons as in continuous time, this is without loss. Fix some small $\delta_\alpha > 0$. Let $\hat{\tau}^{**} = \lceil \tau^{**} \rceil_\Delta$, $\hat{d}^{**}_\tau = d^{**}_\tau$ and $\hat{\alpha}^{**}_\tau = (\alpha^{**}_\tau - \delta_\alpha)\mathbb{1}(u_A(1 - \alpha^{**}_\tau + \delta_\alpha)) > X_{\hat{\tau}^{**}}) + \bar{\alpha}\mathbb{1}(u_A(1 - \alpha^{**}_\tau + \delta_\alpha)) < X_{\hat{\tau}^{**}})$ where $u_A(1-\bar{\alpha}) > \overline{Y}$ (where $\bar{\alpha}$ is defined in the same way as in the diffusion process case). As $\Delta \to 0$, it is straightforward to see that P's discrete time value from $(\hat{\tau}^{**}, \hat{d}^{**}_\tau, \hat{\alpha}^{**}_\tau)$ is bounded below by $J^{**} - u'(0)\delta_\alpha + O(\delta_\alpha^2)$ and, because $\delta_\alpha > 0$, it is dynamically individually rational for small enough Δ. Because δ_α is arbitrary, we can find a discrete-time dynamically individually rational mechanism whose value to P is arbitrarily close to J^{**}, which is strictly greater than $\mathbb{E}^\Delta[e^{-r\tau^\Delta}(d^\Delta_\tau(u_P(\alpha^\Delta_\tau) - \nu) + \nu)]$ for small Δ, a contradiction of the fact that $(\tau^\Delta, d^\Delta_\tau, \alpha^\Delta_\tau)$ delivers an upper-bound on P's value of the discrete time optimal mechanism. Therefore, no such $(\tau^{**}, d^{**}_\tau, \alpha^{**}_\tau)$ exists. $\square$

Role of Assumptions on Stochastic Processes

We assumed that every search process had $\eta(y) < 0\ \forall y$ and $inf\{t : Y_t = \underline{Y}\} = \infty$. These assumptions simplify the statement of the optimal mechanism by allowing us to write the decision of when to stop and reach a split as the first time Y_t crosses $S^*(M_t)$. Suppose that we allowed for Y to reach y' that are absorbing until the arrival of a new $Z_t > y'$-namely, y' such that $\eta(y') = 0$. It is easy to incorporate such absorbing y' into our discrete-time relaxed problem. As we take the limit as $\Delta \to 0$, if y' is absorbing and $S^*(M_t) = y'$, then P will offer a split at a Poisson rate $\gamma^*(M_t) \in \mathbb{R}_+ \cup \{\infty\}$ where $\gamma^*(M_t) = \infty$ implies an immediate split at $S^*(M_t) = y'$.

This observation allows us to extend our results to cover the stationary search without recall model. Once a Z_t offer has been rejected, $Y_{t+} = \underline{Y}$, at which point P proceeds to make a TIOLI offer with demand $\alpha^*(M_t)$ at a rate $\gamma^*(M_t)$. The case of stationary search with recall is less interesting. Because A's outside option is only increasing, the optimal mechanism will have no delay: either P makes an immediate TIOLI offer or A takes his outside option immediately.

The proofs of Theorem 1 its supporting Lemmas rely on several properties that are shared by the diffusion and search processes. These two processes are both approximated by X which satisfy:

1. *Downward Continuity:* if $X_s < X_t$ for $t < s$, then for every $x \in (X_s, X_t) \cap \mathcal{G}^\Delta$, $\exists q \in (t, s)$ such that $X_q = x$.

2. *Upward Crossing Independence:* for every x' and $\tau' := min\{t : \ X_t \geq x'\}$, the distribution of $X_{\tau'}$ is the same for all $X_0 < x'$.

The intuition for our results does not rely on these two properties; their role is purely to simplify the analysis. These two properties are not exclusive to diffusion and search processes. We can extend our results to other Y processes, such as pure-jump processes, which can be approximated by X possessing the two properties above. For a pure-jump process, even though we have an absorbing state at y_-, it will be optimal to stop immediately at y_- because there can be no further change of Y and any additional delay will be purely wasteful. Thus a pure-jump process has an optimal mechanism taking the same form as in Theorem 1.

