INSEAD
The Business School for the World®

Working Paper Series

2016/09/DSC

Be Patient yet Firm: Offer Timing, Deadlines, and the Search for Alternatives

Sasa Zorc
INSEAD, sasa.zorc@insead.edu

Ilia Tsetlin
INSEAD, ilia.tsetlin@insead.edu

20 February 2016

We model two agents who can benefit from a mutual deal or partnership yet are also involved in outside search. This generic situation is observed in various settings (e.g., the job market for experts) and involves several decisions. The proposer decides when to make an offer to the responder, what deadline to set, and how to handle her outside alternatives; the responder decides whether to accept the proposer's offer (if any) and also how to handle his own outside alternatives. We find that the proposer's optimal strategy is to set the shortest possible deadline and that this conclusion is robust to different sets of assumptions. Across all variations of the game, we find that deals are made at a time preferred by the less eager agent (i.e., the one who would rather sign a mutual deal later) unless the more eager agent can make an ultimatum offer.

Keywords: Search; Ultimatum; Exploding offer; Deadlines; Outside Alternatives; Search Deterrence.

Be Patient yet Firm: Offer Timing, Deadlines, and the Search for Alternatives

Saša Zorc, Ilia Tsetlin
INSEAD, 1 Ayer Rajah Avenue, Singapore 138676, Singapore, sasa.zorc@insead.edu, ilia.tsetlin@insead.edu
20 February 2016

We model two agents who can benefit from a mutual deal or partnership yet are also involved in outside search. This generic situation is observed in various settings (e.g., the job market for experts) and involves several decisions. The proposer decides when to make an offer to the responder, what deadline to set, and how to handle her outside alternatives; the responder decides whether to accept the proposer's offer (if any) and also how to handle his own outside alternatives. We find that the proposer's optimal strategy is to set the shortest possible deadline and that this conclusion is robust to different sets of assumptions. Across all variations of the game, we find that deals are made at a time preferred by the less eager agent (i.e., the one who would rather sign a mutual deal later) *unless* the more eager agent can make an ultimatum offer.

Key words: search, ultimatum, exploding offer, deadlines, outside alternatives, search deterrence

1. Introduction

In just July and August 2015, European football clubs paid in excess of $3 billion to acquire new players.[1] This is not an isolated incident because FIFA, football's governing body, requires that international transfers of players occur only during the stipulated bi-annual "transfer windows".[2] During these periods, many clubs are looking to fill positions on their teams for the following season and many are looking to sell some of their players, resulting in a frenetic period of high-stakes decisions. Football clubs must balance their desire to search for the best possible deal against the danger of coveted players signing with other teams – and all under time pressure from the transfer window's approaching end. Plenty of media attention is garnered by the flurry of offers made, the many deals signed, and the many falling through.

[1] "Summer Transfer Window 2015 Analysis," Onefootball Magazine, 1 September 2015, http://bit.ly/1HR4hPI; source data available at http://bit.ly/1N3VdCE.

[2] "Regulations on the Status and Transfer of Players," official FIFA document, 19 October 2003, http://fifa.to/1Obnahh.

In the summer of 2015, stakes were especially high for Real Madrid – the world's most valuable sports team[3] – as they sought to acquire a new goalkeeper. The team set its eyes on David de Gea, then Manchester United's goalkeeper. In most cases, such deals involve lengthy discussions and are not committed to paper until both sides are already in agreement. On this occasion, however, Real faced an ultimatum offer from United early in the transfer period: "£35 [$\sim$ $52] million within 24 hours or we are walking away and any deal is off".[4] Real was unwilling to meet this ultimatum and rejected it. Yet as we shall later recount, the story does not end there.

The situation just described provokes several questions. Did it benefit United to issue such an ultimatum? Could they have timed the offer better? Should they have given a longer deadline? How should have Real reacted to this ultimatum? What factors should be considered when making such decisions? These are the questions that we explore in this paper.

The scenario in question is hardly unique to football. A similar dilemma arises in several job markets for experts, notably the markets for new MBA graduates and assistant professors. Almost any business situation in which one party seeks a partner (e.g., the real estate market, the search for venture capital investors) has similar features as well. In the takeover context, a famous example of an ultimatum offer is Microsoft's bid for Yahoo! in 2008.[5]

In all these situations, there are two agents who could benefit from a mutual deal, with both of them also involved in a search for outside alternatives. One of these agents (the proposer) would like to make an offer, but the structure of that offer presents a challenge because the decision involves multiple trade-offs. She definitely wants to structure the offer in a way that increases the odds of acceptance; however, this goal conflicts with other criteria. On the one hand, the proposer seeks a timely response to her offer; on the other hand, she would like to maximize her own available time for finding a better alternative. The responder in this exchange faces a similar decision: Should he accept the offer, or does he stand more to gain by continuing to look for a better alternative?

The literature on *two-sided matching markets* considers multiple proposers and responders (Roth and Sotomayor 1992, Sönmez and Ünver 2011, Abdulkadiroğlu and Sönmez 2013). This literature aims to provide (centralized or decentralized) matching mechanisms which result in stable and efficient matching (Roth and Xing 1997, Pais 2008). In contrast, we focus on the respective agents' optimal strategies. Research in this area tends also to assume that all alternatives are available simultaneously; hence agents need not employ search strategies or consider stopping rules. Thus our paper contributes more to the literature on search, bargaining, and deadlines.

[3] "Forbes Announces the World's Most Valuable Sports Teams," Forbes, 15 July 2015, http://onforb.es/1X8ZfEX.

[4] The transfer window was from 1 July 2015 to 31 August 2015, and the ultimatum was announced on 13 July 2015: "Paper Round: Real Madrid Handed 24-hour Ultimatum over David de Gea," Eurosport, 13 July 2015, http://bit.ly/1TafifR.

[5] "Microsoft's Letter to Yahoo! Board of Directors," Microsoft News Center, 5 April 2008, http://bit.ly/MicroUltim.

The problem of one agent going through a search process is the subject of considerable literature. McCall (1970), Gronau (1971), Lippman and McCall (1976) and Lancaster and Chesher (1983) focus on the job search context. This question has been addressed more generally by the entire class of optimal stopping problems; reviews of the research on that topic are provided by Shiryaev (1976) and Ferguson (2012). The issues tackled in all of these works correspond closely to the responder's decision problem in our setting, so we will draw from this research when developing our model.

The literature on bargaining in the presence of search shares some of our goals in that it addresses the predicament of agents who must decide between making a mutual deal and searching for alternatives. Muthoo (1995) and Gantner (2008) follow the approach of Rubinstein (1982) in developing a game-theoretic model of bargaining, via alternating offers, in which one of the agents can search for outside alternatives instead of bargaining. Baucells and Lippman (2004) consider a similar setting but, instead of alternating offers, use the Nash bargaining solution's axiomatic approach (Nash 1950). Both of these models allow agents to choose the timing of their offers but *not* the deadlines. One of their key findings is the "no delay" result under complete or symmetric information: any mutual deal will be made at time zero.

In contrast, Lippman and Mamer (2012) and Tang et al. (2009) explore the effects of deadlines but *not* of the offer timing. Lippman and Mamer consider, in a setting with complete recall, the proposer's choice between making an "exploding" offer (i.e., one that expires unless accepted immediately) and one that is valid until the search horizon's end. Tang et al. consider ultimatum deadline games, a setting in which the responder is involved in a standard finite-horizon search with no recall while the proposer sets a deadline in order to maximize the responder's probability of accepting. Both of these papers address the complexity of choosing the correct deadline when the offer time is determined exogenously.

We unify the approaches of these two literature streams by considering a setup in which both time of the offer *and* its deadline are decision variables, exploring how to give offers optimally and how to respond to them. We find that the optimal strategy is starkly different when both offer timing and the deadline are taken into account. In contrast to results from previous research on deadlines and on bargaining in the presence of search, we find that the optimal strategy is always to make an exploding offer – although the best time to make that offer is a complex decision. For most parameter values, the proposer should make an offer only after a delay.

In real life, the exact rules governing how offers are made will vary. In some situations, such as that of the recruiter in a labor market, only one agent can make offers. In other cases, such as football clubs negotiating player transfers, either side can make this first step. Similarly, these offers might be issued as ultimatums or there could be some room for counteroffers. The value of the offer could be part of the decision or it could be fixed. Also, there may be regulations that place restrictions on minimum deadlines. The model we propose allows us to consider all these variants, and answer for every possible change in rules: who benefits, who is harmed, and when does the change even have an impact.

Table 1 serves as a roadmap of all the different settings considered in this paper. The basic model is presented in §2; that simple version relies on several constraining assumptions, which are relaxed later in the paper. In this way we can derive clean, closed-form solutions and develop the underlying intuition. The salient result is that the proposer's optimal strategy is to make an exploding offer. In §3 we explore several variants of the model (as described in the preceding paragraph); here the game no longer needs to have a clearly defined proposer, and counteroffers are possible. Across all of the variants, the mutual deal is made at the time preferred by the less eager agent (one who prefers signing a mutual deal later), unless the more eager agent can make an ultimatum offer.

Table 1 Overview of settings

Section	Additional considerations
2	Basic model
3.1	Either or both players can make an offer and/or commit to an ultimatum
3.2	There exists a min. deadline Δ such that no offer can expire sooner than Δ after being made
4.3	Values of alternatives are drawn from distributions, and their arrival rates are nonstationary
4.4	Distribution of outside alternatives evolves over time
4.5	Value of the offer can be altered by the proposer
4.6	Alternatives found earlier in search can be recalled

Our most general setting is presented in §4, where we show the optimality of making only exploding offers under a robust set of assumptions; that section is the paper's most technically demanding part. In this general model, the values of outside alternatives are drawn from an arbitrary distribution that can evolve over time, and the arrival rate of those alternatives is nonstationary. In §4.5 we allow the proposer to alter the value of her offer, and in §4.6 we allow alternatives that are found earlier in the search process to be recalled.

The conclusion that exploding offers are optimal persists throughout all the different variants and considerations that we undertake to examine. The only exception is that, if the *distribution* of the responder's outside alternatives is not stationary, then longer deadlines may become optimal – although this exception pertains only if that distribution improves over time (in the sense of convex second-order stochastic dominance; see §4.4). We make the concluding remarks and present the managerial implications in §5.

2. Basic Model

We develop a complete-information, game-theoretic model of two agents – for example, a job seeker and an employer – who are each involved in a search process. The proposer (the employer, "she") is searching for the right person to employ while the responder (the job seeker, "he") is searching for a job. Both parties maximize their respective expected utility. The proposer is looking to fill just one position, and the responder can accept only one job. Each party's search for outside alternatives follows a standard finite-horizon search model as characterized by the assumptions to follow. Most of these assumptions are relaxed later in the paper.[6]

Assumptions about the search process

A1. *Each agent's search is defined as a finite-horizon Poisson process, from* 0 *to* T, *parameterized by the arrival rates of outside alternatives* (*i.e., by* λ_P *for the proposer and* λ_R *for the responder*).[7]

A2. *If an agent reaches the search horizon* T *without accepting any alternative, then (s)he receives a fallback utility:* $u_{FP} \geq 0$ *for the proposer or* $u_{FR} \geq 0$ *for the responder.*

A3. *If the proposer* (*resp., responder*) *accepts an outside alternative found in the search process, then that agent receives utility* $u_{AP} > u_{FP}$ (*resp.,* $u_{AR} > u_{FR}$).

A4. *No recall: alternatives found through search require an immediate decision; once rejected, they cannot be recalled at a later time.*

We depart from standard search models by allowing the agents to make a deal with each other, not requiring them to rely exclusively on outside alternatives. At any time during the search, the proposer has the option of making an offer to the responder and setting the time at which the offer expires. We denote the offer time by t_O and its expiration time (deadline) by t_D; Table 2 summarizes the basic notation.

Assumptions about the proposer's offer

A5. *The proposer's offer is a commitment: it remains valid until its deadline* (t_D) *and cannot be withdrawn before that time.*

A6. *The offer is an ultimatum (i.e., nonnegotiable): if the responder rejects the proposer's offer, then no mutual deal can be made.*

A7. *If the offer is accepted, then both agents receive as payoff their respective utilities of making a mutual deal* ($u_{MP} > u_{FP}$ *for the proposer and* $u_{MR} > u_{FR}$ *for the responder*).

[6] Both A1 and A3 are relaxed in §4, where we allow for nonstationary arrival rates and for the value of outside alternatives to be a random draw from a nonstationary distribution; A4 is relaxed in §4.6, where we address search with recall. We relax A6 in §3, which examines nonultimatum offers; A7 is relaxed in §4.5, where we consider the offer's value as an additional decision variable.

[7] The assumption that both agents share the same horizon T is not constraining. One agent having a longer search horizon ($\overline{T} > T$) is equivalent to having a search horizon T and receiving the expected value of the search on $[T, \overline{T}]$ as a fallback value (as defined in A2) when this horizon is reached.

