INSEAD

The Business School for the World

Faculty & Research Working Paper

Ultimatum Deadlines

Wenjie TANG
J. Neil BEARDEN
Ilia TSETLIN
2008/39/DS

Ultimatum Deadlines

by

Wenjie Tang*

J. Neil Bearden**

and

Ilia Tsetlin***

* PhD Candidate in Decision Sciences at INSEAD, 1 Ayer Rajah Avenue 138676 Singapore Tel 6567992538 , e-mail: wenjie.tang@insead.edu

** Assistant Proffessor of Decision Sciences at INSEAD, 1 Ayer Rajah Avenue 138676 Singapore Tel 6567995286 , e-mail: neil.bearden@insead.edu

*** Assistant Proffessor of Decision Sciences at INSEAD, 1 Ayer Rajah Avenue 138676 Singapore Tel 6567995339 , e-mail: ilia.tsetlin@insead.edu

Submitted to *Management Science*
manuscript

Ultimatum Deadlines

Wenjie Tang
INSEAD, 138676 Singapore, wenjie.tang@insead.edu

J. Neil Bearden
INSEAD, 138676 Singapore, neil.bearden@insead.edu

Ilia Tsetlin
INSEAD, 138676 Singapore, ilia.tsetlin@insead.edu

An important characteristic of any offer is the deadline at which it expires. We consider an ultimatum deadline game in which the proposer's decision variable is the offer deadline, while the responder faces a standard finite-horizon search problem. We show that the responder's strategy is characterized by a shortest acceptable deadline: At the time of deadline, he accepts an offer if the deadline is longer than his shortest acceptable deadline, and rejects it otherwise. If the proposer has all information available to the responder, the optimal deadline is the responder's shortest acceptable deadline. If the proposer is uncertain about the responder's situation, the optimal deadline gets longer, unless this uncertainty is very large. After normative analysis of the deadline setting problem, we present results from a behavioral study of the game. The average shortest acceptable deadline set by the responders equals the optimal one, while the proposers tend to set deadlines that are too short. The prescriptive conclusion for a proposer, emerging from the model and the experiment, is that in case of uncertainty it is better to set a deadline longer than what would be optimal if uncertainty were ignored.

June, 2008

Key words: deadlines, job search, ultimatum game

"I love deadlines. I especially like the whooshing sound they make as they go by."

— *Douglas Adams*

1. Introduction

Microsoft made an unsolicited bid for Yahoo! in early 2008. According to media reports, Steve Ballmer, the CEO of Microsoft, issued an ultimatum that gave Yahoo! three weeks to meet at the negotiation table. During the same period, reports showed that Yahoo! was discussing a limited ad-sharing partnership with Google, its (and Microsoft's) big rival, and also discussing deals with

Time Warner's AOL. Perhaps Yahoo! could find an offer better than Microsoft's, if it held out. But perhaps not. Supposing one wanted to issue an ultimatum deadline such as Ballmer's, how long should it be? What trade-offs must be considered? This paper addresses these issues.

Deadlines appear nearly everywhere in business and in life. Freshly minted MBAs receive offers from investment banks and consulting firms that "explode" on a particular date. Airlines offer limited time deals such as "reduced fares" (or sales) for particular dates between two locations. We refer to such exploding offer scenarios as *ultimatum deadline games*. In these games there are decisions that must be made on both sides of the transaction. The *proposer* (the investment bank, the airline) must decide on a deadline to give, and the *responder* (the candidate, the consumer) must decide whether to accept or reject the offer at the time of the deadline. In some settings, both the proposer and the responder would choose from a number of prospects. Here, we study the case in which the proposer has already settled on a particular responder.

In the ultimatum deadline game, the proposer acts as a Stackelberg leader, setting an ultimatum deadline for her offer. The responder, who may yet receive other offers (according to a stochastic arrival process), must then accept or reject the proposer's offer at the time of the deadline if he has not already accepted another one. The responder must balance the risk (or opportunity costs) of accepting a sure-thing (bird-in-the-hand) offer too soon and missing out on better alternatives that might arise in the future, and the risk of passing up a sure-thing and not finding anything better. We show that the responder's optimal strategy involves a *shortest acceptable deadline* (SAD): If no better (preferred) offer is received prior to the proposer's deadline, the responder should accept any offer with a deadline longer than his SAD and reject any offer with a shorter deadline.

