Online Appendix

Part 1: Nonstationarity

Consider a responder who can choose between drawing from two different distributions. The criteria relevant to assessing these distributions are (i) the likelihood of finding an alternative that exceeds the reservation price and (ii) the value of any such alternative that is found. In essence, these criteria express how good is the "right tail" of a distribution; the "left tail" is not relevant because alternatives below the reservation price will be rejected. In order to formalize this preference we first introduce Convex Second Order Stochastic Dominance (CX-SSD).

Definition (Convex Second Order Stochastic Dominance). For any cdf F, let $\bar{F}(x) = 1 - F(x)$. We say that a distribution with cdf G dominates one with cdf H by CX-SSD if

$$\int_y^\infty \bar{G}(x)\,dx \geq \int_y^\infty \bar{H}(x)\,dx \quad \forall y \in \mathbb{R}, \tag{OA.1}$$

where the dominance is strict if there exists at least one y for which this inequality is strict.

Convex second-order stochastic dominance is also known as an *increasing convex order* (Müller and Stoyan 2002, Shaked and Shanthikumar 2007), which is a consequence of the following equivalent definition. Let X and Y be random variables distributed according to (respectively) cdf G and cdf H. Then G dominates H by CX-SSD if, for any increasing convex function ϕ, we have $\mathbb{E}[\phi(X)] \geq \mathbb{E}[\phi(Y)]$. The machine learning literature refers to CX-SSD as "stochastically optimistic dominance" (Osband et al. 2014).

An agent's preference over possible distributions of alternatives is consistent with the ordering of those distributions by CX-SSD. Van den Berg (1990) observes that if the distribution of outside alternatives changes in a way that is either a mean-preserving spread of the original distribution or an improvement (in the sense of first-order stochastic dominance) over the original distribution, both of which are specific cases of CX-SSD, then the searcher will prefer this change and so reservation prices will rise accordingly. Our next lemma, which will be used in the proof of Theorem A1, generalizes this result to CX-SSD.

Lemma A2 (Preferences over Alternative Distributions Are Consistent with CX-SSD) *Let $\xi_F^{SA}(t)$ be the policy that solves the nonstationary single-agent problem given by* (1)*, and let $\xi_G^{SA}(t)$ be the policy solving a modified version of that problem in which the distribution of outside alternatives $F(x;t)$ is replaced by $G(x;t)$ such that, for all $t \in [0,\infty)$ and all $y \in \mathbb{R}$:*

$$\int_y^\infty (\bar{G}(x;t) - \bar{F}(x;t))\,dx \geq 0, \tag{OA.2}$$

i.e., $G(x;t) \succeq_{CX-SSD} F(x;t)$. Then $\xi_G^{SA}(t) \geq \xi_F^{SA}(t)$ for all $t \in [0,\infty)$. If the inequality in (OA.2) *is strict for all $t \in [0,T]$ and all $y \in \mathbb{R}$, then $\xi_G^{SA}(t) > \xi_F^{SA}(t)$ for all $t \in [0,T)$.*

Proof. First we show that under conditions stated in the lemma it holds that $\xi_G^{SA}(t) \geq \xi_F^{SA}(t)$, for all $t \in [T,\infty)$. If the search horizon is finite this statement follows trivially from the fixed points $\xi_G^{SA}(T) = \xi_F^{SA}(T) = u_F$, as given in (1). If the search horizon is infinite, it follows from (1) that the optimal policies on $[T,\infty)$ are monotone threshold ones given by $\xi_F^{SA}(t) = \rho^{-1}\lambda(T)\int_{\xi_F^{SA}(t)}^{\infty}\bar{F}(x;T)dx$ and $\xi_G^{SA}(t) = \rho^{-1}\lambda(T)\int_{\xi_G^{SA}(t)}^{\infty}\bar{G}(x;T)dx$; for the proof of optimality of monotone threshold policies under stationary search please refer to Lippman and McCall (1976). Then, for all $t \in [T,\infty)$ we have

$$\xi_G^{SA}(t) - \xi_F^{SA}(t) = \rho^{-1}\lambda(T)\int_{\xi_G^{SA}(t)}^{\infty}\bar{G}(x;T)dx - \rho^{-1}\lambda(T)\int_{\xi_F^{SA}(t)}^{\infty}\bar{F}(x;T)dx,$$

$$\xi_G^{SA}(t) - \xi_F^{SA}(t) = \rho^{-1}\lambda(T)\left(\int_{\xi_F^{SA}(t)}^{\infty}(\bar{G}(x;t) - \bar{F}(x;t))dx + \int_{\xi_G^{SA}(t)}^{\xi_F^{SA}(t)}\bar{G}(x;t)dx\right). \tag{OA.3}$$

Here, if $\xi_G^{SA}(t) < \xi_F^{SA}(t)$, then the left-hand side (LHS) of (OA.3) is negative. However, if this is the case, then the right-hand side (RHS) of the same equation is positive – the first RHS integral is positive by application of CX-SSD condition ($\int_y^{\infty}(\bar{G}(x;t) - \bar{F}(x;t))\,dx \geq 0$ for all $y \in \mathbb{R}$), and the second is positive directly from $\xi_G^{SA}(t) < \xi_F^{SA}(t)$. Thus, by contradiction, we have $\xi_G^{SA}(t) \geq \xi_F^{SA}(t)$, for all $t \in [T,\infty)$.