Table 2 Basic model notation

Notation	Definition
$\lambda_P > 0$ ($\lambda_R > 0$)	Arrival rate of proposer's (responder's) search process
$u_{FP} \geq 0$ ($u_{FR} \geq 0$)	Proposer's (responder's) fallback utility
$u_{AP} \geq u_{FP}$ ($u_{AR} \geq u_{FR}$)	Proposer's (responder's) utility from the outside alternatives
$u_{MP} \geq u_{FP}$ ($u_{MR} \geq u_{FR}$)	Proposer's (responder's) utility from making a mutual deal
$T > 0$	Horizon of the search process
$t_O \in [0, T]$	Time at which proposer makes her offer to responder
$t_D \in [t_O, T]$	Time at which proposer's offer expires

Let us start by considering how the responder will react to the proposer's offer. We can find the dominant strategy for handling such offers by following the logic of Tang et al. (2009). The responder has nothing to gain by accepting an offer early and will therefore hold it until its deadline, at which point he must either accept the offer (and thus receive u_{MR}) or reject the offer and continue his search until the horizon T. If he continues searching, he will find an outside alternative with probability $1 - e^{-\lambda_R(T-t_D)}$, in which case he receives that alternative's utility u_{AR}; otherwise, he ends up with the lower fallback utility u_{FR}. Hence the responder will accept the proposer's offer if and only if

$$u_{MR} \geq (1 - e^{-\lambda_R(T-t_D)})(u_{AR} - u_{FR}) + u_{FR}. \tag{1}$$

Solving this inequality for t_D yields a single parameter that captures how the responder reacts to the proposer's offer – namely, the responder's *shortest acceptable deadline* (SAD_R). This follows because inequality (1) holds if and only if $t_D \geq \text{SAD}_R$, where SAD_R is given by

$$\text{SAD}_R = \begin{cases} 0 & \text{if } u_{MR} \geq u_{AR}, \\ \left(T - \frac{1}{\lambda_R} \ln\left(\frac{u_{AR} - u_{FR}}{u_{AR} - u_{MR}}\right)\right)^+ & \text{if } u_{MR} < u_{AR}. \end{cases} \tag{2}$$

Thus the responder will reject any proposer's offer that expires before this time ($t_D < \text{SAD}_R$); otherwise (i.e., when $t_D \geq \text{SAD}_R$) he will accept that offer (unless he has already accepted an outside alternative).

We proceed by separating the game into four possible cases that differ according to which of the two agents prefers a mutual deal and which prefers an outside alternative. The case when both agents prefer a mutual deal ($u_{AP} \leq u_{MP}$, $u_{AR} \leq u_{MR}$) is trivial in this setting because all equilibria then result in a mutual deal. We therefore discuss each of the other three cases in §2.1–§2.3; in §2.4, Corollary 1 ties the various cases together in a concise solution.

2.1. Both Agents Prefer an Outside Alternative

The simplest version of our model that still captures the fundamental difficulties involved in this decision is one in which both agents prefer an outside alternative to a mutual deal ($u_{AP} > u_{MP}$ and

$u_{AR} > u_{MR}$). In this case, one of the key decisions is simple: agents who find an outside alternative will accept it, because doing so maximizes their payoffs from this game. Thus our problem is reduced to a degenerate game in which only the proposer has a meaningful decision to make. She must select a time to make the offer (t_O) in conjunction with setting a deadline for that offer (t_D) so as to maximize her expected utility. The proposer would like to maximize her search time – and thereby increase her chances of receiving u_{AP} – before making an offer to the responder. Yet the longer she waits, the more likely it becomes that the responder will reject that offer because he has already accepted an outside alternative. The ideal time for the proposer to "seal the deal" (t_{IP}) balances these two objectives and can be expressed formally as[8]

$$t_{IP} = \begin{cases} 0 & \text{if } u_{AP} \leq u_{MP}, \\ \min\left\{T, \left(T - \frac{1}{\lambda_P}\ln\left(\frac{u_{AP}-u_{FP}}{u_{AP}-u_{MP}}\right) + \frac{1}{\lambda_P}\ln\left(\frac{\lambda_P+\lambda_R}{\lambda_R}\right)\right)^+\right\} & \text{if } u_{AP} > u_{MP}. \end{cases} \tag{3}$$

Proposition 1 *If both agents prefer outside alternatives (i.e., if $u_{AP} > u_{MP}$ and $u_{AR} > u_{MR}$), then the equilibrium is unique. The proposer makes an exploding offer ($t_O = t_D$) at time $t_O = \max\{t_{IP}, \mathrm{SAD}_R\}$, where SAD_R is as given by (2) and t_{IP} by (3); the responder accepts this offer provided he is still available.*

We shall continue referring to offers that expire on the spot ($t_O = t_D$) as *exploding* offers, and we refer to those with longer deadlines as *deadline* offers. Note that this convention differs slightly from the literature, where "exploding" is often used in reference to any kind of ultimatum offer with an extremely short deadline.

2.2. The Proposer Prefers a Mutual Deal but the Responder Does Not

If the proposer prefers a mutual deal but the responder prefers an outside alternative ($u_{AP} \leq u_{MP}$ and $u_{AR} > u_{MR}$), then the decision about when to make an offer is simpler. The proposer will maximize her utility by making an exploding offer at the first possible time at which the responder will accept it (SAD_R). Any offer that expires before this time will be rejected by the responder; any offer that expires afterward will simply give the responder more time to search, thereby reducing the offer's probability of being accepted but yielding no benefit to the proposer.

Although the offer's timing is now straightforward, the proposer has an additional decision to make. Each time the proposer finds an outside alternative before SAD_R, she must decide whether or not to accept it. Such decisions require that the proposer compare the value of the outside alternative to the value of her remaining search – which includes the possibility of making a deal

[8] When $u_{AP} > u_{MP}$, the expression for t_{IP} follows from Proposition 1; see Corollary 1 for the case $u_{AP} \leq u_{MP}$.

with the responder. Denote by $B_P(t)$ the expected payoff to the proposer if she rejects all outside alternatives before $t \geq \mathrm{SAD}_R$ and makes an exploding offer at time t:

$$B_P(t) = e^{-\lambda_R t}(u_{MP} - u_{FP}) + (1 - e^{-\lambda_R t})(1 - e^{-\lambda_P (T-t)})(u_{AP} - u_{FP}) + u_{FP}. \tag{4}$$

A proposer who has not accepted an alternative by time SAD_R can expect to receive $B_P(\mathrm{SAD}_R)$. Hence the optimal policy is to accept any outside alternative for which $u_{AP} > B_P(\mathrm{SAD}_R)$ and to reject when $u_{AP} < B_P(\mathrm{SAD}_R)$; if $u_{AP} = B_P(\mathrm{SAD}_R)$, then the proposer is indifferent between all possible policies for handling outside alternatives. These considerations lead to Proposition 2, which characterizes all the equilibria of this game.

Proposition 2 *If the proposer prefers a mutual deal but the responder does not (i.e., if $u_{AP} \leq u_{MP}$ and $u_{AR} > u_{MR}$), then – in any equilibrium of the game – the proposer makes an offer that expires at $t_D = \mathrm{SAD}_R$.*

(i) *If $u_{AP} > B_P(\mathrm{SAD}_R)$ for B_P given by (4), then the equilibrium is unique. The proposer's strategy is to accept any outside alternative and to make an exploding offer at time $t_O = \mathrm{SAD}_R$.*

(ii) *If $u_{AP} < B_P(\mathrm{SAD}_R)$, then there are multiple equilibria; the proposer rejects all outside alternatives before SAD_R but is indifferent between all $t_O \in [0, \mathrm{SAD}_R]$.*

(iii) *If $u_{AP} = B_P(\mathrm{SAD}_R)$, then any strategy profile in which $t_D = \mathrm{SAD}_R$ is an equilibrium.*

We note that the solution need not be unique. The reason is that, for specific values of model parameters, the proposer is indifferent between several available options. Because these equilibria yield no qualitative insight, it may be desirable to resolve these indifferences – and thereby guarantee the solution's uniqueness – by imposing two weak assumptions as described next.

Remark 1. Let the following optional assumptions hold.

A8. *If an agent derives exactly the same expected utility from search as from closing a deal with the other agent, then (s)he prefers to make the deal.*

A9. *If the proposer derives the same expectation from two different deadlines, then she prefers to set the shorter one.*

Given these assumptions, the equilibrium described by Proposition 2 is unique because A9 guarantees that $t_O = \mathrm{SAD}_R$ when $u_{AP} \leq B_P(\mathrm{SAD}_R)$ and A8 guarantees that the proposer rejects outside alternatives when $u_{AP} = B_P(\mathrm{SAD}_R)$.

2.3. The Responder Prefers a Mutual Deal but the Proposer Does Not

If only the responder prefers a mutual deal ($u_{AP} > u_{MP}$ and $u_{AR} \leq u_{MR}$), then the proposer still faces the same dilemma: when to make the offer (t_O) and when it should expire (t_D). From the responder's perspective, however, he will accept any offer from the proposer because doing so will maximize his payoff. In this case, the responder's additional decision is how to handle the outside alternatives that yield utility u_{AR}. Whenever the responder finds an outside alternative, he faces a choice: he can accept it, receiving its utility for sure but forsaking chances of receiving u_{MR}, or he can reject it in hopes of receiving an offer with value u_{MR}; however, the latter choice exposes the responder to the risk of not finding anything before the end of search horizon. The result is a coordination game, where the responder tries to predict when the proposer will make him an offer and the proposer tries to predict how long the responder will wait before he accepts an outside alternative. As is typical of coordination games, this leads to multiple equilibria. Let

$$B_R(t) = e^{-\lambda_P t}(u_{MR} - u_{FR}) + (1 - e^{-\lambda_P t})(1 - e^{-\lambda_R (T-t)})(u_{AR} - u_{FR}) + u_{FR}. \tag{5}$$

Much like $B_P(t)$ in Proposition 2, $B_R(t)$ is the expected payoff to the responder if he rejects outside alternatives before time t, accepts them after time t, while the proposer makes the responder an offer at time t if she is still available. Then the solution is given by Proposition 3.

Proposition 3 *If the responder prefers a mutual deal but the proposer does not (i.e., if $u_{AP} > u_{MP}$ and $u_{AR} \le u_{MR}$), then the following statements hold.*

(i) *If $u_{AR} > B_R(t_{IP})$ for t_{IP} given by* (3) *and for B_R given by* (5), *then the equilibrium is unique. The proposer's strategy is to make an exploding offer at time t_{IP} while the responder's strategy is to accept any outside alternative.*

(ii) *If $u_{AR} < B_R(t_{IP})$, then define the responder's* longest waiting time *as*

$$\text{LWT}_R = \max\{t \in [0,T] \mid u_{AR} \le B_R(t)\}. \tag{6}$$

Possible equilibria are the following: the proposer makes an offer at time $t_O \in [t_{IP}, \text{LWT}_R]$ and sets a deadline $t_D \in [t_O, T]$ while the responder rejects outside alternatives before t_O and accepts them afterwards.

(iii) *If $u_{AR} = B_R(t_{IP})$, then the set of equilibria consists of those given in* (ii) *as well as any strategy profile in which the proposer makes an exploding offer at* LWT_R.

(iv) *If A8 and A9 hold, then there are no mixed equilibria in this game and the only pure equilibria are those that involve exploding offers.*

Let us elaborate on this solution. The new concept introduced in this proposition is LWT_R, the responder's longest waiting time. Since the responder prefers making a deal with the proposer to his outside alternatives, he is willing to wait up to LWT_R to see if an offer is forthcoming from the proposer, in which case he does not accept any alternatives before that. If $t_{IP} > \text{LWT}_R$, however, then it is in the proposer's best interest to make her offer later irrespective of how the responder handles his search; hence the responder would be better-off not waiting at all, as stipulated in Proposition 3(i). Yet, if $t_{IP} \leq \mathrm{LWT}_R$, then the responder is better-off waiting – but only until t_O, from which follow the coordination game equilibria described in Proposition 3(ii). In this case, the proposer's decision about what deadline to set is moot because the responder is certain to accept her offer. Proposition 3(iii) covers the boundary case in which the responder is indifferent about the handling of his outside alternatives.

2.4. Conclusions from the Basic Model

We are now in a position to summarize the basic model's solution. There are many nuances in possible equilibria across the three cases, arising from different strategies that need to be used to handle outside alternatives as well as possible indifferences between the decisions. Nevertheless, there is one equilibrium that persists across all cases – a finding we formalize in the corollary that follows.

Corollary 1 *In all cases of the basic model, there exists an equlibrium in which the proposer makes an exploding offer at* $\max\{\mathrm{SAD}_R, t_{IP}, \mathrm{LWT}_R\}$; *here* SAD_R, t_{IP}, *and* LWT_R *are given (respectively) by* (2), (3), *and* (6).

Proof. This statement is a direct consequence of Propositions 1–3. □

This equilibrium is a useful focus for discussion because it captures, in a compact way, how the trade-offs involved change over different cases of the model. The proposer's best-case scenario is to close a deal at t_{IP}. Yet if this time is before the responder's shortest acceptable deadline ($t_{IP} < \mathrm{SAD}_R$), then a deal cannot be closed at t_{IP} because the responder would reject it; in that case, the proposer must wait until SAD_R. If the responder is willing to wait even longer before accepting any outside alternative ($\mathrm{LWT}_R > t_{IP}$), then the proposer can make an offer even later (i.e., at LWT_R) because doing so enables continued benefits from her own search without the risk of "losing" the responder.

One result that stands out in all cases is the lack of deadlines. The fundamental reason why deadline offers are dominated is that tweaking the deadline does not alter the responder's decision. On the one hand, if the responder is searching for a better outside alternative (Propositions 1 and 2), then he will react in exactly the same way to an offer made at time t_O and expiring at t_D as he will to an exploding offer made and expiring at t_D. Either way, the responder will continue his search until t_D: accepting an outside alternative if one is found; and accepting the proposer's offer at t_D only if (a) his search has yielded no viable alternative and (b) $t_D \geq \mathrm{SAD}_R$. On the other hand, a responder for whom the proposer's offer is the best possible outcome (Proposition 3) will always accept the proposer's offer – if he is still available – regardless of any deadline. For some values of model parameters, the proposer could be indifferent between different possible deadlines; however, under no game parameters can the proposer be strictly better-off by making a deadline offer rather than an exploding offer.

The line of reasoning just proposed does not hold in a more general setting – that is, when the outside alternative's value is drawn from a distribution. In this case, a responder who has received a deadline offer will become more selective about his outside alternatives, and that dynamic incentivizes the proposer to give longer deadlines. Even so, our "no deadlines" result prevails (as demonstrated by the general model of §4).

3. Variants

3.1. The Ability to Make an Offer and Commit

Our model has so far considered the deal-making process among agents as proceeding in a structured yet asymmetric manner: only one side (the proposer) can make an offer, and the offer is given as an ultimatum. That setup does correspond to some real-life situations, but there are many settings in which the deal-making process is not so structured. The ability to make an offer does not necessarily rest with one side only and either side might not be able to make credible ultimatums, potentially resulting in a series of offers and counter offers.