Proof of Proposition 2

Proof. Suppose that, after some history h_t with $(Y_t, M_t) = (y, m)$, the continuation mechanism was not constrained Pareto efficient. Let $V^c(y, m)$ be the continuation value for A after such a history. By replacing the continuation mechanism after h_t with an optimal dynamically individually rational mechanism for P when we add a promise-keeping constraint to deliver at least $V^c(y, m)$ continuation value to A, we can increase P's payoffs. Moreover, this change would not affect the incentive of A to take his outside option before t because A evaluates the continuation value after h_t in the new mechanism as the same as in the old mechanism. Therefore, A's continuation value at every history h_s that might lead to h_t is exactly the same: if A had no strict incentive to take his outside option at s in the old mechanism, then he will have no strict incentive to take his outside option at s in the new mechanism. □

Comparative Statics Proofs

Proof of Proposition 3

Proof. Consider the discrete time relaxed problem. For a diffusion process, an increase in $\mu(x')$ leads to an increase in $q_+(x')$ and decrease in $q_-(x')$. For a search process, an increase in $\zeta(x')$ leads to an increase in $\zeta^\Delta(x')$. Let $\underline{\mu}, \overline{\mu}$ and $\underline{\zeta}^\Delta, \overline{\zeta}^\Delta$ be two pairs of functions such that $\underline{\mu}(x) \leq \overline{\mu}(x)$ and $\underline{\zeta}^\Delta(x) \leq \overline{\zeta}^\Delta(x)$ for all x. We show that P's utility from the solution to the discrete time problem in 2 increases when we increase μ from $\underline{\mu}$ to $\overline{\mu}$ or ζ^Δ from $\underline{\zeta}^\Delta$ to $\overline{\zeta}^\Delta$. The Proposition then follows by taking the limit as $\Delta \to 0$.

We start with a diffusion process. Let $J(x, m; \mu)$ be P's value of the discrete time relaxed problem when the drift is μ. Let $m = max\{x \in \mathcal{G}^\Delta : x < \overline{R}^\Delta, (x, x) \in \mathcal{R}\}$ be the highest value of x at which $(\tau^\Delta, d^\Delta_\tau, \alpha^\Delta_\tau)$ does not immediately stop, where we define $\overline{R}^\Delta$ as the breakdown threshold in the optimal mechanism when $\mu = \underline{\mu}$. Consider the optimal choice of $(\tau, d_\tau, \alpha_\tau)$ at (m, m) subject to $RDIR(m)$ when we fix continuation value at $\tau_+(m + \epsilon_m)$ for P and A to be ν and $X_{\tau_+(m+\epsilon_m)}$ respectively. P's problem for a fixed μ is

$$\begin{aligned}\widehat{J}(m, m; \mu) = \max_{(\tau, d_\tau, \alpha_\tau)} \; & \mathbb{E}_m[e^{-r\tau} u_P(\alpha) \mathbb{1}(\tau < \tau_+(m + \epsilon_m)) \\ & + e^{-r\tau_+(m+\epsilon_m)} \nu \mathbb{1}(\tau > \tau_+(m + \epsilon_m))|\mu], \\ \text{subject to } m \leq \; & \mathbb{E}_m[e^{-r\tau} u_A(1 - \alpha) \mathbb{1}(\tau < \tau_+(m + \epsilon_m)) \\ & + e^{-r\tau_+(m+\epsilon_m)} X_{\tau_+(m+\epsilon_m)} \mathbb{1}(\tau > \tau_+(m + \epsilon_m))|\mu].\end{aligned} \tag{16}$$

We condition on μ to mean conditional on the law of motion for X given the function μ that determines $q_+(X_t)$ and $q_-(X_t)$.

When $\mu = \underline{\mu}$, the solution to 16 and its value $\widehat{J}(m, m; \underline{\mu}) = J(m, m; \underline{\mu})$. We now modify 16 to allow P to choose the drift $\mu(x) \in [\underline{\mu}(x), \overline{\mu}(x)]$ at each x. Let Ξ be the set of μ such that $\mu(x) \in [\underline{\mu}(x), \overline{\mu}(x)]$ $\forall x$ and let P's choice of μ be $\mu^P \in \Xi$. P's value when he can choose μ is then

$$\widehat{J}^*(m, m) := \max_{\mu^P \in \Xi} \widehat{J}(m, m; \mu^P). \tag{17}$$

Because P could always choose $\mu^P = \underline{\mu}$, we have $\widehat{J}^*(m, m) \geq J(m, m; \underline{\mu})$.

Let λ be the Lagrange multiplier on $RDIR(m)$ in 17 and define $\mathcal{J}^*$ as

$$\begin{aligned}\mathcal{J}^*(x; \mu^P) = \max_{(\tau, d_\tau, \alpha_\tau)} \mathbb{E}_x[e^{-r(\tau \wedge \tau_+(m+\epsilon_m))} \big\{ & [d_\tau(u_P(\alpha) - \lambda u_A(1 - \alpha) - \nu + \lambda X_\tau) \\ & + \nu - \lambda X_\tau] \mathbb{1}(\tau < \tau_+(m + \epsilon_m)) \\ & + [\nu - \lambda X_{\tau_+(m+\epsilon_m)}] \mathbb{1}(\tau > \tau_+(m + \epsilon_m)) \big\} |\mu^P].\end{aligned}$$

For each choice of μ^P, the same arguments as in Lemma A.3 imply that P uses a (S, γ)-split threshold at (m, m) and a constant demand α^m and that $\mathcal{J}^*(x; \mu^P)$ is strictly increasing in x. Let $\mathcal{J}^{**}(m) = \max_{\mu^P \in \Xi} \mathcal{J}^*(m; \mu^P)$ be the Lagrangian for our problem in 17.