In the next step we demonstrate that $\xi_G^{SA}(t) \geq \xi_F^{SA}(t)$ for $t \in [0,T)$ as well. From (1) it follows that the differential equation for ξ_G^{SA} is

$$\xi_G^{SA\prime}(t) = \rho\xi_G^{SA}(t) - \lambda(t)\int_{\xi_G^{SA}(t)}^{\infty}(x - \xi_G^{SA}(t))\,dG(x;t) = \rho\xi_G^{SA}(t) - \lambda(t)\int_{\xi_G^{SA}(t)}^{\infty}\bar{G}(x;t)\,dx. \tag{OA.4}$$

Subtracting (1) from (OA.4) yields

$$\xi_G^{SA\prime}(t) - \xi_F^{SA\prime}(t) = \rho(\xi_G^{SA}(t) - \xi_F^{SA}(t)) + \lambda(t)\int_{\xi_F^{SA}(t)}^{\xi_G^{SA}(t)}\bar{F}(x;t)\,dx - \lambda(t)\int_{\xi_G^{SA}(t)}^{\infty}(\bar{G}(x;t) - \bar{F}(x;t))\,dx. \tag{OA.5}$$

Since $\int_y^{\infty}(\bar{G}(x;t) - \bar{F}(x;t))\,dx \geq 0$ for all $y \in \mathbb{R}$, it follows from equation (OA.5) that

$$\xi_G^{SA\prime}(t) - \xi_F^{SA\prime}(t) \leq \rho(\xi_G^{SA}(t) - \xi_F^{SA}(t)) + \lambda(t)\int_{\xi_F^{SA}(t)}^{\xi_G^{SA}(t)}\bar{F}(x;t)\,dx. \tag{OA.6}$$

If $\xi_G^{SA}(t) < \xi_F^{SA}(t)$ then RHS of inequality (OA.6) is negative, from which we deduce that $\xi_G^{SA}(t) - \xi_F^{SA}(t) < 0 \Rightarrow \xi_G^{SA\prime}(t) - \xi_F^{SA\prime}(t) < 0$. So if there exists a $t^* \in [0,T]$ such that $\xi_G^{SA}(t^*) - \xi_F^{SA}(t^*) < 0$, then $\forall t \geq t^*$: $\xi_G^{SA}(t) - \xi_F^{SA}(t) < 0$. However, from the first part of the proof we have that $\xi_G^{SA}(T) - \xi_F^{SA}(T) \geq 0$, which contradicts the existence of such t^*. As a result, $\xi_G^{SA}(t) - \xi_F^{SA}(t) \geq 0$ for all $t \in [0,T]$.

If $\int_y^{\infty}(\bar{G}(x;t) - \bar{F}(x;t))\,dx > 0$ for all $y \in \mathbb{R}$, then inequality (OA.6) is strict and so $\xi_G^{SA}(t) - \xi_F^{SA}(t)$ is a strictly decreasing function on $[0,T]$. Thus, $\xi_G^{SA}(T) - \xi_F^{SA}(T) \geq 0$ (as shown in first part of the proof) implies $\xi_G^{SA}(t) - \xi_F^{SA}(t) > 0$ for all $t \in [0,T)$. $\square$

Theorem A1 (Conditions for Optimality of Exploding Offers in Absence of Stationarity) *Denote $\bar{u} = \max_{t\in[0,T]}\xi_R^{SA}(t) + u_{MR}$, and let the following two conditions on the responder's distribution of outside alternatives hold*

(i) *Distribution does not change in the vicinity of the reservation price, i.e., $F_R(z;t_1) = F_R(z;t_2)$ for all t_1, t_2, y, z such that $t_1 < t_2, y \in [\bar{u},\infty), z \in [u_{MR},\bar{u}]$.*

(ii) *Expectation of the right tail above $\bar{u}$ deteriorates over time, i.e., $\int_{\bar{u}}^{\infty}\bar{F}_R(x;t)\,dx$ is non-increasing in t. Then, Theorem 1 holds even when the distribution of the responder's alternatives is not stationary.*

Proof. The proof follows the exact same steps as the proof of Theorem 1, with the exception of Part 1 of the proof (showing that $P(\xi_R, t'_O, t'_D) \le P(\xi_R, t_O, t_D)$) which we prove here.