In this section we consider several variants of the game that differ in terms of which agents can make an offer and which ones can make credible ultimatums. Our aim is to develop insight on how a deal's timing and the agents' expected payoffs change under these various circumstances. Tractability of the solutions is ensured by building these variants on the basic model presented in §2 and adopting assumptions A1–A9 (with the possible exception of A6, whereby all offers are ultimatums). For notational consistency we continue to index the agents via P and R; however, our "proposer" and "responder" nomenclature is more usefully interpreted as corresponding to "agent 1" and "agent 2".

A key aspect of our approach will be identifying which of these agents prefers to make a mutual deal earlier. Recall from §2 that the ideal time for the proposer to make a mutual deal is t_{IP}, as given by (3). Analogously, the responder's ideal time to make a deal is t_{IR}, as given by (14) in the Appendix. Therefore, the proposer will prefer to make a mutual deal earlier if and only if $t_{IP} < t_{IR}$, in which case she is referred to as the *more eager* agent.

As a baseline against which the variants can be compared, we use the case in which deal making is the least structured. Thus either agent can propose a deal and neither can commit to an ultimatum. In this variant of the game, each agent can make an offer and set a deadline. Since the offers are not ultimatums, the same agent can later make another offer if his first one is rejected. Agents can likewise reject offers and propose counteroffers at any time. In this case the game is symmetric because both agents have the same available strategies – even as the relevant parameters (arrival rate, value of outside alternatives, value derived from a mutual deal, fallback value) will likely differ by agent.

The results of these variant games are reported in Table 3 and are proved in Proposition 4. In the table, ↑ (resp., ↓) signifies that the agent benefits from (resp., is harmed by) the change in rules; "NC" stands for "no change" in the payoffs.

Table 3 Agents' payoffs compared to the baseline case

Rule set	Less eager agent	More eager agent
Both agents can make offers but no ultimatums (baseline)	NC	NC
Both agents can make offers and ultimatums	↓	↑
Only the less eager agent can make offers or ultimatums	NC	NC
Only the more eager agent can make offers or ultimatums	↓	↑
Only one agent can make offers; no ultimatums	NC	NC

In the original model, the proposer could put the responder in the inconvenient spot of having to choose between the offer and the remainder of his search. In the baseline variant, neither agent has that leverage. An agent who likes an offer but would prefer to search longer before taking it can simply reject the offer and make a counteroffer at a more convenient time.

Proposition 4 *If the more eager agent has the ability to make an offer and can make an ultimatum, the set of equilibria of the game is identical to the basic model game in which the* more *eager agent is the proposer. In all other cases, the set of equilibria is identical to the basic model game in which the* less *eager agent is the proposer. An agent derives higher payoff in equilibrium if (s)he is the one who acts as the proposer in basic model.*

"Eagerness" of the agents parallels the sociological notion of power in relationships. In his seminal work, Emerson (1962) claims that power in a relationship is mainly determined by the availability of outside alternatives, which is precisely the case here. The best times to make a deal (t_{IP} and t_{IR}) are nothing else than measures of how good are the respective agent's outside alternatives as compared with a mutual deal, and the final deal is made at the best time for the agent with respect to whom this measure is higher. It can be seen from (3) that t_{IP} is increasing both in the arrival rate of outside alternatives (λ_P) and in the utility derived from such alternatives (u_{AP}); at the same time, t_{IP} is decreasing both in the other agent's availability of outside alternatives (λ_R) and in the utility derived from a mutual deal (u_{MP}). The results for our baseline case and for most variants are in accord with Emerson's theory in that the less eager agent has more power in the relationship. In particular, any deal that occurs does so at a time of that agent's choice *unless* the more eager agent can credibly commit to an ultimatum.

The key insights of this section are that this game, in all its variants, can play out in only two ways and that each of those outcomes is identical to the basic model's game with one agent acting as the proposer. As a result, the basic model is sufficient to find equilibrium solutions for all the variants considered here. All information relevant for determining who will assume the role of proposer is contained in t_{IP} and t_{IR} – the agents' ideal times to close a mutual deal. In a setting in which no one can make an ultimatum, the mutual deal is consummated at a time preferred by the less eager agent – that is, irrespective of which party is empowered to make offers. The more eager agent could force an earlier deal only by making an ultimatum offer. It follows that less eager agents are indifferent about their own offer-making abilities yet would like to deny a more eager agent the ability to make ultimatums.

With this insight in hand, let us revisit our introductory example. Recall that Manchester United made an ultimatum offer to Real Madrid: "£35 million for David de Gea within 24 hours or we are walking away and any deal is off". The problem with this ultimatum is that United did not credibly commit to it. Suppose, for instance, that Real came back – after the 24-hour deadline had passed – and offered the same amount (or more); then it would still be in United's interest to make the deal. Thus, the actual situation in this example corresponds more closely to the "natural" baseline variant in which neither party can commit but either one can make a proposal. This is indeed what happened. After rejecting United's apparent ultimatum, Real came back on the very last day of the transfer window and made another offer for David de Gea. Despite earlier claims that any such deal would be off the table, United agreed to the new offer.[9]

3.2. Minimum Deadline

A fundamental issue with any kind of "exploding offer" prescription is that instantaneous responses are simply not feasible. Even if the other side made has made a decision in advance, some time is needed to convey the response and/or process the documentation. In practice, then, all offers incorporate some minimum deadline. The United–Real story is a good example of what can happen if these procedural realities are not taken into account. We stated that, at the very end of transfer window, Real came back with an offer to which United agreed. Yet because the deal was concluded so late, there was not enough time to formalize the terms and hence the final documentation arrived at FIFA's offices two minutes *after* the transfer window had closed. As a consequence, the deal fell through in a spectacular way.[10]

[9] "David de Gea Finally Signs for Real Madrid," The Sport Bible, 31 August 2015, http://bit.ly/DeGeaAcc.

[10] "David de Gea's Real Madrid Transfer Collapsed at the Last Minute . . . ," Daily Mail, 1 September 2015, http://dailym.ai/1ljLoKG.

Proposition 5 gives the solution of the game described in §2 but with an added minimum deadline constraint. According to this result, the proposer will always find it optimal to make an offer with the minimum possible deadline.

Proposition 5 (Solution under a minimum deadline) *If there exists a minimum deadline $\Delta \leq T$ such that the proposer is constrained to choosing $t_D \in [t_O + \Delta, T]$, then all equilibria of this game have the proposer setting exactly the minimum deadline: $t_D - t_O = \Delta$.*

If $u_{AR} > u_{MR}$ or if $t_{IP} > \mathrm{LWT}_R + \Delta$ with t_{IP} and LWT_R given by (3) and (6) respectively, then the equilibrium is unique and the equilibrium offer expires at $t_D = \max\{\Delta, \mathrm{SAD}_R, t_{IP}\}$.

If $t_{IP} \leq \mathrm{LWT}_R + \Delta$ then, for all $t_D \in [\max\{t_{IP}, \Delta\}, \min\{\mathrm{LWT}_R + \Delta, T\}]$, the proposer's offer expiring at t_D is an equilibirum.

We thus have a simple prescription for when the responder prefers an outside alternative (as in Propositions 1 and 2). If there were no constraints on deadlines, then the proposer would make an offer to the responder at $\max\{t_{IP}, \mathrm{SAD}_R\}$ and have it expire on the spot. If there is a minimum required deadline Δ such that $t_{IP} \geq \Delta$, then the proposer should make her offer Δ *before* she otherwise would yet have it expire at the same time as in the previous case. If $t_{IP} < \Delta$, then the proposer should make an offer immediately and set the minimum possible deadline. In all cases, increasing the minimum deadline benefits the responder and harms the proposer.

Corollary 2 *In equilibrium, either imposing or increasing the minimum deadline Δ will reduce (resp., increase) the proposer's (resp., the responder's) expected utility.*

Although the complete proof of this corollary is tedious, the main intuition is simple: the proposer faces the same optimization problem but now with an additional constraint. If she wants to make the offer expire at the same time as it would without the restriction, then the offer must be made earlier – thus limiting her search. That limitation, in turn, reduces her expectation and makes it more likely that the responder will receive the proposer's offer.

This last result comes with a caveat. Our model assumes the proposer will – at some point – make an offer to the responder. In the context of the basic model, this assumption is a natural one because the proposer is always better off making an exploding offer at T (once her outside alternatives have been exhausted) than making no offer at all. However, that assumption may be unwarranted when there is a minimum deadline. Because the proposer can no longer make exploding offers, she might actually prefer to make no offer – which can occur when Δ is high. In the extreme case where $\Delta = T$, the proposer has but two options: make the responder an offer at $t_O = 0$ with deadline $t_D = T$, thus forfeiting all of her own search possibilities; or decline to make any offer and instead rely on outside alternatives.

4. General Model

The basic model provides the convenience of tractability by assuming that all outside alternatives yield the same utility. So that our conclusions will be as general and robust as possible, in this section we adopt one of the weakest sets of assumptions common to the search literature. Thus we allow for nonstationarity in arrival process and also – and more importantly – in the values of outside alternatives, which are drawn from a distribution with a continous cumulative distribution function (cdf). This means that our general model is based on most of the same assumptions as the basic model: A2 and A4–A7. However, we replace A1 and A3 with their general model counterparts, as follows.

General model assumptions

A1*. *Each agent's search is a nonhomogenous, finite-horizon Poisson process from* 0 *to* T *parameterized by the continuous arrival rate of alternatives:* $\lambda_P(t) > 0$ *for the proposer and* $\lambda_R(t) > 0$ *for the responder.*

A3*. *Once the proposer (resp. responder) finds an alternative in her (resp. his) search process, the value of that alternative is drawn from a continuous distribution with cdf* F_P *(resp.* F_R*), support* $[u_{FP}, \infty)$ *(resp.* $[u_{FR}, \infty)$*), and finite expectation.*

We will build up the model gradually. After revisiting standard search in §4.1, we consider how agents react to deadline offers in §4.2 and develop the resulting full model in §4.3.

4.1. Standard Finite-Horizon Search without Recall

Let us first consider the case of a single agent (the responder) searching for alternatives. This is the standard finite-horizon job search model (Lippman and McCall 1976, Mortensen 1986), which we use as a building block.

The agent has a finite search period $[0, T]$, during which he searches for alternatives that arrive according to a Poisson process with rate $\lambda(t)$. If an alternative is found, its utility (u_A) is drawn from a distribution with cdf F. The agent must then either accept the alternative, receiving its value, or reject it and continue the search. If he arrives at the horizon of his search without accepting an alternative, the agent receives fallback utility u_F.

In short, the agent accepts an alternative if and only if its value is greater than the expectation of the rest of his search under the optimal search policy – a monotone threshold policy captured by the agent's reservation price $\bar{\xi}(t)$. As per Van den Berg (1990, Thm. 1), this reservation price is given by the differential equation

$$\bar{\xi}'(t) = -\lambda(t) \int_{\bar{\xi}(t)}^{\infty} (x - \bar{\xi}(t))\, dF(x) \tag{7}$$

with border condition $\bar{\xi}(T) = u_F$. At time t, the agent accepts any alternative with value at least $\bar{\xi}(t)$ and rejects all others. This reservation price is decreasing with time: the responder will become increasingly less selective about alternatives as the end of his search approaches.

4.2. Reacting to Deadline Offers

Now we consider the situation of our responder as he follows the search process just described. As in the basic model, we use subscript R to indicate the responder's search parameters (thus $\lambda = \lambda_R$, $F = F_R$, and $u_F = u_{FR}$). We shall use $\bar{\xi}_R$ to denote the optimal single-agent search policy, as given by (7), for these parameters.

How does the responder react when given a deadline offer? An offer received at time t_O that expires at t_D will force him to make a decision at t_D: he must either accept the offer (and receive its value, u_{MR}) or reject the offer and continue with his search (which yields $\bar{\xi}_R(t_D)$ in expectation). Therefore, we can once again condense the responder's decision – on whether or not to accept an offer – into a single parameter: his shortest acceptable deadline. That time is now given by

$$\text{SAD}_R = \min\{t \in [0,T] \mid u_{MR} \geq \bar{\xi}_R(t)\}.$$

However, SAD_R is no longer sufficient to capture completely the responder's reaction to an offer. The additional factor now involved is how best to handle outside alternatives when an offer is already in hand. If the offer's value is high enough ($u_{MR} \geq \bar{\xi}_R(t_D)$) then the responder will never search all the way to the horizon: once he reaches time t_D, he will accept the deadline offer. As a result, his search horizon changes to t_D and his fallback utility to u_{MR}. Applying (7) then yields an expression for the responder's optimal policy when he holds this offer:

$$\bar{\xi}'_R(t; u_{MR}, t_D) = -\lambda_R(t) \int_{\bar{\xi}_R(t;u_{MR},t_D)}^{\infty} (x - \bar{\xi}_R(t; u_{MR}, t_D))\, dF_R(x) \tag{8}$$

with border condition $\bar{\xi}_R(t_D; u_{MR}, t_D) = \max\{u_{MR}, \bar{\xi}_R(t_D)\}$. It follows that holding an offer of sufficient value ($u_{MR} > \bar{\xi}_R(t_D)$) changes how the responder reacts to outside alternatives. More specifically, he will become more selective as his threshold for accepting outside alternatives increases, which we refer to as the *deterrence effect*.

The deterrence effect is an important feature of our model because it allows the proposer to influence how the responder handles his search for outside alternatives. She can induce the responder to become more selective by giving a longer deadline, thereby reaping a tangible benefit from making a deadline offer rather than an exploding one.

4.3. Complete General Model

We model the complete game as a strategic-form game in which the responder chooses an integrable function $\xi_R(t) : [0, T] \to \mathbb{R}$, which is his threshold policy for handling outside alternatives. Likewise, the proposer chooses $t_O \in [0, T]$ (the timing of her offer), $t_D \in [t_O, T]$ (that offer's expiration time), and an integrable function $\xi_P(t) : [0, T] \to \mathbb{R}$ (her own threshold policy for handling outside alternatives). If the proposer makes an offer that is rejected by the responder, then both agents use the single-agent policy given by (7) for the rest of their search because that is the sole optimal policy once a mutual deal is off the table. The payoffs in this game are the expected utilities when following the policies determined by the decision variables. Table 4 summarizes our notation.