The proposer's optimal deadline maximizes the probability that the responder accepts her offer. When the responder's SAD is known to the proposer, the optimal ultimatum deadline is equal to the responder's SAD. When the proposer is unsure about the responder's situation, she is better off setting longer deadlines (relative to the case where the SAD is known) when her uncertainty is relatively small, and setting shorter deadlines when her uncertainty is substantial.[1] Theoretical

[1] Loosely speaking, if the proposer believes that the SAD is 6 or 8 days, the optimal deadline is 8 days, whereas if she believes it is 2 or 12 days, the optimal deadline is 2 days.

analysis is in Section 3, and Section 4 presents a behavioral experiment in which we examine the behavior of flesh-and-blood players (MBA participants) in an ultimatum deadline game.

2. Setup and Background

For focus — to anchor and aid the narrative — in most of what follows, we will frame the ultimatum deadline game as a hiring problem. Imagine the proposer as a firm making a job offer, and the responder as a job applicant receiving the offer. We suppose that the firm would like to maximize its probability of successfully hiring the applicant, and that the applicant wishes to maximize his expected utility. He is engaged in a job search process that can be represented by a generic sequential search (or optimal stopping) problem. In this section, we review the literature that provides the most important foundation for our work.

2.1. Related Theoretical Work

There is now a sizable literature on the economics of job search. The foundational work of Stigler (1961, 1962) has been extended in a number of ways, and includes both theoretical models and empirical tests of those models using data from actual labor markets (more recent developments are reviewed in Mortensen, 1986). Typically, job search models are formulated as dynamic models of optimal search in the face of uncertainty, i.e., as optimal stopping problems. A basic assumption underlying these models is that the job seeker wishes to maximize some function of his future income stream. Under standard models (e.g., Gilbert and Mosteller, 1966; Freeman, 1983; Gronau, 1971; Lippman and McCall, 1976; Mortensen, 1970; Presman and Sonin, 1972), in each period of the search process the job seeker receives an offer from a single firm with some specified probability (most often with probability one). When an offer is received, it must be accepted or rejected on the spot. Most commonly, the offers are treated as i.i.d. draws from a known wage distribution. Under many formulations of these search problems, the optimal policy for the searcher is represented by a set of decreasing reservation values (Van Den Berg, 1990). For instance, in many finite-horizon formulations, as the search process moves closer to the end of the horizon, the searcher becomes willing to accept progressively poorer alternatives, since the option value of continued search is progressively decreasing. This type of decreasing option value (through time) will play an important role in our model.

Implicit in most job search models formulated as optimal stopping problems is that the firms give the searcher "exploding" offers that must be accepted or rejected on the spot. There are other formulations of optimal stopping problems in which the searcher can recall previously rejected offers (e.g., Karni and Schwartz, 1977). However, the length of the window given to the searcher is exogenous and not treated as a decision variable for the offering firm. In reality, in some markets (e.g., for newly minted MBAs entering consulting), firms often do give job applicants offers with deadlines, but generally the offers do not explode on the spot. Instead, an applicant is given, say, two weeks to make an up or down decision. On the date of the deadline, the applicant, if he has not accepted another position already, can accept the offer, or reject it and wait for a better one, which may or may not come. In some cases, there is room for negotiation, and the applicant may be able to get additional time to make his decision. Here, however, we will focus only on cases in which the deadline is issued as an *ultimatum* — where the decision date is non-negotiable.[2]

Two-sided matching markets bear some similarity to the problem we consider here (Roth and Sotomayor, 1990). In these markets, the sets of firms (proposers) and applicants (responders) are fixed *a priori*. A major objective of this line of work has been to find methods for creating stable matchings (as conceived by Gale and Shapley, 1962). Some actual markets (e.g., the National Resident Matching Program, for new physicians) employ a centralized clearinghouse where firms and applicants are matched simultaneously by a central authority. There are also decentralized markets in which they are matched sequentially.