Although the translation property (A.6) does not hold without stationarity of the responder's search, we can derive and rely on a weaker inequality version of it. Let $G_R(x;t)$ be a distribution which is identical to $F_R(x;t)$ for $x \in [u_{MR}, \bar{u}]$, and consists of two mass points outside of this region: when drawing from $G_R(x;t)$ there is a $F_R(u_{MR};t)$ probability of drawing an alternative of value 0, and $\bar{F}_R(\bar{u};t)$ probability of drawing value $\frac{1}{\bar{F}_R(\bar{u};t)} \int_{\bar{u}}^{\infty} \bar{F}_R(x;t)\, dx$. Notice that a responder who holds the proposer's offer which he plans to accept at t_{ED}, will be indifferent about drawing from either F_R or G_R as his optimal strategy at that point never involves accepting offers with value below u_{MR}, nor rejecting offers above $\bar{u}$ (both probability of finding an alternatives with value in access of $\bar{u}$ and conditional expectation of the draw when this occurs are also the same). Because of that, it follows from (2) that such a responder will use the same policy no matter which of these two distributions he is drawing from. Notice also that $\forall t_1, t_2 \in [0,T] \mid t_1 < t_2 : G_R(x;t_1) \succeq_{CX-SSD} G_R(x;t_2)$. This enables us to use Lemma A2 to show that for $t_{ED} < t'_{ED}$:

$$\xi_R^{HO}(t; u_{MR}, t_{ED}) \ge \xi_R^{HO}(t - t_{ED} + t'_{ED}; u_{MR}, t'_{ED}), \forall t \in (t_O, t_{ED}), \tag{OA.7}$$

which is the inequality variant of the translation property. Applying (OA.7) to (A.5) yields that the proposer's offer will be accepted with probability

$$P(\xi_R, t_O, t_D) \le \exp\left\{ -\int_0^{t_O} \lambda_R(u) \bar{F}_R(\xi_R(u); u)\, du \int_{t_O}^{t_{ED}} \lambda_R(v) \bar{F}_R(\xi_R^{HO}(v - t_{ED} + t'_{ED}; u_{MR}, t'_{ED}); v)\, dv \right\}.$$

As the optimal threshold when holding the proposer's offer is in $[u_{MR}, \bar{u}]$ and λ_R is stationary is follows that

$$P(\xi_R, t_O, t_D) \le \exp\left\{ -\int_0^{t_O} \lambda_R(u) \bar{F}_R(\xi_R(u); u)\, du \int_{t'_{ED} - t_{ED} + t_O}^{t'_{ED}} \lambda_R(v) \bar{F}_R(\xi_R^{HO}(v; u_{MR}, t'_{ED}); v)\, dv \right\},$$

$$P(\xi_R, t_O, t_D) \le \exp\left\{ -\int_0^{t_O} \lambda_R(u) \bar{F}_R(\xi_R(u); u)\, du \int_{t_O}^{t'_{ED}} \lambda_R(v) \bar{F}_R(\xi_R^{HO}(v; u_{MR}, t'_{ED}); v)\, dv \right\} = P(\xi_R, t'_O, t'_D),$$

which demonstrates the claim. □

While the distribution of alternatives can indeed change over time under the conditions of Theorem A1, it changes in a way that it becomes less attractive over time for the responder. More precisely, changes can occur in the tails of the distribution, but if they occur in the right tail, they need to be such that they decrease the expected value of drawing from the right tail. No such condition is needed for the left tail; any kind of changes in the distribution are allowed to occur in the region below u_{MR}, but such changes are inconsequential in situations where the responder already holds the proposer's offer, as optimal policy necessitates rejection of all offers with value below u_{MR}. The technically restrictive part of these conditions is that changes in the distribution cannot occur in the vicinity of the reservation value when holding the proposer's offer (the $[u_{MR}, \bar{u}]$ region).

Part 2: Bargaining

Bargaining in the presence of search has been the focus of Muthoo (1995), Gantner (2008), and Baucells and Lippman (2004). Bargaining can also be incorporated into our model by embedding the Nash bargaining solution, as in Baucells and Lippman (2004) – we give the results for that variant here.

According to the Nash solution, the total value of the mutual deal ($u_{MR}+u_{MP}$) is split in such a way that each agent receives the net present value of the rest of their search ($\xi_R^{SA}(t)$ or $\xi_P^{SA}(t)$), with remaining value split evenly between the agents. This gives us the expressions for the value of a mutual deal for each agent as a function of time at which the mutual deal is concluded:

$$u_{MR}(t_D) = \frac{1}{2}\Big(u_{MR} + u_{MP} - \xi_P^{SA}(t_D) + \xi_R^{SA}(t_D)\Big), \tag{OA.8}$$

$$u_{MP}(t_D) = \frac{1}{2}\Big(u_{MR} + u_{MP} + \xi_P^{SA}(t_D) - \xi_R^{SA}(t_D)\Big). \tag{OA.9}$$

Thus, how much each agent will make in a mutual deal evolves over time, depending on the responder's bargaining power $\alpha_R(t) := \xi_R^{SA}(t) - \xi_P^{SA}(t)$. Note that if both agents' search process is stationary with infinite horizon, then the bargaining power will remain constant. As in the basic model, we assume existence of a time t such that $u_{MR} + u_{MP} > \xi_R^{SA}(t) + \xi_P^{SA}(t)$ (if this is not true, we are in a degenerate situation where the agents never make a deal and just follow a single-player search policy). We also assume that if either agent withdraws from bargaining due to accepting an outside alternative, the other party will be informed immediately.