Table 4 General model notation

Notation	Definition
$\lambda_P(t)$ ($\lambda_R(t)$)	Arrival rate of proposer's (responder's) search process
$u_{FP} \geq 0$ ($u_{FR} \geq 0$)	Proposer's (responder's) fallback utility
$F_P(x)$ ($F_R(x)$)	Distribution from which the value of the proposer's (responder's) outside alternatives is drawn
u_{MP} (u_{MR})	Proposer's (responder's) utility from making a mutual deal
$T > 0$	Horizon of the search process
$\xi_P(t)$ ($\xi_R(t)$)	Proposer's (responder's) policy for handling outside alternatives
$t_O \in [0, T]$	Time at which proposer makes her offer to responder
$t_D \in [t_O, T]$	Time at which proposer's offer expires
$\bar{\xi}_P(t)$ ($\bar{\xi}_R(t)$)	Proposer's (responder's) single-agent search policy as given by (7)
$\bar{\xi}_R(t; u_{MR}, t_D)$	Responder's optimal policy, as given by (8), when holding an offer with deadline t_D and value u_{MR}

This is a game of complete but imperfect information, since each agent remains unaware of the alternatives that the other agent has found. What makes the game challenging to analyze is that payoff functions rely on recursive expressions with no closed form. As an illustration of the difficulties involved: even though the Caylay–Moser problem studied in mathematics (Moser 1956, Guttman 1960) is a specific and simpler case of the recursions (7) and (8), it remains intractable. Because of the intractability due to recursive payoffs and of the need to optimize over a function space, this game's equilibria cannot be characterized analytically. Yet such a characterization is fortunately not needed for our primary goal in this section, which is to examine whether the basic model's qualitative conclusions – in particular, the no-deadlines result – hold in a more general setting. For that purpose we can rely instead on best-response functions.

We start by examining the responder's rationalizable strategies, which amount to best responses to particular beliefs about the proposer's future actions (Mas-Colell et al. 1995, p. 242). The crucial property of such strategies is that they incorporate a deterrence effect (as described by the following lemma).

Lemma 1 (Deterrence effect for rationalizable policies) *Let $\xi_R(t)$ be a rationalizable policy for the responder, and let $\bar{\xi}_R(t)$ be the single-agent search policy given by (7). Then, for all $t \in [0,T]$, we have $\xi_R(t) \geq \bar{\xi}_R(t)$.*

Proof. The expectation of the rest of the search from t onwards under an optimal policy is $\xi_R(t)$.[11] However, the responder always has the option of rejecting any proposer's offer and using the single-agent search policy given by (7) to handle outside alternatives – a strategy that yields $\bar{\xi}_R(t)$ in expectation. It then follows from the optimality of ξ_R that $\xi_R(t) \geq \bar{\xi}_R(t)$. □

A responder who believes that the proposer might make him an offer will be more selective when evaluating his outside alternatives than will a responder who holds no such belief. Theorem 1 uses this fact to show that the basic model's most salient feature holds also in the general model. Namely, the proposer is never better-off from making a deadline (rather than an exploding) offer.

Theorem 1 (No deadlines) *Let the responder use a rationalizable policy ξ_R. Then, for any times t_D^{**} and t_O^{**} such that $t_D^{**} > t_O^{**}$, there exists a t_O^* such that the proposer receives higher expected utility from making an exploding offer at t_O^* than from making an offer at t_O^{**} with deadline t_D^{**}.*

As seen in the previous section, setting a deadline benefits the proposer because it deters the responder from accepting some of his outside alternatives. Yet a deadline comes at a cost, since the proposer cannot then withdraw the offer and so cannot accept her own outside alternatives until the offer expires. As Theorem 1 shows, the benefits gained are never worth this cost.

There are several insights to be drawn from this theorem; of these, the most important is that, when the responder is employing any rationalizable policy ξ_R, the proposer's best response is always to make an exploding offer. As an extension, any equilibrium of this game must consist of the proposer making an exploding offer. Note that Theorem 1 does not rely on the responder having complete information. Suppose the responder has some *mistaken* beliefs about when the proposer might make him an offer. Provided the responder is attempting to maximize expected utility according to his beliefs, the policy he uses is a rationalizable one even if those beliefs are mistaken. Even for responders of this (misguided) type, the proposer's best course of action is still to make an exploding offer.

The foregoing analysis is independent of the proposer's search process, and it remains valid for a nonstationary distribution of the proposer's alternatives. For a proposer who plans to make an offer at t_O^{**} with deadline t_D^{**}, the proof of Theorem 1 prescribes how to improve on that offer by making it an exploding one:

[11] Lippman and McCall (1976) and Mortensen (1986) provide details about such policies.

- if t_D^{**} is such that proposer prefers that the responder reject the offer (i.e., if $\bar{\xi}_P(t_D^{**}) > u_{MP}$), then she is better-off making an exploding offer at T;
- otherwise, the proposer is better-off making an exploding offer at $\max\{t_O^{**}, \mathrm{SAD}_R\}$.

4.4. Nonstationary Distribution of Outside Alternatives

Finally, we consider the situation where the distribution of outside alternatives found by each agent changes over time. When an alternative is found, its value is drawn from a distribution with cdf $F(x;t)$, which is a continuous function of both time. As a function of x, $F(x;t)$ is also continuous, with finite mean and with support $[u_F, \infty)$, on which it is strictly increasing. In such a setting, the single-agent reservation price given by (7) becomes

$$\bar{\xi}'(t) = -\lambda(t) \int_{\bar{\xi}(t)}^{\infty} (x - \bar{\xi}(t))\, dF(x;t), \tag{9}$$

with the usual border condition $\xi(T) = u_F$. This type of nonstationarity is of special interest because it can violate the no-deadlines property stated in Theorem 1. Whether or not it remains optimal to make only exploding offers depends on how the responder's distribution of outside alternatives ($F_R(x;t)$) evolves over time. Properties of $F_P(x;t)$ do not affect this result, since the proof of Theorem 1 does not require that $F_P(x;t)$ be stationary.

If $F_R(x;t)$ evolves over time, but in such a way that it becomes worse in the sense of convex second-order stochastic dominance (CX-SSD), it is still optimal to issue only exploding offers. For any cdf F, let $\overline{F}(x) = 1 - F(x)$. By definition, a distribution with cdf G dominates the one with cdf H by CX-SSD if

$$\int_y^{\infty} \overline{G}(x)\, dx \geq \int_y^{\infty} \overline{H}(x)\, dx \quad \forall y \in \mathbb{R}, \tag{10}$$

with the dominance being strict if there exists at least one y for which this inequality is strict.[12]

Theorem 2 *Theorem 1 extends to the case of nonstationary distribution of outside alternatives if, for any $t_1, t_2 \in [0,T]$ such that $t_1 < t_2$, $F_R(x;t_1)$ dominates $F_R(x;t_2)$ by CX-SSD.*

Let us provide some intuition for why CX-SSD is the relevant ordering here. When an agent evaluates a distribution of outside alternatives, there are two relevant criteria: (i) the probability of finding an alternative that exceeds the reservation price; and (ii) the value of such alternative when found. In essence, these criteria capture how good is the "right tail" of a distribution. The "left tail" is not relevant because alternatives below the reservation price will be rejected. Both

[12] Convex second-order stochastic dominance is also known as an increasing convex order (Müller and Stoyan 2002, Shaked and Shanthikumar 2007), a consequence of the following equivalent definition. Let X and Y be random variables distributed according to the cdfs G and H, respectively. 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).

of these criteria are contained in the integral condition (10). Consequently, 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,[13] then the searcher will prefer the new to the old distribution causing reservation prices to rise accordingly. Our Lemma 3 (in the Appendix) generalizes this result to CX-SSD and is used in the proof of Theorem 2.

It is often reasonable to expect that the distribution of outside alternatives will become less desirable over time, in which case Theorem 2 holds. For example, Lippman and Mamer (2012) observe that in real estate markets there is initially strong interest in newly listed properties but that this interest wanes over time. Along these lines, Weitzman (1979) demonstrates that a strategic agent who can choose the order in which to draw from different distributions should pursue them in decreasing order of preference.

If the conditions of Theorem 2 are not fulfilled, then it might be possible for the proposer to benefit from making a deadline offer. This claim is illustrated by the following example.

Example 1A. Suppose the search horizon is divided into two periods. The *proposer* can find outside alternatives only in the first period, during which she finds one (yielding \$2) with probability 0.2. In this period, the *responder* can nearly always (with probability 0.9) find an outside alternative, which also yields \$2 if accepted. In the second period, the responder's outside alternatives have high value (\$10) but are found only with low probability (0.1). The proposer must to decide when to make her offer to the responder; she can do this before any search time has elapsed (i.e., at time 0), in between the two search periods (at time 1), or at the search horizon (time 2). The proposer can also set the deadline on her offer to any of these three times. The responder must decide how to handle his outside alternatives in the event no offer is forthcoming from the proposer (ξ_R); in this example, that decision boils down to whether or not the first-period alternative should be accepted. If the agents agree on a mutual deal then it will be worth \$1.5 to each. Any agent who reaches the search horizon (i.e., without accepting an alternative) receives the fallback value of \$0.

In the sole equilibrium of this game, the proposer makes an offer at time 0 that lasts until the search horizon (i.e., a time-2 deadline). Table 5 reports the game's payoffs for all the proposer's

[13] Inequality (10) is implied by G dominating H (in the sense of first-order stochastic dominance) and is implied also by G being a mean-preserving spread of H. Therefore, either of the following two conditions is sufficient for Theorem 2 to hold: (i) for all $x \in \mathbb{R}$ and $t_1, t_2 \in [0, T]$ such that $t_1 < t_2$, we have $F_R(x; t_1) \leq F_R(x; t_2)$; (ii) for all $t_1, t_2 \in [0, T]$ such that $t_1 < t_2$, we have that $F_R(x; t_1)$ is a mean-preserving spread of $F_R(x; t_2)$.

strategies when the responder uses the single-agent search policy; underlying calculations and proof of the equilibrium are given in the Appendix. □

Table 5 Payoffs for Example 1A

t_O	t_D	Proposer's payoff	Responder's payoff	Probability of mutual deal
0	0	\$0.4	\$1.9	0
1	1	\$0.52	\$1.94	0.08
2	2	\$0.508	\$2.008	0.072
0	1	\$0.15	\$1.95	0.1
0	2	\$1.35	\$2.35	0.9
1	2	\$0.508	\$2.008	0.072

In Example 1A, the proposer can use the deterrence effect of deadline offers to make the responder reject low-value offers (which he can easily find) in favor of chasing high-value (albeit unlikely) offers later on. This effect can be achieved only by making a deadline offer at time 0 with deadline at time 2. The reasoning behind this conclusion is that (a) deadline offers made at time 0 but expiring at time 1 do not have a sufficiently strong deterrence effect, and (b) an exploding offer made at time 0 will be rejected outright (because $\text{SAD}_R = 1$). This result is driven by the responder's chance to find high-value alternatives later in his search – ones that are unavailable early in the search. However, this scenario cannot occur under the conditions of Theorem 2.

We conclude this section by noting an insight derived from the proof of Theorem 2. Even if that theorem's conditions do not hold, a deadline offer that expires at SAD_R cannot be better (from the proposer's standpoint) than an exploding offer made at SAD_R; nor can a deadline offer made at $t_O \geq \text{SAD}_R$ be better than an exploding offer made at t_O. So when deadline offers are optimal, they are always made before SAD_R and expire after SAD_R (just as in Example 1A).

4.5. Adjusting the Offer's Value

Let us now consider another possibility. In addition to controlling when to make an offer and what deadline to set, the proposer may also be able to adjust the offer's value. We can account for this possibility by introducing value of the offer x as an additional variable in the proposer's decision process. The utility derived from making a mutual deal then becomes a function of the offer value. The proposer receives $u_{MP}(x)$, which is a decreasing differentiable function, and the responder receives $u_{MR}(x)$, an increasing differentiable function. Hence SAD_R is also a function of the offer value: $\text{SAD}_R(x) = \min\{t \in [0,T] \mid u_{MR}(x) \geq \bar{\xi}_R(t)\}$. In order to guarantee that both agents can benefit from a mutual deal, we assume that there exists an x such that $u_{MP}(x) > u_{FP}$ and $u_{MR}(x) > u_{FR}$.

The framework developed here accounts for a wide variety of situations – including employment benefits, which yield high utility for the employee while costing the employer little – as well as different degrees of risk aversion among the agents. Under this framework the game becomes simpler, as the proposer's dominant strategy is to make only exploding offers and in such a way that the responder receives the same utility from accepting this offer as he does from continuing his search. That dynamic is formalized in the following remark.

Remark 2. Consider the setting of Theorem 2, but suppose that the proposer can adjust the offer's value as described previously. If the responder is using a rationalizable strategy, then it follows from Theorem 2 that, for any fixed offer value x, there exists a time $t_O^* \geq \mathrm{SAD}_R(x)$ such that the proposer's expectation is higher from making an exploding offer at t_O^* than from making any deadline offer. The proposer can do better still by lowering the offer's value to $x^* \leq x$ such that $t_O^* = \mathrm{SAD}(x^*)$, which is the lowest possible value acceptable to the responder.

There are several implications of this remark. First, the responder derives no utility from an offer of value x that expires at $\mathrm{SAD}_R(x)$, since he can get as much from his own search. That being said, any proposer's best response – and, by extension, every equilibrium of the game – consists exclusively of such offers. Thus in any equilibrium the responder simply maximizes the value of his own search (i.e., as if the proposer did not exist) by using the single-agent search policy (7).