Markets for clerks for Federal appellate judges, for example, employ such a decentralized matching system. Law school students apply for clerkship positions, and then they can receive offers from the judges themselves. Exploding offers are the norm in these markets, and offers, once made, are generally not rejected (Niederle and Roth, 2007). Importantly, quite often the offers explode on

[2] Baucells and Lippman (2004) examined a somewhat related problem in which a buyer and a seller must negotiate a price, and the seller has the outside option of sequentially searching for a buyer with a higher willingness to pay. One extension of our research is a setting in which the searcher and the firm can bargain over the length of the deadline.

the spot (Avery et al, 2007). The applicants are not given the opportunity to see what other offers might present themselves in the future.[3] One unfortunate consequence of this market design is that it leads to significant unravelling. The market encourages judges to make offers progressively earlier in order to get the best applicants. In some cases the markets start to clear in the Fall after the first year of law school. The result is that judges hire students on the basis of very noisy signals, since they have very little information about their law acumen. (Roth and Xing (1994), Suen (2000), Li and Rosen (1998), and Li and Suen (2004) provide richer descriptions of unravelling in matching markets.)

Our work is distinguished from the work on two-sided matching markets in at least two important ways. First, the matching market literature has focused on settings in which the applicants and the set of open positions become available at the same time (simultaneously). This holds in a number of entry-level labor markets such as those for new physicians, new junior professors, and appellate law clerks. But in many, if not most, markets applicants and positions become available sequentially. For example, the set of administrative assistants looking for employment in Singapore, as well as the set of open positions for them, changes daily. In such markets it is not possible to employ a centralized matching mechanism. Nor can methods based on matching firms and applicants according to their rank-ordered preferences for one another (the most common mechanism employed by two-sided matching algorithms) be used — whether it be done simultaneously or sequentially (e.g., the decentralized procedures used for matching clinical psychologists and positions, Roth and Xing, 1997). This is because the rank orderings are dynamic: They depend on what applicants and positions are present in the market at any one time, as well as on the participants' expectations about what new alternatives might arise in the future. To capture this dynamic feature we suppose that the alternatives (the positions) available to an applicant arrive according to a generic stochastic process akin to those one finds in standard job search problems (as discussed above).

Second, and most importantly, the work on two-sided matching has not treated the length of

[3] One respondent in a survey of applicants for judicial clerkships reported (Avery et al., 2007): "I had ten minutes to respond."

the deadline given to applicants as a decision variable — the nature of the deadline has been imposed *exogenously*. Many markets have their own "culture" that dictates whether exploding offers are permissible or whether offers are open-ended, meaning they do not expire prior to the close of the market, which is at some future date (e.g., April 15 for graduate school admissions). There has been some work comparing the outcomes of markets with exploding offers to those with open-ended offers (e.g., Niederle and Roth, 2007). But there are many deadline possibilities that lie between on-the-spot exploding offers and those that are open-ended. Here, we make the length of the deadline *endogenous* to the problem — we allow the firms to set the time at which the offer explodes. Showing how to set optimal deadlines is our main contribution.

2.2. Related Behavioral Work

From a prescriptive point of view, it is important to understand both the normative and the descriptive aspects of decision making. Normative models deepen our understanding of the ways in which features of the decision problem should influence decision behavior, and they also provide a basis for evaluating behavior, serving as standards by which one can judge the quality of decision making. Quite often, actual decision behavior departs systematically from the dictates of normative, rational models. By understanding the ways in which actual behavior departs from ideal behavior (as given by the models), we can see where decision makers can be helped. Now, we briefly review some relevant behavioral literature, and show that there is good reason to anticipate that normative models might not perfectly predict behavior in ultimatum deadline games.

In addition to the substantial body of theoretical work on optimal stopping problems, there have been a number of studies of the behavior of actual human decision makers in sequential search problems (e.g., Bearden, Murphy, and Rapoport, 2005; Bearden, Rapoport, and Murphy, 2006; Cox and Oaxaca, 1989; Seale and Rapoport, 1997, 2000). This research demonstrates that the optimal models often fail to predict the behavior of actual flesh-and-blood decision makers. For instance, human subjects tend to stop searching sooner than is optimal. Importantly, though, in these problems, the subjects faced games against nature. They were not playing against people.