A natural question is what does the deadline even mean with bargaining? We take it to mean a time at which the bargaining process will conclude, and payoffs will be realized. By giving a shorter deadline, the proposer can force the bargaining process to conclude faster. Another relevant question is, can the agents search while they are bargaining? We considered three different possible setups for this variant, which are detailed below.

An observation that can be made across all three settings is that stationarity results in relatively simple solutions, which are consistent with other work on bargaining. However, nonstationarity causes difficult tradeoffs, as it creates an incentive for the proposer to strategically time and prolong negotiations in order to reach an agreement when her bargaining power is higher.

Setup 1: One-sided commitment. This setup is a straightforward extension of our model, in which the only thing which is different is the endogenous offer value. One peculiar feature of this setup is that during the negotiation process, the proposer is prevented from doing outside search while the responder is not. Yet, this is not without meaning, as in many markets (job market and real estate come to mind), the buyer or the recruiter are the only ones which have to formally commit by issuing offers while no commitment is required from the other side prior to signing. In this setup, the proposer faces a complex tradeoff between extracting more value out of the responder and increasing the probability the responder accepts, which is shown in the following lemma.

Lemma A3 (The Proposer's Tradeoff Under Bargaining with One-sided Commitment) *Let conditions of Theorem A1 hold and let α_R be weakly decreasing. Then for any offer (t_O, t_D), the probability that the responder accepts it is weakly decreasing in t_D but the value the proposer gets from a mutual deal is weakly increasing in t_D.*

Proof. $u_{MP}(t)\uparrow t$ follows directly from (OA.9) and $\alpha_R(t)\downarrow t$. Denote by $P(\xi_R, t_O, t_D, u_{MR})$ the probability that the responder accepts an offer where bargaining begins at t_O, ends at t_D, and pays the responder u_{MR} if accepted. Applying Van den Berg (1990, eq. 8), we can express the this probability as

$$P(\xi_R, t_O, t_D, u_{MR}) = \exp\left\{ -\int_0^{t_O} \lambda_R(u)\bar{F}_R(\xi_R(u); u)\, du \int_{t_O}^{t_D} \lambda_R(v)\bar{F}_R(\xi_R^{HO}(v; u_{MR}, t_D); v)\, dv \right\}. \quad \text{(OA.10)}$$

Analogously to the proof of Theorem A1 it holds that $P(\xi_R, t_O, t_D, u_{MR})\downarrow t_D$. However, in this setting u_{MR} is not fixed but depends on t_D, or more precisely: from (OA.8) and $\alpha_R(t)\downarrow(t)$ we have that $u_{MR}(t_D)\downarrow t_D$. From $\xi_R^{HO}(v; u_{MR}, t_D)\uparrow u_{MR}$ also we have that $P(\xi_R, t_O, t_D, u_{MR})\uparrow u_{MR}$, since u_{MR} only appears in (OA.10) as an argument of ξ_R^{HO}. Finally, from $u_{MR}(t_D)\downarrow t_D$ and $P(\xi_R, t_O, t_D, u_{MR})\downarrow t_D \uparrow u_{MR}$ it follows that $P(\xi_R, t_O, t_D, u_{MR}(t_D))\downarrow t_D$, completing the proof. □

This tradeoff makes the full game hard to resolve, however some specific versions are quite tractable. In particular, if we use the standard stationary search model, one part of the tradeoff above disappears: the bargaining power no longer fluctuates over time, which enables the following result.

Proposition A1. (Solution of the Stationary Bargaining Game with One-sided Commitment) *Let both agents' searches be stationary with infinite horizon. Then bargaining both commences and ends at time 0* ($t_O = t_D = 0$).

Proof. In the stationary infinite horizon search, the values of residual search are constant, i.e., $\exists \xi_R, \xi_P \in \mathbb{R}^+$ s.t. $\forall t_D \in \mathbb{R}^+ : \xi_R^{SA}(t_D) = \xi_R$ and $\xi_P^{SA}(t_D) = \xi_P$. Thus $\alpha_R(t_D)$ is also constant and value of the mutual deal will not depend on the time the deal is concluded. As $\exists t \in \mathbb{R}^+$ s.t. $u_{MR} + u_{MP} > \xi_R + \xi_P$, the responder will accept the deal at any t_D if he is still available. It is hence in the proposer's interest to set $t_O = t_D$, since that will result in weakly higher probability that the offer is accepted (for $t_O < t_D$, the proposer can find an alternative with higher value with positive probability during bargaining), while also giving the proposer a higher payoff (because of discounting and constant bargaining power).