One might well think an adjustable offer value would benefit the responder for at least some parameter values, since the proposer might increase that value as a means of closing the deal sooner (and thus with greater likelihood) than she otherwise could. Yet even though the proposer sometimes makes that adjustment, it never benefits the responder. The reason is that, because these higher-value offers are always made earlier in the search period, accepting one requires the responder to forgo a larger part of his search. At first glance, it would seem like the proposer can only benefit from being able to adjust the offer's value, since it gives her an another degree of freedom over which she can optimize. However, the next example demonstrates that this is not the case.

Example 2. Consider the setting of Proposition 3, with the following parameters. The responder has a fast arrival rate of outside alternatives ($\lambda_R = 3$), derives low utility from those alternatives ($u_{AR} = 5$), and benefits considerably from making a deal with the proposer ($u_{MR} = 15$). The proposer has a low arrival rate ($\lambda_P = 1$) and derives more utility from her outside alternatives ($u_{AP} = 20$) than from a mutual deal ($u_{MP} = 5$). The horizon is $T = 1$, and each agent's fallback value is zero ($u_{FR} = u_{FP} = 0$). Applying Proposition 3 yields this game's lone equilibrium, in which the proposer makes her offer at the search horizon (since $t_{IP} = T$) while the responder rejects all outside alternatives before then (since $\mathrm{LWT}_R = T$). The proposer's equilibrium payoff is then $(1 - e^{-\lambda_P T})(u_{AP} - u_{MP}) + u_{MP} \approx 14.48$.

Next we consider the same situation except that the proposer can alter the offer's value x so that she receives $u_{MP} = 20 - x$ from a mutual deal and the responder receives $u_{MR} = x$. The game in the preceding paragraph corresponds to x being fixed to 5. By Remark 2, the responder will accept all outside alternatives in any equilibrium of the game, whereas the proposer will make an offer at some time t_O and set its value x such that $\text{SAD}_R(x) = t_O$. Solving this equation for x yields $x(t_O) = 5 - 5e^{-3(T-t_O)}$. The game's equilibrium can then be found by maximizing the proposer's expected utility over possible offer times. Applying equation (11) (from the Appendix) shows that this expected utility is equal to $(1 - e^{-\lambda_P t})u_{AP} + e^{-\lambda_P t}e^{-\lambda_R t}(20 - x(t)) + e^{-\lambda_P t}(1 - e^{-\lambda_R t})(1 - e^{-\lambda_P (T-t)})u_{AP}$; this expression is maximized at $t = 1$, where it yields a payoff of 12.66 to the proposer. Thus, when the proposer can control the offer's value, the game's only equilibrium yields a lower payoff for *both* agents than when the offer's value is a fixed amount. □

The intuition behind Example 2 is straightforward. If the proposer cannot adjust the offer's value, then the responder is willing to reject outside alternatives in order to wait and see whether he receives the proposer's offer – knowing that the proposer's offer, if made, will exceed the value of his outside alternatives. If the proposer can adjust the offer's value, she derives no benefit from setting it higher than the minimum needed to induce the responder's acceptance. Hence the responder no longer waits for the proposer, which harms the proposer more than being able to adjust the value of the offer benefits her.

Lastly, let us comment on the validity of no-deadlines result. As shown in Remark 2, no-deadlines holds under conditions of Theorem 1 or, more generally, Theorem 2. If the conditions are different, that is, if the distribution of the responder's outside alternatives $F(x;t)$ is nonstationary but in a way that is not captured by Theorem 2, then the proposer may be better served by leaving a longer deadline. This possibility is demonstrated by our last example.

Example 1B. Consider the game described in Example 1A, but now suppose that the proposer can choose the offer's value – a dollar amount x. If accepted, the offer pays $u_{MR}(x) = x$ to the responder and $u_{MP}(x) = \$3 - x$ to the proposer.[14] In this case, the game's only equilibrium will have the proposer making an offer at time $t_O = 0$ with deadline $t_D = 2$ and value $x = \$10/9$. Table 6 summarizes expected payoffs from the proposer's different strategies when the responder uses the single-agent search policy; calculations and a proof of the equilibrium are given in the Appendix. □

Let us examine why this is the case. A proposer who makes an exploding offer should set its value equal to the expected value of the rest of the responder's search under an optimal policy. Any offer amounting to less than that expected value will be rejected, and any offer of higher value

[14] Example 1A corresponds to x being fixed at $x = \$1.5$.

Table 6 Payoffs for Example 1B

t_O	t_D	x	Proposer's payoff	Responder's payoff	Probability of mutual deal
0	0	\$1.9	\$1.1	\$1.9	1
1	1	\$1	\$0.56	\$1.9	0.08
2	2	\$0	\$0.616	\$1.9	0.072
0	1	\$2	\$1	\$2	1
0	2	\$10/9	\$1.7	\$2	0.9
1	2	\$0	\$0.616	\$1.9	0.072

will be accepted but will give the proposer less utility. Table 6 reveals that the best option for a proposer who favors an exploding offer is to make a high-value offer early in order to prevent the responder from searching for outside alternatives. However, such an offer yields \$1.1 to the proposer, which is less than her payoff from making the same deadline offer as in Example 1A (at time 0, with deadline 2, of value \$1.5, and yielding \$1.35). Yet the proposer can improve upon this outcome by changing the offer's value x. The deadline offer in Example 1A was successful because it deterred the responder from accepting period-1 alternatives, which he finds with high probability. The proposer can lower the value of the offer to the minimum that still achieves this effect, which in this case is $x = \$10/9$. Observe that only the deadline offers are beneficial to the responder. With exploding offers, the responder never receives more than the expectation of his search for alternatives (in line with Remark 2).

4.6. Search with Recall

Until now we have assumed that the search is without recall: all outside alternatives found are exploding ones. It is worthwhile to consider what happens when this is not the case.

We shall use the setting of Lippman and Mamer (2012) as the basis for our discussion. Those authors consider a situation in which the responder is undertaking search *with recall* – that is, all outside alternatives that he finds remain viable until the search horizon – and there is a single proposer who seeks to maximize the probability of the responder accepting her offer. The proposer is faced with a choice: Does she make an exploding offer, or does she make an offer that is valid until the horizon is reached? Lippman and Mamer successfully illustrate the great difficulty of making this decision, which can change multiple times as the search nears its horizon and depends on many factors (including the shapes of the distributions from which alternatives are drawn). In that setting, the responder need not strategize his decisions about outside alternatives; he simply collects them and, when the search horizon is reached, accepts the best one.

If the proposer makes an offer at time t_O that expires at t_D, then the responder will hold the offer until time t_D – when he will accept it provided the utility from doing so exceeds the expected value of his search from t_D onward. The responder will therefore treat any offer that expires at time t_D

exactly the same (i.e., regardless of when it was made). Unlike the no-recall setting, which featured a deterrence effect, in this scenario there is absolutely no benefit to setting a deadline because the responder is not required to decide about any outside alternatives prior to T. Consequently, if a proposer can control the timing of her offer, rather than it being exogenous, then our no-deadlines result holds in this setting as well.

5. Summary and Discussion

Once a recruiting party has identified a suitable candidate, the next step is to make an offer. The offer's value (e.g., the salary) may or may not be a decision variable. For instance, salaries on the academic job market for rookies are often fixed (or amenable to only slight adjustment); hence the offer's attractiveness is mostly exogenous and largely determined by the candidate's own perception of the school. Additional decision variables include the deadline (the time past which the offer is no longer valid) and the offer's timing (it need not be made as soon as a suitable candidate is identified). The proposer who hopes to find other alternatives therefore faces a difficult trade-off: making an offer earlier (for a given expiration date) makes it more likely that the responder will accept, but such offers reduce the amount of time available for the proposer to search. And even though setting a later deadline deters the responder from accepting some of his outside alternatives, that strategy also gives the responder more time to search. Finally, an early offer with too short a deadline might lead to immediate rejection. Our paper addresses these difficult yet ubiquitous decisions: the timing of an offer and its deadline.

Our most general model is in §4, where both the proposer and the responder are searching for outside alternatives, which come from a process with a nonstationary arrival rate and have a nonstationary distribution of the payoffs. The main result is that it is not optimal to leave any deadline, if the proposer is free to choose the timing of the exploding offer. More precisely, Theorems 1 and 2 show that any deadline offer is dominated by an offer expiring on the spot – provided the former's timing is appropriately adjusted. Compared to an optimal exploding offer, a longer deadline will not improve the chances of responder acceptance; however, it will leave the proposer with less time to search. Regarding nonstationary distribution of the outside alternatives, Theorem 2 shows that this result holds as long as this distribution (of the responder's outside search) gets worse over time in the sense of the second-order convex stochastic dominance. Example 1A shows that this assumption is cruical for our no-deadlines result, for if the distribution improves over time then the proposer may be better-off setting a longer deadline.

Other than this lone exception, the no-deadlines property is quite robust: it holds when the proposer is allowed to change the offer's value (§4.5 and Remark 2) and when the responder's outside search is with full recall (§4.6). An interesting practical implication of the no-deadlines result is that a job candidate will be hard pressed to justify demanding that the recruiter's offer incorporate a substantially delayed deadline. Many of us on the recruiting side (e.g., of the academic job market) have faced the situation of a candidate calling to rush us for an offer because a competing offer is about to expire. An appropriate response to such a call is either no offer or an offer with a short deadline, since the candidate evidently prefers our offer (or would not have bothered to call). A job candidate could more convincingly justify an extended deadline by acknowledging that he prefers a competing offer to ours (and so our short-deadline offer would be rejected) and then agreeing, if only we extend our offer's deadline, to reject that competing offer in hopes of finding something better (while using our offer as a fallback option) – exactly as in Examples 1A and 1B.

A simplified version of our general model is considered in §2, where we assume that all outside alternatives have exactly the same value and arrive at a constant rate (assumptions A1 and A3). We can then solve for an equilibrium explicitly and, in §3, proceed to studying comparative statics: what happens if both agents are able to make offers, when is it beneficial to have an opportunity (reputation) to credibly commit to an ultimatum offer, and related issues (Table 3 and Proposition 4). In most of these variants we find that the mutual deal is concluded at the time preferred by the less eager agent; the deal happens earlier only in variants where the more eager agent can make an ultimatum offer.

In §3.2 we examine how the conclusions change if there is a constraint on the minimal deadline. Business schools often impose such constraints on firms that engage in on-campus recruitment, restricting how soon the offers made to new MBA graduates can expire.[15] If this is the case, deals still get made at the same time but the offers might need to be extended earlier in order to accomplish this (Proposition 5). In general, a longer minimum deadline benefits the responder but not the proposer; however, a minimum deadline that is too long could be detrimental for both agents because it may result in the proposer declining to make any offer.

Even when market rules do not mandate a minimum deadline, the proposer must take into account that short deadlines (and, in the limit, exploding offers) might be perceived as bullying and unfair and thus be rejected for these reasons alone (and not because the responder actually expects to find a better alternative). It is also worth remarking that, after an exploding offer has been accepted, the offer's nature may be perceived negatively by the responder and lead to negative reciprocation (Lau et al. 2014). For this reason we would advise setting not the minimal possible deadline but rather the minimal *socially acceptable* deadline; thus the proposer's attention should

[15] "Offer Policy, University of Chicago Booth School of Business," http://bit.ly/Boothop; "Recruiting Policies," Yale School of Management, http://bit.ly/YaleSOMrp; "Recruiting Policies, INSEAD Career Development Centre," http://bit.ly/INSEADcrp.

be focused on the optimal time to make an offer while realizing that the objective of the deadline is to make the responder feel well treated.

To the best of our knowledge, this paper is the first one to study the optimal timing of an ultimatum offer. Simplified versions of our setup might be amenable to experiments. We anticipate that participants will likely not set the timing of offers and deadlines in an optimal way – since this decision involves a delicate balance between their own search opportunities and the odds of responder acceptance – and that the beginning and end of the search horizon may act as strong anchors. It would be interesting to look for systematic deviations in the participants' choices from the optimal prescription. Future research could also extend our theoretical model in several ways: there could be uncertainty about agent types, learning about the search process, and/or multiple agents (several responders and proposers). A multi-agent model in which agents still have access to exogenous search could provide insights about functioning of decentralised matching markets. Such research may well require a creative modeling approach to reduce the corresponding dynamic programs' high number of state variables.

As indicated by our paper's title, the prescriptive advice from our analysis is to be "patient yet firm": *patient* in the sense of thinking hard about the timing of the offer, typically delaying it and not pulling the trigger immediately; *firm* in the sense that, once an offer is issued, it should have a deadline no longer than the minimum one allowed by social norms. Looking back at the Manchester United and Real Madrid story, United was in line with our "firm" advice when making an offer with a short deadline; however, it failed the "patient" criterion by making the offer too early in the transfer window. It turned out that Real rejected the initial offer because of its timing – and not because of its high price – as the team later agreed to very similar terms.[16] It is difficult remaining patient enough to make an optimally timed offer when the stakes are high, so we hope that our research provides a useful way of thinking about such decisions.