We next show that $\mathcal{J}^{**}(x) = \mathcal{J}^*(x;\overline{\mu})$ for all x. Let (S^m, γ^m) be the optimal (S,γ)-threshold in 17. At $x \le S^m$ we have $\mathcal{J}^{**}(x) = \mathcal{J}^*(x;\mu^P) = u_P(\alpha^m)$ for any choice of μ^P. At $x > S^m$ we have[26]

$$\mathcal{J}^{**}(x) = \max_{\mu^P(x)\in[\underline{\mu}(x),\overline{\mu}(x)]} \mathbb{1}(x = S^m + \epsilon_{S^m})\gamma^m u_P(\alpha^m) + e^{-r\Delta}\big[\frac{1}{2}(\frac{\sigma^2(x)}{\sigma_0^2} + \frac{\mu^P(x)}{\sigma_0})\mathcal{J}^{**}(x+\epsilon)$$
$$+ \frac{1}{2}(\frac{\sigma^2(x)}{\sigma_0^2} - \frac{\mu^P(x)}{\sigma_0})\mathcal{J}^{**}(x-\epsilon) + (1 - \frac{\sigma^2(x)}{\sigma_0^2})\mathcal{J}^{**}(x)\big].$$

Because $\mathcal{J}^{**}(x)$ is strictly increasing in x, it is clear that the uniquely optimal choice of $\mu^P(x)$ is $\overline{\mu}(x)$. Because this holds for all x, $\overline{\mu}$ is P's optimal μ^P.

Now consider $m_- = max\{x \in \mathcal{G}^\Delta : \ x < m, \ (x,x) \in \mathcal{R}\}$. Replacing m with m_- in the problem for $\widehat{J}^*(m,m)$ and let $\widehat{J}^*(m,m)$ replace ν as P's continuation value at $\tau_+(m_- + \epsilon_{m_-})$, we can apply the same arguments to conclude that $\mu^P = \overline{\mu}$ is optimal for m_-. Proceeding by induction, we conclude that choosing $\overline{\mu}$ is always optimal.

Let us return of our original relaxed problem when the choice of μ is fixed to $\overline{\mu}$. We can define a mechanism when $\mu = \overline{\mu}$ by, at each (m', m'), using the corresponding $(S^{m'}, \gamma^{m'})$-split thresholds and split amount $\alpha^{m'}$ found above until $\tau_+(m' + \epsilon_{m'})$. This new mechanism gives P a time zero continuation value of $\widehat{J}^*(X_0, X_0; \overline{\mu})$, which is higher than $\widehat{J}^*(X_0, X_0; \underline{\mu}) = J(X_0, X_0; \underline{\mu})$. This new mechanism is easily seen to be dynamically individually rational by the imposition of $RDIR(m')$ at each (m', m'). Therefore, P's utility from our discrete time relaxed problem when $\mu = \overline{\mu}$ is higher than when $\mu = \underline{\mu}$.

We can repeat the same argument for a search process, replacing the choice of $\mu^P \in \Xi$ with $\zeta^{\Delta,P} \in \Xi$ when we redefine Ξ to be the set of functions ζ^Δ such that $\zeta^\Delta(x) \in [\underline{\zeta}^\Delta(x), \overline{\zeta}^\Delta(x)]$ for all x. At $x \le S^m$ we have $\mathcal{J}^{**}(x) = \mathcal{J}^*(x;\zeta^{\Delta,P}) = u_P(\alpha^m)$ for any choice of $\zeta^{\Delta,P} \in \Xi$. Our Lagrangian $\mathcal{J}^{**}$ at $x > S^m$ is

$$\mathcal{J}^{**}(x) = \max_{\zeta^{\Delta,P}(x)\in[\underline{\zeta}^\Delta(x),\bar{\zeta}(x)]} \mathbb{1}(x = S^m + \epsilon_{S^m})\gamma^m u_P(\alpha)$$
$$+ e^{-r\Delta}\big[\zeta^{\Delta,P}(x)\mathbb{E}_Z[\mathcal{J}^{**}(max\{Z,x\})] + (1 - \zeta^{\Delta,P}(x))\mathcal{J}^{**}(x - \epsilon_x)].$$

Again, by the same arguments as in Lemma A.3, $\mathbb{E}_Z[\mathcal{J}^{**}(max\{Z,x\})] > \mathcal{J}^{**}(x - \epsilon_x)$, so the optimal choice of $\zeta^{\Delta,P}$ is $\overline{\zeta}^\Delta$ The same arguments used for a diffusion process imply that P's utility in our discrete time relaxed problem is higher when $\zeta^\Delta = \overline{\zeta}^\Delta$. □

[26]Because $\mathcal{J}^{**}(x) \ge J(x,x;\underline{\mu})$, it is never optimal to stop below $\overline{R}^\Delta$.