It remains to show that $t_O = t_D = 0$ is optimal, which we will do by contradiction. Assume $t_O = t_D > 0$ is optimal. Then, the proposer prefers searching till t_O and making an offer at t_O to the certain payoff she can receive at time 0 by making an offer immediately (this payoff is $(u_{MR} + u_{MP} - \xi_R + \xi_P)/2$ as per bargaining solution). However at time t_O the proposer will find herself in an almost identical situation to one at time 0 (the only difference is that she is unsure if the responder is still available). Thus, if she preferred delaying her offer at time 0, she will also prefer delaying it once at t_O, contradictory to optimality of t_O. □

This result is consistent with the search bargaining literature: bargaining both starting and ending at time 0 is one of the main results in Muthoo (1995), Baucells and Lippman (2004), and Gantner (2008). While this is the result of the main model in all three papers, some extensions in this literature identify situations where this conclusion is invalidated. Gantner (2008) shows that information asymmetry can cause delays, while Baucells and Lippman (2004) show that the same can occur if the bargaining parties are forced to use actual offers as a disagreement threat (rather than expectation of the search).

Setup 2: Mutually committing. Alternatively, we can consider a setup where neither of the agents can search during the bargaining process. While we would consider this setup to be generally less realistic than the first one, there certainly do exist situations where the bargaining parties are physically in the same room till bargaining completes, preventing any other activities. At a glance, there are strong incentives to make the bargaining process as short as possible in this setting, as prolonging bargaining destroys value by "shrinking the pie." Indeed, that is true under stationarity where the solution is identical to the one with one-sided commitment.

Proposition A2. (Solution of the Stationary Bargaining Game with Mutual Commitment) *Let both agents' searches be stationary with infinite horizon. Then bargaining both commences and ends at time 0 ($t_O = t_D = 0$).*

Proof. Analogous to the proof of Proposition A1. □

Yet, in absence of stationarity other incentives also emerge. As in the first setup, the responder has an incentive to conclude a deal at a time where her bargaining power is larger. Additionally, there is a benefit to starting the bargaining process as early as possible: blocking the responder from conducting his search.

Lemma A4 (The Proposer's Tradeoff Under Bargaining with Mutual Commitment) *Given any t_O, the optimal deadline is the one which solves*

$$\underset{t_D \in [t_O, \infty)}{\arg\max} \quad e^{-\rho(t_D - t_O)}(u_{MR} + u_{MP} - \alpha_R(t_D))/2, \tag{OA.11}$$

$$s.t. \quad e^{-\rho(t_D - t_O)}(u_{MR} + u_{MP} + \alpha_R(t_D))/2 \geq \xi_R^{SA}(t_O). \tag{OA.12}$$

In case there is no t_D which satisfies constraint (OA.12)*, it is optimal to set $t_D = t_O$.*

Proof. Follows directly from the objective of (OA.11) being the expected payoff to the proposer, conditional on the responder accepting the offer. The constraint (OA.12) is that the payoff to the responder (given by the left-hand side), needs to be at least as high as the value of responder's search if he rejects the offer. □

Setup 3: No Commitment. Lastly, we can consider the setup in which both agents are free to continue their search during bargaining process. Here the results are much simpler than in the previous two setups.

Proposition A3. (Equilibrium with No Commitment while Bargaining). *In any equilibrium, the proposer issues an offer at time 0, but can leave a longer deadline.*

Proof. During bargaining, both players are aware if the other accepts an outside alternative, and will switch to single agent search policy if that occurs. Analogously to the full information model of Section 3, keeping the deadline t_D fixed, this causes the payoffs of both players and their optimal policies to be decreasing in t_O (by application of Van den Berg 1990, Theorem 3). Hence, the proposer's payoff is always maximized at $t_O = 0$. □

This is a consequence of there being no cost of giving out an offer (otherwise there is an opportunity cost of forgoing search), but there being two benefits: a) the responder will become more selective during the bargaining process, and b) the proposer will be able to switch to a more profitable search policy if the responder withdraws from bargaining. Lastly, stationary search results in an identical result as in the two other setups.

Proposition A4. (Solution of the Stationary Bargaining Game with No Commitment) *Let both agents' searches be stationary with infinite horizon. Then bargaining both commences and ends at time 0* $(t_O = t_D = 0)$.