Appendix

Proof of Proposition 1. The responder will accept the proposer's offer at time t_D only if $t_D \geq \mathrm{SAD}_R$. Therefore, it is optimal for the proposer to set $t_O = t_D \geq \mathrm{SAD}_R$. The proposer will search for outside alternatives from time 0 to t_O, finding one with probability $1 - e^{-\lambda_P t_O}$, in which case she will accept it and receive u_{AP}. If she does not find any, she will make an offer to the responder at time t_O. The responder will accept it if he did not find any outside alternative by t_O (probability of that is $e^{-\lambda_R t_O}$). If the offer is accepted, the proposer will receive u_{MP}. If it is rejected, she will continue searching for outside alternatives

[16] The ultimatum issued by United on 13 July 2015 demanded £35 million, whereas the deal finalized on 31 August 2015 was for £29 million and included also the transfer of Keylor Navas to United from Real: "Paper Round: Real Madrid Handed 24-hour Ultimatum over David de Gea," Eurosport, 13 July 2015, http://bit.ly/1TafifR; "David de Gea Finally Signs for Real Madrid," The Sport Bible, 31 August 2015, http://bit.ly/DeGeaAcc.

until T and find one with probability $(1-e^{-\lambda_P(T-t_O)})$. If she reaches the end of her search horizon without finding one, she will receive the fallback utility u_{FP}. Therefore, the proposer will receive utility u_{AP} with probability $1-e^{-\lambda_P t_O}+(1-e^{-\lambda_R t_O})(e^{-\lambda_P t_O}-e^{-\lambda_P T})=1-e^{-(\lambda_P+\lambda_R)t_O}-(1-e^{-\lambda_R t_O})e^{-\lambda_P T}$, receive u_{MP} with probability $e^{-(\lambda_P+\lambda_R)t_O}$, and receive u_{FP} otherwise. The proposer's expected payoff is then

$$\pi_P(t_O)=\left(1-\left(1-e^{-\lambda_R t_O}\right)e^{-\lambda_P T}\right)(u_{AP}-u_{FP})-e^{-(\lambda_P+\lambda_R)t_O}(u_{AP}-u_{MP})+u_{FP}. \tag{11}$$

We first solve the unconstrained maximisation problem $\max_{t_O\in\mathbb{R}}\pi_P(t_O)$. Taking the derivative of π_P yields

$$\pi_P'(t_O)=e^{-T\lambda_P-t_O\lambda_R}\left((\lambda_R+\lambda_P)e^{\lambda_P(T-t_O)}(u_{AP}-u_{MP})-\lambda_R(u_{AP}-u_{FP})\right).$$

From the equation above, we can see that the first derivative will cross over 0 at most once, as $e^{-T\lambda_P-t_O\lambda_R}>0$ and the expression in brackets is a decreasing function function of t_O. Since $\lim_{t_O\to-\infty}\pi_P'(t_O)=\infty$ and $\lim_{t_O\to\infty}\pi_P'(t_O)=0^-$, $\pi_P'(t_O)$ crosses over 0 exactly once and the unconstrained function $\pi_P(t_O)$ is unimodal. Hence, FOC is sufficient and solving for it yields

$$\arg\max_{t_O\in\mathbb{R}}\pi_P(t_O)=\left\{T-\frac{1}{\lambda_P}\ln\left(\frac{u_{AP}-u_{FP}}{u_{AP}-u_{MP}}\right)+\frac{1}{\lambda_P}\ln\left(\frac{\lambda_P+\lambda_R}{\lambda_R}\right)\right\}. \tag{12}$$

Taking the constraint $0\le t_O\le T$ into account we are left with $\arg\max_{t_O\in[0,T]}\pi_P(t_O)=\{t_{IP}\}$, with t_{IP} given by (3). The optimal time to make an offer is then $\max\{t_{IP},\mathrm{SAD}_R\}$. □

Proof of Proposition 3. Analogous to Proposition 2, the responder has just three possible strategies for handling outside alternatives that can plausibly be best responses. He either rejects all alternatives until t_O (best response when $u_{AR}<B_R(t_O)$), accepts all of them (when $u_{AR}>B_R(t_O)$) or is indifferent about what to do (when $u_{AR}=B_R(t_O)$). In all of the cases, the responder accepts all outside alternatives after t_O unless he received the proposer's offer.

First, consider the case when the responder accepts all alternatives. Then, the proposer is in the same situation as in Proposition 1, and thus makes an exploding offer at time t_{IP}. If $u_{AR}\ge B_R(t_{IP})$, the responder has no incentive to deviate so this is an equilibrium. If $u_{AR}<B_R(t_{IP})$ there are no equilibria in which the responder accepts all alternatives.

Next, consider the case where the responder rejects all alternatives until t_O. For this to be an equilibrium, it is necessary that this is really the responder's best response, i.e., $u_{AR}\le B_R(t_O)$. This in turn implies $t_O\le \mathrm{LWT}_R$ since $B_R(t)$ is decreasing. The proposer cannot increase her expectation by making an offer earlier, since it will not increase the chances that the responder accepts the offer, but will decrease the chances of getting an outside alternative with $u_{AP}>u_{MP}$. She can increase her expectation by making an offer later if and only if $t_O<t_{IP}$. This stems from the fact that the proposer's expectation as a function of her offer time is increasing up to t_{IP} and decreasing afterwards when the responder is accepting outside alternatives, as shown in the proof of Proposition 1. Consequently, for each t_O such that $t_{IP}\le t_O\le \mathrm{LWT}_R$, the proposer making an offer at t_O and the responder rejecting all outside alternatives before t_O is an equilibrium, irrespective of the deadline. There are no other equilibria in this case, as for $t_O<t_{IP}$ the proposer has an incentive to deviate and for $t_O>\mathrm{LWT}_R$ the responder has.

Lastly, consider the case where the responder is indifferent about what to do with outside alternatives. This is truly the responder's best response only when $u_{AR} = B_R(t_O)$ which implies $t_O = \text{LWT}_R$. We only need to consider equilibria in which there is a strictly positive probability the responder accepts an outside alternative before LWT_R, since we already covered the case when he rejects all alternatives until then. Then, as already demonstrated, the proposer has an incentive to deviate unless $\text{LWT}_R = t_{IP}$. Hence, if $t_{IP} = \text{LWT}_R$ then the proposer making an exploding offer at LWT_R is an equilibrium, irrespective of the responder's strategy and there are no other equilibria in which the responder is indifferent about his strategy.

It remains to show that there are no mixed-strategy equilibria. Assume the opposite: there exists a mixed strategy equilibrium. If A9 holds, making a deadline offer is a strictly dominated strategy, thus any equilibrium strategy for the proposer has only exploding offers (with $t_O = t_D$) in its support. Then, in equilibrium the proposer randomises over t_O drawing it from a distribution with some cdf F_O. Denote by $B_R^*(t)$ the expected payoff to the responder from time t onwards under an optimal policy given the proposer randomises in this way and no proposer's offer was received before t. Then, given A8 holds, when an outside alternative is received at time t_1, the responder should accept (reject) it if $u_{AR} > (\leq) B_R^*(t_1)$. Since $B_R^*(t)$ is non increasing, the responder follows a pure strategy of rejecting alternatives up to t^* and accepting them afterwards, where $t^* = \sup\{t \in [0,T] \mid u_{AR} \leq B_R^*(t)\}$. However, as already discussed, there is just one possible offer time for the responder which is a best response to this proposer's strategy (that is $t_O = \max\{t^*, t_{IP}\}$). By virtue of A8, $t_O = t_D$, thus the proposer strategy in equilibrium is pure as well. □

Proof of Proposition 4. We present the proof of Proposition 4 for the setting of Proposition 1 only ($u_{AP} > u_{MP}$ and $u_{AR} > u_{MR}$). The proposition holds in other cases as well and is derived following analogous procedures. The proof is done in three parts.

Part 1: When either or both agents can make offers but neither can commit to an ultimatum, the set of equilibria is the same as in the basic model game in which the less eager agent acts as the proposer. Analogously to (11), we can define R's expected utility as a function of offer time which is then given by

$$\pi_R(t_O) = \left(1 - \left(1 - e^{-\lambda_P t_O}\right) e^{-\lambda_R T}\right)(u_{AR} - u_{FR}) - e^{-(\lambda_P + \lambda_R) t_O}(u_{AR} - u_{MR}) + u_{FR}. \tag{13}$$

Notice that $\pi_P(t_O)$ and $\pi_R(t_O)$ are expected utilities of the agents when both of them are accepting any outside alternative found and concluding a mutual deal at t_O if both of them are still available. These expectations are the same irrespective of which agent is the one who made the offer at t_O. As shown in Proposition 1, these two functions are unimodal and $\pi_R(t_O)$ will be maximised at t_{IR}, the R's ideal time to make a deal which is defined analogously to t_{IP} and given by

$$t_{IR} = \begin{cases} 0 & \text{if } u_{MR} \geq u_{AR} \\ \min\left\{T, \left(T - \frac{1}{\lambda_R}\ln\left(\frac{u_{AR}-u_{FR}}{u_{AR}-u_{MR}}\right) + \frac{1}{\lambda_R}\ln\left(\frac{\lambda_P+\lambda_R}{\lambda_P}\right)\right)^+\right\} & \text{if } u_{MR} < u_{AR}. \end{cases} \tag{14}$$

Without loss of generality (w.l.o.g.), we can assume that P is the less eager agent ($t_{IP} \geq t_{IR}$). At t_{IP}, R's strictly dominant strategy is to accept any exploding offer given at that time, due to $\pi_R(t_O)$ being a decreasing function after t_{IR} and $t_{IP} \geq t_{IR}$. Thus the P can attain maximum possible payoff from the game ($\pi_P(t_{IP})$)

by rejecting any offer made by R before t_{IP} and making an exploding offer then. Since $t_{IP} \geq t_{IR} \geq \text{SAD}_R$, P making an exploding offer at t_{IP} was also the sole equilibrium in the basic model.

Similarly, if P does not have the ability to make an offer but R does, P can count on the fact that R's strictly dominant strategy at t_{IP} will be to give an exploding offer then (for reasons already described). Thus, P can reject any R's offer before t_{IP}, resulting in the same equilibrium as in other cases.

Part 2: If both agents can make offers and the more eager agent can commit to an ultimatum, the set of equilibria is the same as in the basic model game in which the more eager agent acts as the proposer. Consider the agent who is more eager, the one who's ideal time to make an offer is sooner. We can take P to be this agent w.l.o.g., so $t_{IP} < t_{IR}$. Using Proposition 1: if R never makes an offer, P maximizes expected utility by making an exploding offer at $\max\{\text{SAD}_R, t_{IP}\}$).

We still need to show that R cannot make himself better off by making an offer. Since this offer is issued as an ultimatum, it also prevents R from issuing an offer later on, thus the decision when facing P's offer is the same one as before: R accepts it if its value is higher than the expectation of the rest of the search for outside alternatives (condition which is true from SAD_R onward). Thus, R can only make an offer before $\max\{\text{SAD}_R, t_{IP}\}$, however that will decrease R's expected utility since the expected payoff (π_R) is increasing up to t_{IR} and $t_{IR} \geq \max\{\text{SAD}_R, t_{IP}\}$. Thus, this game will have the same equilibrium as a game where the more eager agent is the proposer.

Part 3: an agent derives higher payoff in equilibrium if he or she is the one who acts as the sole proposer. In this part, we consider how the payoffs change in the basic model if the agents switch roles, i.e., only R is able to make offers and can make them ultimatums. Denote the basic model as BM, and the one with same parameters but switched roles as SM (switched model).

Since $u_{AP} > u_{MP}$ and $u_{AR} > u_{MR}$, conditions of Proposition 1 still hold in SM, with both agents preferring outside alternatives to a mutual deal. Analogously to SAD_R, we can look at SAD_P, P's shortest acceptable deadline, given by

$$\text{SAD}_P = \begin{cases} 0 & \text{if } u_{MP} \geq u_{AP} \\ \left(T - \frac{1}{\lambda_P} \ln\left(\frac{u_{AP}-u_{FP}}{u_{AP}-u_{MP}}\right)\right)^+ & \text{if } u_{MP} < u_{AP}. \end{cases} \tag{15}$$

Consider $\pi_P(t_O)$ and $\pi_R(t_O)$, the expected utilities of the agents as a function of offer time, as given by (11) and (13), which are then maximised at t_{IP} and t_{IR} respectively. Applying Proposition 1, there will be a sole equilibrium of SM in which R makes an offer at $\max\{t_{IR}, \text{SAD}_P\}$.

First, let us show that P will be worse off in SM. If $t_{IP} \geq \text{SAD}_R$, in BM the offer was made at t_{IP} and any change can only make P worse off, since $t_{IP} \in \arg\max \pi_P(t_O)$. If $t_{IP} < \text{SAD}_R$, the offer was originally made at SAD_R, whereas in SM it will be made at $\max\{t_{IR}, \text{SAD}_P\}$. However, from (2) and (14) we have $t_{IR} \geq \text{SAD}_R$ and from (15) and (3) we have $t_{IP} \geq \text{SAD}_P$. Thus if R is the one making an offer, it will be made at $t_{IR} \geq \text{SAD}_R > t_{IP}$. Since π_P is maximised at t_{IP} and decreasing afterwards, this will decrease P's expectation.

Similarly, the switch of roles with benefit R. If $t_{IR} \geq \text{SAD}_P$, in equilibrium of SM, R will be making an offer at t_{IR} which maximizes his expected utility. On the other hand, if $t_{IR} < \text{SAD}_P$, the new equilibrium will have an offer at SAD_P. Since R's expected utility is maximised at t_{IR} while being decreasing after that and $\text{SAD}_R \le t_{IR} < \text{SAD}_P \le t_{IP}$, switch of equilibrium from t_{IP} to SAD_P will benefit R. □

Proof of Proposition 5. (Solution under a minimum deadline.) If $u_{AP} > u_{MP}$ and $u_{AR} > u_{MR}$, this game corresponds to the one in Proposition 1, with the only difference being the additional restriction on deadline times ($t_D \in [t_O + \Delta, T]$ instead of $t_D \in [t_O, T]$). Whether the responder accepts an offer depends only on t_D and not on t_O, thus for a fixed t_D, the proposer prefers to give the offer as late as possible, which in this case is at $t_O = t_D - \Delta$. This allows us to us to reduce the proposer's problem to univariate maximisation,[17]

$$\underset{t_D \in [\max\{\text{SAD}_R, \Delta\}, T]}{\arg\max} \left(1 - \left(1 - e^{-\lambda_R t_D}\right) e^{-\lambda_P (T-\Delta)}\right)(u_{AP} - u_{FP}) - e^{-\lambda_P (t_D - \Delta) - \lambda_R t_D}(u_{AP} - u_{MP}) + u_{FP}, \quad (16)$$

which simplifies to

$$\underset{t_D \in [\max\{\text{SAD}_R, \Delta\}, T]}{\arg\max} u_{AP}(1 - e^{\lambda_P \Delta}) + e^{\lambda_P \Delta} \pi_P(t_D),$$

where π_P is given by (11) in the proof of Proposition 1. Thus, the unique unconstrained maximiser of π_P, that is $t^* = T - \frac{1}{\lambda_P} \ln\left(\frac{u_{AP} - u_{FP}}{u_{MP} - u_{FP}} \frac{\lambda_R}{(\lambda_P + \lambda_R)}\right)$ as given by (12), is the unconstrained solution of (16) as well. Notice from (3) that $t^* \in [0, T] \Rightarrow t^* = t_{IP}$. Since the objective function in (16) is unimodal with the sole peak at $t = t^*$ (follows from the same property of π_P), it is increasing before t^* and decreasing afterwards. Hence, the solution to the constrained problem is $t_D = \max\{\Delta, \text{SAD}_R, t_{IP}\}$, $t_O = t_D - \Delta$.