Other research has shown that whether one plays games against nature or against flesh-and-blood opponents can influence behavior significantly (e.g., Sanfey et al., 2003). The bias to accept early offers might not necessarily occur if the offers are generated by people rather than by chance.

There is good reason to expect that psychological — and not just purely "rational" — considerations come into play in the context of offers with ultimatum deadlines. From the enormous body of literature on the ultimatum game, we know that the standard normative assumptions underlying the game-theoretic predictions do not hold (Bearden, 2001; Roth, 1995). In the standard ultimatum game, the proposer offers the responder some fraction of a fixed pie. If the responder accepts the offer, the pie is split accordingly; if the responder rejects the offer, then neither player earns anything. The subgame perfect equilibrium is that the proposer offers the smallest possible slice of the pie and the responder accepts it. This is because the responder is strictly better off with something rather than nothing, which is what he gets if he rejects the offer. The proposer maximizes her own payoff by offering the smallest slice that the responder will accept — the responder's *minimal acceptable offer*. In contrast to the game-theoretic prediction, typically one finds that small offers (say \$1 from a \$10 pie) are rejected, and that most proposers offer more than the smallest possible slice. One explanation for this is that the human players have fairness concerns that trump pure pecuniary payoffs: Small offers are perceived as unfair and are therefore rejected. The proposers seem to anticipate that.

It is not hard to imagine that offers with very short deadlines might be perceived as unfair, or as bullying. Thus, purely normative considerations that only take into account material payoffs might not provide the best prescriptive basis. One can imagine that a model might dictate a firm giving a very short deadline, and yet find that the offer is rejected because the deadline is seen as unfair by an applicant. Of course, there are any number of psychological reasons why model predictions based on standard assumptions of rationality might fail to provide good predictions of actual behavior. Notions such as the anticipated regret of rejecting an offer and not finding a better one, or pure optimism that "something better will come along" could also play a role in the ultimatum deadline game and the settings it is intended to model. The psychology involved in these settings is undoubtedly rich. The results from our behavioral study, reported in Section 4, provide a straightforward test of the normative (model) predictions. Now, we describe the game and its solution formally.

3. The Ultimatum Deadline Game and Its Solution

We consider a situation where, at time zero, the proposer makes an offer that expires at deadline t_D. The deadline is the point of time at which the offer explodes, and it is issued as an ultimatum (i.e., nonnegotiable). The responder follows a standard finite-horizon search model, looking for other alternatives. If the responder accepts the proposer's offer, the proposer gets payoff one, and if the responder rejects the offer, the proposer gets zero. If the responder accepts the proposer's offer, he gets utility u_O, and if he rejects the proposer's offer, he gets the utility resulting from continuation of the search.

The proposer moves first, by setting deadline t_D, $t_D \geq 0$. If the responder accepts the proposer's offer, he cannot accept any other alternative. Therefore, the responder will hold the proposer's offer until the deadline, and decide whether to accept the proposer's offer at time t_D. At that point, the responder is facing a choice between accepting the proposer's offer and thus receiving u_O or rejecting the proposer's offer and thus receiving the expected utility of continued search.

To summarize the game, the responder must decide whether to accept or reject the proposer's offer at time t_D in order to maximize his expected utility. The proposer must decide which deadline t_D to set in order to maximize the probability that the responder accepts her offer. Our model of the responder's search process is specified in Assumption 1.

Assumption 1

(*i*) *The search process is such that all alternatives have the same utility for the responder,* u_H*,* $u_H > u_O$*, where* u_O *is the value of the proposer's offer to the responder. The responder can accept only one alternative. (Thus, once an alternative is accepted, the search stops.)*

(*ii*) *If the responder ends up with no offer (i.e., the responder rejects the proposer's offer and no alternative from the search comes), the responder's utility is* u_L*,* $u_L < u_O$*.*

(iii) The alternatives with utility u_H arrive stochastically from time zero until time T. This stochastic process is specified by $G(t)$, the probability that at least one alternative arrives during the time window from zero to t.