Proof of Proposition 4

Proof. Fixing the mechanism, the result is shown clearly through the HJB equation for P's continuation value. For a diffusion process, when $J(y,m)$ is P's continuous time continuation value at $(Y_t, M_t) = (y, m)$, we have

$$rJ(y,m) = \frac{\sigma^2(y)}{2}\frac{\partial^2 J(y,m)}{\partial y^2}.$$

By Lemma A.5, $J(y,m) \geq \nu \geq 0$, so we have $\frac{\partial^2 J(y,m)}{\partial y^2} \geq 0$. Therefore, an increase in $\sigma(y)$ increases $J(y,m)$. For a search process, when $y < \widetilde{Y}$ we take $\zeta(y) = \frac{-\eta(y)}{\mathbb{E}_Z[max\{Z,y\}]}$ to ensure Y is a martingale. P's value function is then

$$\begin{aligned} rJ(y,m) &= \eta(y)\frac{\partial J(y,m)}{\partial y} + \zeta(y)\big[\mathbb{E}_Z[J(max\{Z,y\}, max\{Z,m\})] - J(y,m)\big] \\ &= \eta(y)\big(\frac{\partial J(y,m)}{\partial y} - \frac{\mathbb{E}_Z[J(max\{Z,y\}, max\{Z,m\})] - J(y,m)]}{\mathbb{E}_Z[max\{Z, min\{\widetilde{Y}, y\}\}]}\big) \end{aligned}$$

which is decreasing in $\eta(y)$ because, by $J(y,m) \geq 0$ and $\eta(y) < 0$,

$$\frac{\partial J(y,m)}{\partial y} - \frac{\mathbb{E}_{\mathbb{Z}}[J(max\{Z,y\}, max\{Z,m\})] - J(y,m)]}{\mathbb{E}_Z[max\{Z,y\}]} < 0.$$

Replacing J with V, we get that A's continuation value is also increasing in $\sigma(y), -\eta(y)$. By reoptimizing the mechanism, P may do even better. □

Proof of Proposition 5

Proof. We first consider an upper-bound on the value of J^* for a diffusion or search process. Let us replace DIR in P's problem 1 with only $RDIR(Y_0)$:

$$\begin{aligned} &\sup_{(\tau, d_\tau, \alpha_\tau)} J(\tau, d_\tau, \alpha_\tau) \\ &\text{subject to } V(\tau, d_\tau, \alpha_\tau) \geq Y_0. \end{aligned} \tag{18}$$

It is straightforward to show by the arguments in the proof of Theorem 1 that the solution to 18 will take a stationary form: for some constants $\widetilde{S}$, $\widetilde{R}$ and $\widetilde{\alpha}$, the solution is

$$\widetilde{\tau} = inf\{t: \ Y_t \notin (\widetilde{S}, \widetilde{R})\}, \ \widetilde{d}_\tau = \mathbb{1}(Y_\tau \leq \widetilde{S}), \ \widetilde{\alpha}_\tau = \widetilde{\alpha}.$$

In order to satisfy $RDIR(Y_0)$, it must be that $u_A(1-\widetilde{\alpha}) \geq \widetilde{S}$.

Take the pure jump process which jumps from Y_0 either to $\widetilde{S}$ or $\widetilde{R}$, where the jump rates are pinned down by the martingale restriction and capacity constraint. By Theorem 3 of Zhong (2017), this pure jump process gives P the highest payoff among all martingale processes that satisfy the capacity constraint when we fix the mechanism to be $(\widetilde{\tau}, \widetilde{d}_\tau, \widetilde{\alpha}_\tau)$. A similar observation holds for A's continuation value. All that is left is the verify is that $(\widetilde{\tau}, \widetilde{d}_\tau, \widetilde{\alpha}_\tau)$ is dynamically individually rational for a pure jump process. Because the mechanism satisfies $RDIR(Y_0)$, A's continuation value is weakly greater than Y_0 as long as $Y_t = Y_0$. Because $u_A(1-\widetilde{\alpha}) \geq \widetilde{S}$, the split at $\widetilde{S}$ is individually rational for A, so the mechanism is dynamically individually rational. $\square$