If $u_{AP} \le u_{MP}$, $u_{AR} > u_{MR}$, this game corresponds to the one in Proposition 2. Following the steps of Proposition 2, the proposer always makes an offer at the earliest time such that it will be accepted. In this case, this time is $t_D = \max\{\Delta, \text{SAD}_R\} = \max\{\Delta, t_{IP}, \text{SAD}_R\}$, $t_O = t_D - \Delta$.

If $u_{AP} > u_{MP}$, $u_{AR} \le u_{MR}$, this game corresponds to the one in Proposition 3. In this case, the responder always accepts the proposer's offer if he is available, thus the proposer has nothing to gain from leaving a deadline longer than the minimum required once. By A9, any equilibrium offer will thus have $t_D = t_O + \Delta$. The equilibria can then be characterised by following the steps of the proof of Proposition 3, which yields the following. If $t_{IP} \le \text{LWT}_R + \Delta$, then $\forall t_D \in [\max\{t_{IP}, \Delta\}, \min\{\text{LWT}_R + \Delta, T\}]$, the proposer making an offer at $t_O = t_D - \Delta$ with deadline t_D and the responder rejecting all outside alternatives before then is an equilibrium. Otherwise (if $t_{IP} > \text{LWT}_R + \Delta$), there is a sole equilibrium of the game in which the proposer makes an offer at $t_O = t_{IP} - \Delta$ with deadline $t_D = t_{IP} = \max\{\Delta, t_{IP}, \text{SAD}_R\}$ and the responder accepts all outside alternatives from the beginning. □

Proof of Corollary 2. Denote basic model as BM, and the same model but with a minimum deadline Δ as DM. If $u_{AR} > u_{MR}$ then the equilibrium of BM is the result of expected utility maximisation for the proposer, as given in Propositions 1,2. Under DM, the objective function being maximised is the same, but over a smaller domain (the proposer is restricted in his choice of t_O, t_D). Thus the proposer's expected utility in equilibrium is equal or lower in DM compared to BM. Likewise, the responder's expected utility is higher since a minimum deadline causes him to receive his offer earlier, thus with higher probability (since $\max\{\Delta, t_{IP}, \text{SAD}_R\} - \Delta \le \max\{t_{IP}, \text{SAD}_R\}$), while it expires at the same time or later than in the basic model (since $\max\{\Delta, t_{IP}, \text{SAD}_R\} \ge \max\{t_{IP}, \text{SAD}_R\}$).

[17] The proposer's expected utility (16) is derived following the method of Proposition 1.

If $u_{AR} \leq u_{MR}$, we will have to separate the problem in 3 parts and account for multiple equilibria.

Part 1: Consider $t_{IP} \leq \mathrm{LWT}_R$. By Proposition 3, in BM the set of all equilibria consists of the proposer making an exploding offer at any $t_O \in [t_{IP}, \mathrm{LWT}_R]$ and the responder rejecting all outside alternatives until t_O. By Proposition 5, in DM the set of all equilibria consists of the proposer making an offer at any $t_O \in [(t_{IP} - \Delta)^+, \min\{\mathrm{LWT}_R, T - \Delta\}]$, setting a minimum possible deadline while the responder rejects outside alternatives until t_O.

Let us look at the payoffs in these equilibria. In both BM and DM settings, the proposer searches until t_O accepting outside alternatives (if found) and makes an offer to the responder at t_O (if still available), with this offer always being accepted by the responder, giving the proposer a total expectation of $(1 - e^{-\lambda_P t_O})(u_{AP} - u_{MP}) + u_{MP}$, which is an increasing function of t_O. The responder on the other hand rejects all alternatives until t_O, at which point he either receives the proposer's offer which he accepts or in absence of such offer starts accepting outside alternatives which gives him total expectation of $e^{-\lambda_P t_O}(u_{MR} - u_{FR}) + (1 - e^{-\lambda_P t_O})e^{-\lambda_R (T - t_O)}(u_{AR} - u_{FR}) + u_{FR}$, a decreasing function of t_O. Notice that for both agents, the payoffs in these equilibria depend solely on t_O.

Finally, for all $t_O \in [t_{IP}, \min\{\mathrm{LWT}_R, T - \Delta\}]$, there exists an equilibrium where an offer is made at t_O in both BM and DM with payoffs being the same in either model. For $t_O \in [(t_{IP} - \Delta)^+, t_{IP}]$, in DM, but not in BM, there exists an equilibria where offer is made at t_O. Since the payoff functions are increasing for the proposer and decreasing for the responder in t_O, all of these equilibria have higher payoff for the responder and lower for the proposer than any equilibria of BM. Likewise, for $t_O \in [\min\{\mathrm{LWT}_R, T - \Delta\}, \mathrm{LWT}]$, there exist only BM equilibria where offer is made at t_O, with them paying higher to the proposer and lower to the responder than any equilibria of DM.

Part 2: Consider $\mathrm{LWT}_R + \Delta \geq t_{IP} > \mathrm{LWT}_R$. Using Propositions 3 and 5, in BM there is a sole equilibrium in which the proposer makes an exploding offer at $t_O = t_{IP}$ while the responder accepts all outside alternatives, while in DM any offer with minimum deadline given at $t_O \in [(t_{IP} - \Delta)^+, \min\{\mathrm{LWT}_R, T - \Delta\}]$ and the responder rejecting outside alternatives until then will be an equilibrium. The DM equilibria all have the same payoffs to the proposer as exploding offers made at the same time (shown in Part 1 of this proof) and, since $t_{IP} > \mathrm{LWT}_R$, this payoff is lesser than one the proposer can get by making offer at t_{IP}, even if the responder accepts all alternatives before this (shown in the proof of Proposition 3). Thus, the proposer is better off in BM equilibrium. Likewise, the DM equilibria give higher payoff to the responder than the BM one as all of them include the responder receiving proposer's offer sooner and thus with higher probability than in BM.

Part 3: Consider $t_{IP} > \mathrm{LWT}_R + \Delta$. Here, BM equilibrium has the offer being made at t_{IP} as an exploding offer, while DM equilibrium has the offer being made at $t_{IP} - \Delta$ and expiring at t_{IP}. The responder accepts all outside alternatives in both equilibria. The responder receives his offer sooner and thus with higher probability in DM, while it expires at the same time, making his expectation in DM higher. For the proposer, given the responder is accepting all outside alternatives, her expectation from making an exploding offer at t_{IP} is higher than with any other offer, as shown in Proposition 1. Therefore, the proposer is better off in BM. □

Lemma 2 *By defining reservation price as a function of "search distance" instead of time, every single agent search problem as given by (7) or (9) can be reduced to a search problem with a homogenous arrival rate of* 1.

Proof. Define a bijection $S(t):[0,T]\to[0,\int_0^T \lambda(x)\,dx]$ with mapping rule

$$S(t)=\int_0^t \lambda(x)\,dx, \tag{17}$$

as done in (Ross 1996, p. 79). Also, define function $\overline{\xi^S}:[0,\int_0^T \lambda(x)\,dx]\to\mathbb{R}$ with mapping rule

$$\overline{\xi^S}(s)=\bar{\xi}\left(S^{-1}(s)\right). \tag{18}$$

This will allow us to express single agent search policy ($\bar{\xi}$ given by (7)) in terms of $s\in[0,\int_0^T \lambda(x)\,dx]$ instead of $t\in[0,T]$. Intuitively, we can think of arrival rate λ as speed of search, in which case s represents "search distance" covered, equal to the expected number of arrivals during this period. Taking a derivative of (18) yields

$$\left(\overline{\xi^S}\right)'(s)=-\int_{\overline{\xi^S}(s)}^{\infty}(x-\overline{\xi^S}(s))\,dF(x;t). \tag{19}$$

Hence, instead of solving the problem (7) directly we can find the reservation price by solving the dual problem (19) with border condition $\overline{\xi^S}\left(\int_0^T \lambda(x)\,dx\right)=u_F$. □

Proof of Theorem 1. Let us show that for any choice of $\xi_R(t)$, t_O^{**} and t_D^{**} there exists t_O^* such that making an exploding offer at t_O^* will give the proposer a higher expected payoff. If the proposer derives higher utility from the responder rejecting this offer than accepting it ($\bar{\xi}_P(t_D^{**})>u_{MP}$) or if the deadline is so short that the offer will always be rejected ($t_D^{**}<\text{SAD}_R$), the situation is trivial as the proposer can do better by giving an exploding offer at the end of the search horizon ($t_O^*=T$). The proof is based on looking at the situations where this is not the case ($t_D\ge\text{SAD}_R$ and $u_{MP}\ge\bar{\xi}_P(t_D^{**})$) and constructing t_O^* for three separate cases depending on the values of t_O^{**} and t_D^{**}. For any choice of t_O and t_D, the proposer's optimal policy can be constructed in the following way. Denote by $P(\xi_R,t_O,t_D)$ the probability the responder accepts the proposer's offer given at t_O and expiring at t_D when following policy ξ_R. The proposer's problem before t_O is then reduced to a single agent problem with horizon t_O and the fallback value equal to the proposer's expectation from t_O onwards given by $\bar{\xi}_P(t_D)+P(\xi_R,t_O,t_D)(u_{MP}-\bar{\xi}_P(t_D))$. Application of (7) then yields the optimal ξ_P.

In each of the following cases we will construct t_O^* in such a way that it increases this fallback value by simultaneously increasing probability the responder will accept the proposer's offer ($P(\xi_R,t_O^*,t_O^*)\ge P(\xi_R,t_O^{**},t_D^{**})$), as well the time left to search when the responder rejects the offer ($(t_O^*\le t_D^{**})\Rightarrow\bar{\xi}_P(t_O^*)\ge\bar{\xi}_P(t_D^{**})$), while also giving the proposer more time to search before this "horizon" ($t_O^*\ge t_O^{**}$), with at least one of these three inequalities being strict.

Case 1: $t_D^{**}=\text{SAD}_R$. From (8), the responder uses policy $\bar{\xi}_R(t;u_{MR},\text{SAD}_R)$ while he holds the proposer's offer. However, at time SAD_R the responder derives the same utility from accepting the proposer's offer as he does from continuing his search (from the definition of SAD_R). Thus, from (7) and (8), we have

$\bar{\xi}_R(t;u_{MR},\text{SAD}_R)=\bar{\xi}_R(t)$. Since all rationalizable policies have reservation prices at least as high as $\bar{\xi}_R(t)$ (using Lemma 1), the responder has lower chances of still being available at time SAD_R if he uses policy $\bar{\xi}_R(t)$ with any other rationalizable policy. Thus, an exploding offer at $t_O^*=\text{SAD}_R$ will have have the responder using $\xi_R(t)$ instead of $\bar{\xi}_R(t)$, increasing the chances the responder accepts while also giving the proposer additional time $[t_O^{**},\text{SAD}_R]$ to search.

Case 2: $t_O^{**}\geq \text{SAD}_R$, $t_D^{**}>\text{SAD}_R$. Setting $t_O^*=t_O^{**}$ gives the proposer additional $[t_O^{**},t_D^{**}]$ to search when her offer is rejected, while also increasing the chance the responder accepts as exploding offer will always be accepted if the responder is still available, but a deadline offer will give him additional search time.

Case 3: $t_O^{**}<\text{SAD}_R, t_D^{**}>\text{SAD}_R$. We first prove this case for the situation where $\lambda_R(t)$ is constant. Notice that if $t_D\geq\text{SAD}_R$ then at any given time t, an agent who holds the proposer's offer expiring at t_D is in exactly the same situation as an agent at time $t+\Delta$ who holds an offer expiring at $t_D+\Delta$. Thus the policies of two agents who hold the proposer's offer but with different deadlines are horizontal translations of each other. Specifically, comparing the offer from t_O^{**} to t_D^{**} to the one made at the same time but expiring at SAD_R we have

$$\forall t\in[t_O^{**},t_D^{**}]\mid t-t_D^{**}+\text{SAD}_R\geq 0:\bar{\xi}_R(t;u_{MR},t_D^{**})=\bar{\xi}_R(t-t_D^{**}+\text{SAD}_R;u_{MR},\text{SAD}_R). \tag{20}$$

Let us show that making an offer at $t_O^*=t_O^{**}$ but having it expire at SAD_R will give the proposer a higher payoff than the one made at a same time but expiring at t_D^{**}. In both cases the responder will search until t_O^{**} using policy ξ_R at which point he will change policies according to the offer he is given. Using (20), the probability that a responder holding an offer which expires at SAD_R accepts an outside alternative between t_O^{**} and SAD_R is the same as the probability of the responder holding an offer which expires at t_D^{**} accepting an outside alternative between $t_D^{**}-\text{SAD}_R+t_O^{**}$ and t_D^{**}. Consequently, making the offer expire at SAD_R gives higher chance of acceptance as well as more time to search if responder rejects the offer, yielding higher expectation. From case 2 of this proof we know that an exploding offer at $t_O^*=\text{SAD}_R$ will yield even higher expectation.