For example, if alternatives arrive according to a Poisson process with rate λ, the probability of at least one arriving before time t is $G(t) = 1 - e^{-\lambda t}$. For another example, suppose there is only one party that can issue an alternative with utility u_H to the responder. The probability that this will occur is p, and the moment at which the alternative will arrive is drawn from a distribution with cumulative distribution function (hereafter, cdf) $F(t)$ with support $[0, T]$. In this case, $G(t) = pF(t)$.[4] Note that $G(0) = 0$, and $G(T)$ is the probability that the responder gets an alternative with utility u_H if he waits till the very end of the search process. This setting is quite general, since it allows for any pattern of the alternative arrival rate. It is more restrictive than the standard search setting since we assume that all alternatives have exactly the same utility for the responder, u_H. That would make the standard search problem uninteresting, because the searcher's strategy would be simply to accept the first alternative that comes along. However, our focus is on the proposer's strategy. Therefore, the assumption that all alternatives have the same utility seems appropriate, since it allows us to solve the problem analytically while maintaining the main property of search process, namely the decreasing (option) value of continuing the search.

Sections 3.1 and 3.2 consider the responder's and the proposer's strategies under Assumption 1. Section 3.3 discusses generalizations to more standard search processes where the alternatives do not necessarily have the same utility u_H but are drawn from a known distribution.

3.1. The Responder's Strategy

The standard finite-horizon search model is characterized by a reservation value (expected value of continuing search) that is decreasing in time, so that the responder accepts any offer that is above the reservation value and rejects any offer below it. When facing the proposer's offer with

[4] One way to visualize this is to think of time t, drawn from distribution $F(t)$, as the moment when the third party decides about issuing its alternative. Once it decides, the offer will be issued with probability p, and the responder will learn about the decision only if it is positive (i.e., only if he receives the alternative with utility u_H).

deadline t_D, the responder should reject it if the offer is below the reservation value (at time t_D), and accept it otherwise. Proposition 1 gives an analytical solution for the responder's shortest acceptable deadline when Assumption 1 holds.

PROPOSITION 1. *Denote*

$$u = \frac{u_O - u_L}{u_H - u_L}, \tag{1}$$

$$t_D^* = \begin{cases} G^{-1}\left(\frac{G(T)-u}{1-u}\right) & \textit{if } u \le G(T) \\ 0 & \textit{if } u > G(T) \end{cases}, \tag{2}$$

where $G^{-1}(\cdot)$ is the inverse of $G(\cdot)$.[5] *Then the responder should accept the proposer's offer with deadline t_D if and only if $t_D \ge t_D^*$.*

Proof. Suppose the proposer makes an offer with deadline t_D. To compute the expected utility of continuing the search we need to find the probability that a better alternative will arrive after t_D, given that one has not arrived before. By Bayes Theorem, this probability equals $\frac{G(T)-G(t_D)}{1-G(t_D)}$, since $G(T) - G(t_D)$ is the probability of "at least one alternative arriving after t_D and no alternative arriving before t_D," and $1 - G(t_D)$ is the probability of no alternative arriving before t_D. Then the expected utility of continuing the search at time t_D equals $u_H \frac{G(T)-G(t_D)}{1-G(t_D)} + u_L \left(1 - \frac{G(T)-G(t_D)}{1-G(t_D)}\right)$. The proposer's offer will be accepted if and only if $u_O \ge u_H \frac{G(T)-G(t_D)}{1-G(t_D)} + u_L \left(1 - \frac{G(T)-G(t_D)}{1-G(t_D)}\right)$. □

The critical deadline t_D^*, defined in Proposition 1, can be thought of as the shortest acceptable deadline (SAD), in the sense that an offer with deadline shorter than the SAD will be rejected by the responder. The value u, defined in (1), can be interpreted as the relative utility of the proposer's offer. The SAD decreases with u, which in turn implies that it decreases with u_O, and it increases with u_H and u_L. This agrees with intuition: The SAD is shorter when the proposer's offer is more attractive, compared to the prospect of continuing to search. In addition, if the proposer's offer results in a sure payoff to the responder, then the SAD is shorter (longer) for the risk averse (risk seeking) responder, as stated by Observation 1.[6]

[5] By definition, $G(t)$ is increasing. If $G(t)$ is strictly increasing and continuous, then $G^{-1}(y) = \{t : G(t) = y\}$ for $0 \le y \le G(T)$. Otherwise, $G^{-1}(\cdot)$ is defined as $G^{-1}(y) = \sup\{t : G(t) \le y\}$.