Proof. Optimality of $t_O = 0$ is given by Proposition A3, so it remains to show optimality of $t_D = 0$. As in Proposition A1, stationary infinite horizon search implies constant bargaining power α_R. Denote by $S(t_D)$, the proposer's expectation of searching if he starts bargaining at $t_O = 0$, sets deadline $t_D > 0$ and uses an optimal policy. Then the proposer's expected payoff is a convex combination of $S(t_D)$ and $e^{-\rho t_D}(u_{MP} + u_{MR} - \alpha_R)/2$ (the latter here is her possible payoff from making a deal with the responder). However, it holds that $(u_{MP} + u_{MR} - \alpha_R)/2 \geq \xi_P^{SA}(0) \geq S(t_D)$, and thus $(u_{MP} + u_{MR} - \alpha_R)/2$ – which the proposer can get by setting $t_D = 0$ – is always greater than convex combination above. □

Part 3: Nonstationary Multiplayer Game

Here we consider the same multi-proposer model as in Section 4.4 of the paper, but without making the stationarity assumption as in Proposition 3. There are technical limitations to obtaining the same results if the search process is not stationary. These limitations stem from the (in)ability to solve explicitly for the function which maps each proposer's type to the optimal timing of the offer. However, there are multiple ways in which assumptions can be tightened in order to enable tractability of this function. For example, if the proposers do not have outside alternatives (other than the responder), then this mapping becomes tractable as the optimal timing only depends on maximizing the probability the responder accepts, without the need to balance it with other alternatives. This in turn enables us to express the arrival of the offers from such strategic proposers as a Poisson process.

Proposition A5. (Multiple Proposers with No Outside Alternatives). *Let ξ_R be a differentiable decreasing threshold policy for the responder, and let the probability density function $f_R(x;t) := F_R'(x;t)$ exist. Then, if $\lambda_P = 0$, the optimal exploding offer of proposer i will be given at time $t_{Oi} = \min\{t \geq t_i \mid u_{MRi} \geq \xi_R(t)\}$. Consequently, the arrival stream of offers coming from such strategic proposers can be described by a nonstationary process with arrival rate $\lambda_O(t)$ given by*

$$\lambda_O(t) = -\xi_R'(t) \int_0^t f_R(\xi_R(t); s)\lambda_R(s)ds + \lambda_R(t)\bar{F}_R(\xi_R(t); t). \qquad \text{(OA.13)}$$

The first term of (OA.13) *is the arrival rate of strategically delayed offers, which all have value $\xi_R(t)$. The second term is the arrival rate of high value offers, which are made as soon as that proposer enters the game, and their value is the result of a random draw from $F_R(x;t)$ conditional on the draw being at least $\xi_R(t)$.*

Proof. Since $\lambda_P = 0$, the proposers maximize their expected payoff by maximizing the probability that the responder accepts. For each responder, if the value of her offer is sufficiently high so that the responder is willing to accept it $(u_{MRi} \geq \xi_R(t))$, it is optimal for the proposer to give an exploding offer immediately upon entering the game. However, if $u_{MRi} < \xi_R(t)$, the proposer will need to delay her offer till the first time at which the responder is willing to accept, hence the optimal offer timing for proposer i is given by $t_{Oi} = \min\{t \geq t_i \mid u_{MRi} \geq \xi_R(t)\}$.

Consider temporarily a version of this model where the time is divided into periods of width h, and the proposers can only make their offers at beginning of each of these periods. Then, at any time t which is a multiple of h, the responder will receive offers from all proposers who entered the game thus far and whose value of offers falls between $\xi_R(t)$ and $\xi_R(t-h)$; those are the ones who strategically delayed their offers. Thus the probability that the responder receives at least one strategically delayed offer at time t is given by $1-\exp\{-\int_0^t \lambda_R(s)\,(F_R(\xi_R(t-h);s)-F_R(\xi_R(t);s))\,ds\}$. Hence, the arrival rate of strategically delayed offers (λ_O^{SD}) in the full model will be given by the following limit when the width of these periods approaches 0:

$$\lambda_O^{SD}(t)=\lim_{h\to 0}\frac{1-\exp\{-\int_0^t \lambda_R(s)\,(F_R(\xi_R(t-h);s)-F_R(\xi_R(t);s))\,ds\}}{h}.$$

Applying L'Hôpital rule yields

$$\lambda_O^{SD}(t)=-\xi_R'(t)\int_0^t f_R(\xi_R(t);s)\lambda_R(s)ds.$$

Finally, summing this expression with the arrival rate for offers without strategic delay $(\lambda_R(t)\bar{F}_R(\xi_R(t);t))$ yields (OA.13). $\square$

Thus, even in this setting, it holds that the arrival stream of offers coming from strategic proposers can be represented by a Poisson process. Yet, this is only a partial result as there are still missing steps to establishing an equilibrium as in Proposition 3. The limiting factor here stems from us restricting proposers to only exploding offers in Proposition A5, which then yielded the arrival process described above. Had this arrival process satisfied the conditions of Theorem A1, then it would indeed be optimal for all of these proposers to make exploding offers even if they had other options. However, as can be seen from the proposition, this process is a peculiarly nonstationary one, where both value and arrival change over time in a way which violates conditions of Theorem A1. This nonstationarity in the induced arrival process is driven by the proposers who are less valued by the responder (have lower u_{MRi}) having a stronger incentive to delay their offers, lest it be rejected by the responder.