If $\lambda_R(t)$ is not constant, (20) does not hold so in order to prove this step we first need to apply the method of Lemma 2 to express the policy of an agent holding an offer in terms of "search distance" with a function $\overline{\xi_R^S}$ such that $\bar{\xi}_R(t;u_{MR},t_D^{**})=\overline{\xi_R^S}(S(t);u_{MR},S(t_D^{**}))$, where $S(t)$ is given by (17). Using this transformation, the following variant of the translation property holds:

$$\overline{\xi_R^S}(s;u_{MR},S(t_D^{**}))=\overline{\xi_R^S}(s-S(t_D^{**})+S(\text{SAD}_R);u_{MR},S(\text{SAD}_R)). \tag{21}$$

From (21), the probability that the responder holding an offer which expires at SAD_R accepts an outside alternative between t_O^{**} and SAD_R is the same as the probability of the responder holding an offer which expires at t_D^{**} accepting an outside alternative between $S^{-1}(S(t_D^{**})-S(\text{SAD}_R)+S(t_O^{**}))>t_O^{**}$ and t_D^{**}. Hence, we arrive to the same conclusion and complete the proof. $\square$

Lemma 3 *Let $\bar{\xi}_F(t)$ be the policy which solves the nonstationary single agent problem given by (9) and let $\bar{\xi}_G(t)$ be the policy which solves 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 $t\in[0,T], y\in\mathbb{R}$:*

$$\int_y^\infty(\overline{G}(x;t)-\overline{F}(x;t))\,dx\geq 0, \tag{22}$$

then $\forall t \in [0,T] : \bar{\xi}_G(t) \geq \bar{\xi}_F(t)$*. If the inequality in (22) is strict* $\forall t \in [0,T], y \in \mathbb{R}$*, then* $\forall t \in [0,T) : \bar{\xi}_G(t) > \bar{\xi}_F(t)$*.*

Proof. From (9), the differential equation for $\bar{\xi}_G$ is

$$\bar{\xi}'_G(t) = -\lambda(t) \int_{\bar{\xi}_G(t)}^{\infty} (x - \bar{\xi}_G(t))\, dG(x;t) = -\lambda(t) \int_{\bar{\xi}_G(t)}^{\infty} \overline{G}(x;t)\, dx. \tag{23}$$

Subtracting (9) from (23) yields

$$\bar{\xi}'_G(t) - \bar{\xi}'_F(t) = \lambda(t) \int_{\bar{\xi}_F(t)}^{\bar{\xi}_G(t)} \overline{F}(x;t)\, dx - \lambda(t) \int_{\bar{\xi}_G(t)}^{\infty} (\overline{G}(x;t) - \overline{F}(x;t))\, dx, \tag{24}$$

with border conditions $\bar{\xi}_F(T) = \bar{\xi}_G(T) = u_F$. Since $\int_y^{\infty} (\overline{G}(x;t) - \overline{F}(x;t))\, dx \geq 0, \forall y \in \mathbb{R}$, from (24) it follows that

$$\bar{\xi}'_G(t) - \bar{\xi}'_F(t) \leq \lambda(t) \int_{\bar{\xi}_F(t)}^{\bar{\xi}_G(t)} \overline{F}(x;t)\, dx. \tag{25}$$

If $\bar{\xi}_G(t) < \bar{\xi}_F(t)$ the right hand side of (25) is negative giving us $\bar{\xi}_G(t) - \bar{\xi}_F(t) < 0 \Rightarrow \bar{\xi}'_G(t) - \bar{\xi}'_F(t) < 0$. Because of this, if there exists $t^* \in [0,T]$ such that $\bar{\xi}_G(t^*) - \bar{\xi}_F(t^*) < 0$, then $\forall t \geq t^* : \bar{\xi}_G(t) - \bar{\xi}_F(t) < 0$. However, from border conditions we known that $\bar{\xi}_G(T) - \bar{\xi}_F(T) = 0$, which is contradictory to existence of such t^*. Consequently, $\forall t \in [0,T] : \bar{\xi}_G(t) - \bar{\xi}_F(t) \geq 0$. If we have $\int_y^{\infty} (\overline{G}(x;t) - \overline{F}(x;t))\, dx > 0, \forall y \in \mathbb{R}$, then inequality (25) is strict and thus $\bar{\xi}_G(t) - \bar{\xi}_F(t) = 0 \Rightarrow \bar{\xi}'_G(t) - \bar{\xi}'_F(t) < 0$, hence $\bar{\xi}_G(t) - \bar{\xi}_F(t)$ can have no more than one zero. From border conditions we know that $\bar{\xi}_G(T) - \bar{\xi}_F(T) = 0$, therefore $\forall t \in [0,T) : \bar{\xi}_G(t) - \bar{\xi}_F(t) > 0$. □

Proof of Theorem 2. We prove this theorem for stationary arrival rate λ_R. If this is not the case, responder's search process first needs to be transformed to the one with stationary arrival rate, as done in Lemma 2, after which the proof proceeds in the same way.

The proof follows the steps of Theorem 1, with the same argument being valid for cases 1 and 2. Case 3 ($t_O^{**} < \mathrm{SAD}_R, t_D^{**} > \mathrm{SAD}_R$) requires a different approach. As in Theorem 1, we will demonstrate that there is higher probability that the responder accepts an offer expiring at SAD_R than the one expiring at t_D^{**}. The responder who holds the proposer's offer which expires at t_D^{**} will use a policy $\bar{\xi}_R(t; u_{MR}, t_D^{**})$ which solves

$$(\bar{\xi}_R)'(t; u_{MR}, t_D^{**}) = -\lambda_R \int_{\bar{\xi}_R(t; u_{MR}, t_D^{**})}^{\infty} (x - \bar{\xi}_R(t; u_{MR}, t_D^{**}))\, dF_R(x;t),$$

with the border condition $\bar{\xi}_R(t_D^{**}; u_{MR}, t_D^{**}) = u_{MR}$, while one who holds an offer which expires at SAD_R will use a policy which solves the same differential equation with a different border condition: $\bar{\xi}_R(\mathrm{SAD}_R; u_{MR}, \mathrm{SAD}_R) = u_{MR}$. If $\forall t_1, t_2 \in [0,T] \mid t_1 < t_2 : F_R(x;t_1)$ dominates $F_R(x;t_2)$, then from the integral condition (10) we have $\forall y \in \mathbb{R} : \int_y^{\infty} \overline{F}_R(x;t) dx \downarrow t$ and from Lemma 3 it follows that

$$\forall t \in (t_O^{**}, \mathrm{SAD}_R) : \bar{\xi}_R(t - \mathrm{SAD}_R + t_D^{**}; u_{MR}, t_D^{**}) \leq \bar{\xi}_R(t; u_{MR}, \mathrm{SAD}_R). \tag{26}$$

Then, the responder who is still available at t_O^{**} will accept the proposer's offer with deadline SAD_R with probability $\exp\{-\lambda_R \int_{t_O^{**}}^{\mathrm{SAD}_R} \overline{F}_R(\bar{\xi}_R(t; u_{MR}, \mathrm{SAD}_R); t)\, dt\}$. Applying (26), this probability is greater than

$$\exp\{-\lambda_R \int_{(t_O^{**} - \mathrm{SAD}_R + t_D^{**})}^{t_D^{**}} \overline{F}_R(\bar{\xi}_R(t; u_{MR}, t_D^{**}); t)\, dt\} \geq \exp\{-\lambda_R \int_{t_O^{**}}^{t_D^{**}} \overline{F}_R(\bar{\xi}_R(t; u_{MR}, t_D^{**}); t)\, dt\},$$

where the right hand side is the probability that the responder who is still available at t_O^{**} accepts the proposer's offer expiring at t_D^{**}, which completes the proof. □

Example 1A calculations. First, consider the responder's single-agent policy: the expectation of his search in the second period is \$1, thus he should accept any offer he finds in the first period as it gives \$2. This policy gives him an expectation of $0.9 \times \$2 + 0.1 \times \$1 = \$1.9$, thus the responder's shortest acceptable deadline will be $\text{SAD}_R = 1$ since that is the first time at which the value of the proposer's offer is higher than the rest of his search. Exploding offer at time 0 will just be rejected by the responder, meaning the proposer only receives expectation of her own search which is equal to $0.2 \times \$2 = \$0.4$. Exploding offer at time 1 will be accepted with probability 0.1 (chance that the responder is still available), by making such an offer the proposer also has the opportunity to go through the whole search giving her a total expectation of $0.2 \times \$2 + 0.8 \times 0.1 \times \$1.5 = \$0.52$. Lastly, giving an offer at time 2 is worse as it gives the responder the opportunity to search in period 2, for total proposer expectation of $0.2 \times \$2 + 0.8 \times 0.09 \times \$1.5 = \$0.508$.

The proposer can do a lot better by giving a deadline offer, creating a deterrence effect and thus inducing the responder to reject outside alternatives in the first period. The deadline offer expiring at time 1 will not be enough to cause this effect. The responder holding such an offer will receive \$1.5 at time 1, so it will still be in his interest to accept period one offers since they yield more than this. However, this will change if the responder receives an offer which expires at time 2. The responder holding such an offer receives \$1.5 at the end of the search horizon, thus the expectation of his search in the second period is $0.1 \times \$10 + 0.9 \times \$1.5 = \$2.35$, so he prefers to reject the first period alternatives which give only \$2, meaning he will accept the proposer's offer with probability 0.9 (chance he does not find anything in period 2). The proposer who makes such an offer has to forsake her own search, but due to deterrence effect has a higher chance of making a deal with the proposer, for a total expectation of $0.9 \times \$1.5 = \$1.35$, more than double than what she can make with the best possible exploding offer. Since neither agent can benefit from deviating, this is an equilibrium. Note that an exploding offer at time 2 (or a deadline one made at time 1 but expiring at time 2) cannot be an equilibrium since the responder's strictly dominant strategy in the second period is to accept outside alternatives, thus no matter what either of them do in period 1, the proposer is strictly better off giving an exploding offer at time 1 than giving a deadline offer or waiting for time 2 before making one.

Example 1B calculations. Assume the responder uses the optimal single-agent policy, which in this case is to accept the first outside alternative he finds. If the proposer wants to make an exploding offer, she has only 3 strategies available which are not strictly dominated. She can make an offer at time 0 setting its value to \$1.9. This offer is always accepted and yields \$1.1 to the proposer. Alternatively, she can make an offer at 1 and set its value to \$1. The responder will accept this offer if he is still available at time 1 (probability 0.1), but the proposer can also search during the first period for the offer with value \$2 which she will find with probability 0.2. Thus the proposer's expectation when making such offer is $0.2 \times \$2 + 0.8 \times 0.1 \times \$2 = \$0.56$. Lastly, she can make an offer at time 2 and set its value to \$0. The responder will only accept it if he finds nothing, yielding an expectation of $0.2 \times \$2 + 0.8 \times 0.1 \times 0.9 \times \$3 = \$0.616$ to the proposer.

Since deadline offers from t_O to t_D which do not cause deterrence effect are dominated by exploding offers at t_D, the proposer's only other non-dominated option is to use deadline offers which set x to the minimum possible value which will cause the deterrence effect. For an offer made at 0 with deadline 2, we can find such x by solving $\$2 = 0.1 \times \$10 + 0.9x$, yielding $x = \$10/9$. Making that offer will give the proposer expectation $0.9 \times (\$3 - \$10/9) = \$1.7$. A check of other possible deadline offers needs to be done in order to show that this offer is indeed the best response to the responder using the single-agent policy. Deterring the responder from accepting offers in period 2 is not feasible since that would require setting $x = \$10$ which will result in negative payoff to the responder. Thus the only other deadline offer which is not eliminated is one made at time 0, expiring at 1 with value \$2 (lowest value which causes deterrence). Such an offer is always accepted, but yields only \$1 to the proposer, thus is not her best response either. Since neither agent can benefit from deviating, this is an equilibrium.

References

Abdulkadiroğlu, A., T. Sönmez. 2013. Matching markets: Theory and practice. *Advances in Economics and Econometrics* **1** 3–47.

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.

Emerson, R. M. 1962. Power-dependence relations. *American Sociological Review* **27**(1) 31–41.

Ferguson, T. S. 2012. *Optimal stopping and applications*. Ebook at http://www.math.ucla.edu/~tom/.

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

Gronau, R. 1971. Information and frictional unemployment. *American Economic Review* **61**(3) 290–301.

Guttman, I. 1960. On a problem of L. Moser. *Canadian Mathematical Bulletin* **3** 35–39.

Lancaster, T., A. Chesher. 1983. An econometric analysis of reservation wages. *Econometrica* **51**(6) 1661–76.

Lau, N., Y. Bart, J. N. Bearden, I. Tsetlin. 2014. Exploding offers can blow up in more than one way. *Decision Analysis* **11**(3) 171–188.

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.

Mas-Colell, A., M. D. Whinston, J. R. Green. 1995. *Microeconomic theory*, vol. 1. Oxford university press New York.

McCall, J. J. 1970. Economics of information and job search. *The Quarterly Journal of Economics* **84**(1) 113–26.

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

Moser, L. 1956. On a problem of Cayley. *Scripta Mathematica* **22** 289–292.

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.

Nash Jr, J. F. 1950. The bargaining problem. *Econometrica* **18**(2) 155–162.

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

Pais, J. 2008. Incentives in decentralized random matching markets. *Games and Economic Behavior* **64**(2) 632–649.

Ross, S. M. 1996. *Stochastic processes*, vol. 2. John Wiley & Sons New York.

Roth, A. E., M.A.O. Sotomayor. 1992. *Two-sided matching: A study in game-theoretic modeling and analysis*. 18, Cambridge University Press.

Roth, A. E., X. Xing. 1997. Turnaround time and bottlenecks in market clearing: Decentralized matching in the market for clinical psychologists. *Journal of Political Economy* **105**(2) 284–329.

Rubinstein, A. 1982. Perfect equilibrium in a bargaining model. *Econometrica* **50** 97–110.

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

Shiryaev, A. N. 1976. *Optimal stopping rules*. Nauka Moscow.

Sönmez, T., M. U. Ünver. 2011. Matching, allocation, and exchange of discrete resources. *Handbook of Social Economics* **1** 781–852.

Tang, W., J. N. Bearden, I. Tsetlin. 2009. Ultimatum deadlines. *Management Science* **55**(8) 1423–1437.

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

Weitzman, M. L. 1979. Optimal search for the best alternative. *Econometrica* **47**(3) 641–54.