[6] In our model, u_O corresponds to the expected utility of accepting the proposer's offer, so the actual payoff might be uncertain.

Observation 1 *Suppose that if the responder accepts the proposer's offer, he gets a sure payoff. Consider responder 1 with utility function for wealth $u_1(w)$ and responder 2 with utility function $u_2(w)$. If responder 2 is more risk averse than responder 1 (i.e., $u_2(w) = \phi(u_1(w))$, where $\phi(\cdot)$ is increasing and concave), then the SAD for responder 2 is shorter.*

To summarize, the responder's strategy is to accept the proposer's offer at the time of the deadline if the offer value is higher than his reservation value (at that time). Since the reservation value is decreasing in time, we can define the shortest acceptable deadline (SAD) as the maximal time at which the reservation value is greater than the value of the proposer's offer, and then the responder's strategy is to accept an offer with deadline longer than SAD and reject an offer with deadline shorter than SAD. To finalize the analysis of the game, we now turn to the proposer's strategy.

3.2. The Proposer's Strategy

First, we consider a situation where the proposer knows the responder's SAD. Then, in Section 3.2.1, we consider the case where the responder can be of one of two types, i.e., where the responder's SAD has two possible values. Section 3.2.2 deals with a more general case of a continuous distribution over the responder's SAD.

If the proposer has all information that is available to the responder (i.e., u_O and the parameters of the search process specified in Assumption 1), then the proposer can rationally derive the responder's SAD. In that case the optimal strategy for the proposer is to set the deadline equal to the responder's SAD (i.e., t_D^* given by (2)). This result is intuitive: If the deadline is shorter than the SAD, the proposer's offer will be rejected. If the deadline t_D is longer than t_D^*, there is a chance that a better alternative will arrive between t_D^* and t_D, and thus the probability that the responder will accept the proposer's offer is lower than if the deadline is set at t_D^*.

Observe that by setting the deadline at $t_D \geq t_D^*$, the proposer's offer will be accepted with probability $1 - G(t_D)$, i.e., the probability that a better alternative does not arrive before time t_D. When $t_D = t_D^*$, given by (2), this probability equals $min\left(1, \frac{1-G(T)}{1-u}\right)$, which is greater than

$1-G(T)$, probability that an alternative with utility u_H does not come at all. This illustrates that the proposer benefits by imposing a deadline, rather than passively waiting. Also, if $G(T) \le u$, where u is defined by (1), the optimal deadline is zero (i.e., on the spot) and thus the proposer's offer is accepted with probability one. This result is also intuitive: If $G(T) \le u = \frac{u_O - u_L}{u_H - u_L}$, the responder prefers getting u_O (i.e., accepting the proposer's offer) to the lottery of receiving u_H with probability $G(T)$ and u_L otherwise, and therefore setting the deadline at zero (i.e., giving take-it-or-leave-it offer) is optimal for the proposer.

3.2.1. The Responder is of One of Two Types

We now consider a situation in which the responder is of one of two types, i.e., the responder's SAD is either t_1 or t_2, $0 \le t_1 < t_2 \le T$, with probabilities π_1 and $\pi_2 = 1 - \pi_1$, correspondingly. For instance, this would be the case if the proposer is uncertain about at least one of the responder's utilities u_L, u_O, u_H, and believes that with probability π_1 these utilities are such that t_D^*, given by (2), equals t_1, and with probability π_2 it equals t_2. We assume that the alternatives arrive at the same rate, i.e., $G(t)$ is the same for both types.[7] Which deadline should the proposer set in this case?

First observe that the optimal deadline t_D should equal either t_1 or t_2: If $t_D < t_1$, the proposer's offer will be rejected for sure, so the proposer would do better by setting deadline t_1. If $t_1 < t_D < t_2$, the proposer again would do better by setting deadline t_1. Finally, if $t_D > t_2$, the proposer is better off by setting deadline t_2. Therefore, the proposer's choice is between setting deadline at t_1 or t_2. Proposition 2 provides conditions for the shorter or longer deadline to be optimal.