Part 4: Example Calculations

Proof of Equilibrium for Example 1. This is an extensive-form game with imperfect information but perfect recall; hence the appropriate solution concept is a sequential equilibrium (Osborne and Rubinstein 1994, chap. 12). We proceed by reducing the game via a sequential elimination of strategies that are strictly dominated for all nodes of an information set.

The proposer (if available) always accepts period-1 alternatives, and the responder (if available) always accepts period-2 alternatives, because those alternatives yield the game's strictly highest possible payoffs. Likewise, the responder has a strictly dominant strategy for responding to a proposer's offer that expires at time 1: he should accept it, receiving \$130, because rejecting it yields only \$100 in expectation (i.e., an 0.1 probability of receiving \$1000). These considerations imply that, if the proposer finds herself still available at time 1, then she will make an exploding offer to the responder at that point; after all, any other strategy gives her a strictly lower expected payoff owing to the possibility of the responder finding a period-2 alternative. With the dominated strategies thus eliminated, if the proposer made no offer at time 0 then it becomes a strictly dominant strategy for the responder to accept a period-1 alternative. The reason is that doing so yields him greater payoff (\$200) than both possible scenario which can occur if he rejects that alternative: \$130 if the proposer makes him an offer at time 1 or \$100 if she does not.

Thus in any sequential equilibrium, unless the proposer makes an offer at time 0, events can unfold in only one way: both agents accept the first outside alternative they find and, if neither finds anything before time 1, the proposer then makes an exploding offer that the responder accepts. If at least one of these two agents found an outside alternative, there is no mutual deal but the responder will still try to find an alternative in period 2. Hence the proposer has an expectation of $0.2 \times \$200 + 0.8 \times 0.1 \times \$130 = \$50.4$ and the responder has an expectation of $0.9 \times \$200 + 0.1 \times (0.8 \times \$130 + 0.2 \times 0.1 \times \$1000) = \$192.4$.

It remains to check for whether the proposer can improve on this expectation by making an offer at time 0. She cannot make an exploding offer, as it would be rejected by the responder (it yields \$130 to the responder whereas the responder's search yields \$190 in expectation). Also, making an offer that expires at time 1 is a strictly dominated strategy for the proposer because it forfeits her own period-1 search but does not increase the probability of acceptance (as compared with an exploding offer made at time 1). Yet the proposer can improve her expectation if she makes a time-0 offer that expires at time 2; we shall now establish that claim.

Consider the responder's optimal policy when holding such an offer. At time 1, our responder has an expected payoff of \$217: with probability 0.1 he will find the alternative valued at \$1000; otherwise, he will accept the proposer's \$130 offer. Therefore, a responder who finds an alternative of value \$200 in the first period is better-off rejecting it. Hence an offer that expires at the search horizon yields expected payoffs of $0.9 \times \$130 = \$117$ for the proposer and $0.1 \times \$1000 + 0.9 \times \$130 = \$217$ for the responder. Such an offer gives the proposer a higher expectation than does forgoing a time-0 offer. Thus, in any sequential equilibrium of the game, the proposer makes an offer at time 0 whose deadline is the search horizon (time 2). The existence of this equilibrium is guaranteed as every finite extensive-form game with perfect recall contains a sequential equilibrium (Kreps and Wilson 1982). □

Proof of Equilibrium for Example 2. As shown in the example itself, a dominant counter-offer policy for the proposer is to counter period-1 alternatives for the responder with a deadline offer of her own, setting the deadline to 2, if and only if the proposer has not found an alternative in the first period. We now enumerate payoffs for all the possible offer and deadline times, when the proposer is using this counter-offer policy and the responder reacts optimally.

If $t_O = 0$, the proposer commits to her offer before there is an opportunity for the responder to find any outside alternatives, thus the counter-offer policy becomes irrelevant and the payoffs are identical to those in Example 1.

If $t_O = t_D = 1$ the sequence of actions is as follows. If the proposer finds an outside alternative in the first period she will just accept it, leaving the responder to his own search. If the proposer does not find an alternative, but the responder does, the proposer will give a counter-offer with deadline 2. The responder will then reject the outside alternative (as it pays \$200 compared to \$217 in expectation from searching), search in period 2, accepting the alternative of value \$1000 if it is found, or accepting the proposer's counter-offer otherwise. Lastly, if the responder does not find an alternative in the first period, the proposer will make an exploding offer at time 1, which the responder will accept (payoff of \$130, whereas searching more yields \$100 in expectation). This gives a total payoff of $0.2 \times \$200 + 0.8 \times 0.9 \times 0.9 \times \$130 + 0.8 \times 0.1 \times \$130 = \$134.64$ to the proposer. Here the first additive term is the situation where proposer finds an outside alternative, second one is when she gives a counter-offer which is accepted, and the last one is where she gives an exploding offer at time 1. Analogously, the responder's payoff is $0.1 \times (0.2 \times 0.1 \times \$1000 + 0.8 \times \$130) + 0.9 \times (0.2 \times \$200 + 0.8 \times 0.1 \times \$1000 + 0.8 \times 0.9 \times \$130) = \$204.64$.

If $t_O = 1$ and $t_D = 2$ or if $t_O = 2$ and $t_D = 2$, the sequence of actions prior to time 1 will be the same as described in the preceding paragraph, while from time 1 onward in both of these cases the responder will search in period 2, accepting the outside alternative if found, or otherwise accepting the proposer's offer at time 2. This gives an expected payoff of $0.2 \times \$200 + 0.8 \times 0.9 \times \$130 = \$133.6$ to the proposer and $0.2(0.9 \times \$200 + 0.1 \times 0.1 \times \$1000) + 0.8(0.1 \times \$1000 + 0.9 \times \$130) = \$211.6$ to the responder.

Finally, notice that in all of these cases, the responder's optimal strategy is the same: the responder should always accept period-2 alternatives and accept period-1 alternatives unless an offer with deadline 2 is received from the proposer. A consequence of this is that the responder changing his offer time and/or deadline will just change the game payoff to one of the other ones enumerated here, hence the proposer cannot increase her payoff by deviating from the strategy $t_O = t_D = 1$. Since the responder is using his best response policy, neither can he; thus this is an equilibrium. □

Proof of Equilibrium for Example 3. Assume that the responder follows the optimal single-agent policy, which in this case is to accept the first outside alternative he finds. Then, if the proposer wants to make an exploding offer, she has only three strategies that are not strictly dominated. (i) She can make an offer at time 0 and set its value to \$190; this offer is always accepted and yields the proposer \$70. (ii) The proposer can make an offer at time 1 and set its value to \$100. The responder will accept this offer if he is still available at time 1 (which occurs with probability 0.1); however, the proposer can also undertake a first-period search for the \$200 alternative, which she will find with probability 0.2. Thus the proposer's expectation when making such an offer is $0.2 \times \$200 + 0.8 \times 0.1 \times \$160 = \$52.8$. (iii) She can make an offer at time 2 and set its value to 0. The responder will accept this offer only if he has found no alternative, hence this strategy yields an expectation $0.2 \times \$200 + 0.8 \times 0.1 \times 0.9 \times \$260 = \$58.72$ to the proposer.

Since deadline offers – made from time t_O to time t_D – that do not cause acceptance deterrence are dominated by exploding offers at t_D, the proposer's only other nondominated option is to use deadline offers in which x is set to the minimum possible value that *will* deter the responder from accepting an outside alternative. For an offer made at time 0 with deadline 2, we can find such x by solving $\$200 = 0.1 \times \$1000 + 0.9x$; re-arranging this expression yields $x = \$1000/9$. A proposer who makes this offer has expectation $0.9 \times (\$260 - \$1000/9) = \$134$. Other possible deadline offers must be checked in order to show that this one is indeed the best strategy when the responder employs the single-agent policy. Deterring the responder from accepting offers in period 2 is not feasible because that would require setting $x = \$1000$, which would result in a negative payoff for the responder. Thus the only deadline offer not eliminated is the one made at time 0, expiring at time 1, and with value \$200 (i.e., the lowest value that leads to acceptance deterrence). Such an offer is always accepted, but it yields the proposer only \$60 and so cannot qualify as her best response. Since neither agent can benefit from deviating, we have an equilibrium. □

References

Baucells, M., S. A. Lippman. 2004. Bargaining with search as an outside option: The impact of the buyer's future availability. *Decision Analysis* **1**(4) 235–249.

Gantner, A. 2008. Bargaining, search, and outside options. *Games and Economic Behavior* **62**(2) 417–435.

Kreps, D. M., R. Wilson. 1982. Sequential equilibria. *Econometrica* **50**(4) 863–894.

Lippman, S. A, J. W. Mamer. 2012. Exploding offers. *Decision Analysis* **9**(1) 6–21.

Lippman, S. A., J. McCall. 1976. The economics of job search: A survey. *Economic Inquiry* **14**(2) 155–189.

Mortensen, D. T. 1986. Job search and labor market analysis. *Handbook of Labor Economics* **2** 849–919.

Müller, A., D. Stoyan. 2002. *Comparison methods for stochastic models and risks*, vol. 389. Wiley.

Muthoo, A. 1995. On the strategic role of outside options in bilateral bargaining. *Operations Research* **43**(2) 292–297.

Osband, I., B. Van Roy, Z. Wen. 2014. Generalization and exploration via randomized value functions. *arXiv preprint* arXiv:1402.0635.

Osborne, M. J., A. Rubinstein. 1994. *A course in game theory*. MIT press.

Shaked, M., J. G. Shanthikumar. 2007. *Stochastic orders*. Springer Science & Business Media.

Van den Berg, G. J. 1990. Nonstationarity in job search theory. *The Review of Economic Studies* **57**(2) 255–277.