Proposition 2. *Suppose that the responder's SAD is either t_1 or t_2, $0 \le t_1 < t_2 \le T$, with probabilities π_1 and $\pi_2 = 1 - \pi_1$, and $G(t)$ is the same for both types. Then:*

1) If the two possible SADs are close enough, the proposer should set the longer deadline. Formally, if $G(t_2) - G(t_1) < \frac{\pi_2}{\pi_1}\big(1 - G(T)\big)$, then the proposer should set the deadline to t_2.

[7] This assumption makes exposition easier, but the results generalize to the case in which the alternatives arrive at different rates, i.e., where the two types have different $G_1(t)$ and $G_2(t)$. This more general setting is discussed in Section 3.3, Proposition 7.

2) If $G(t_2) < \pi_2$ (in particular, if $G(T) < \pi_2$), the proposer should set the longer deadline (for any t_1).

3) If $G(t_2) > \pi_2$ and the deadline t_1 is close enough to zero, the proposer should set the shorter deadline.

Proof. Denote by $Pr(t_D)$ the probability that the responder accepts the offer with deadline t_D. When $t_D = t_1$, the probability that the responder accepts the offer equals $Pr(t_1) = \pi_1(1 - G(t_1))$, i.e., the probability that the responder's SAD is t_1 times the probability that a better alternative did not arrive before t_1. When $t_D = t_2$, the probability that the responder accepts the offer equals $Pr(t_2) = \pi_1(1 - G(t_2)) + \pi_2(1 - G(t_2))$. Then $Pr(t_1) < Pr(t_2)$ if and only if

$$\begin{aligned} \pi_1(1 - G(t_1)) &< \pi_1(1 - G(t_2)) + \pi_2(1 - G(t_2)) \Leftrightarrow \\ \pi_1(G(t_2) - G(t_1)) &< \pi_2(1 - G(t_2)) \Leftrightarrow \\ G(t_2) - G(t_1) &< \frac{\pi_2}{\pi_1}(1 - G(t_2)). \end{aligned} \tag{3}$$

The first statement follows from (3) and $G(t_2) \le G(T)$. The second statement follows from (3) and noticing that $G(t_2) - G(t_1) \le G(t_2)$, since $G(t_1) \ge 0$. Then $G(t_2) < \frac{\pi_2}{\pi_1}(1 - G(t_2))$ is equivalent to $G(t_2) < \pi_2$. The proof of the third statement is similar: If $G(t_2) > \pi_2$, then $Pr(t_1) > Pr(t_2)$ when $G(t_1)$ is close enough to zero. □

The first claim of Proposition 2 states that if the uncertainty about the responder's SAD is small (in the sense that t_1 and t_2 are close to each other), then the proposer should set the longer deadline. The second is that if the probability of the longer SAD (π_2) is greater than the probability of losing the responder by the time of the longer SAD ($G(t_2)$), then setting the longer deadline is optimal. The third states that if the shorter SAD is close enough to zero, then setting the shorter deadline is optimal unless chances that the responder's SAD is the longer one are very high.

Setting the longer deadline implies that the proposer makes her offer acceptable to both types. Proposition 2 shows that doing so is optimal if a) the two types are not very far from each other or b) the probability of a better alternative coming is lower than the probability that the responder's

SAD is the longer one. Focusing on the type with the shorter deadline is optimal only if the probability that the responder's SAD is the longer one is not very high, and the two responder's types are far enough from each other.

To conclude, if the two possible responder's SADs are close enough to one another, the proposer should set the longer deadline in order to make her offer acceptable to either type. If the possible SADs are far enough from each other, the proposer should set the shorter deadline if the probability of the responder having the shorter SAD is not very small. The next section reveals that this intuition extends to a more general case in which the distribution of the responder's SAD is continuous.

3.2.2. Responder's Type is Continuously Distributed

This section considers a situation where the distribution of the responder's SAD is continuous. To simplify the solution, it is convenient to parameterize uncertainty about the responder's SAD via a distribution over a dummy variable z, which can be interpreted as the responder's type. The distribution of the responder's SAD will be specified by mapping from z to SAD. Define $z = \frac{u_H - u_L}{u_H - u_O}$. Observe that $z > 1$, since $u_L < u_O < u_H$. By (1), $z = \frac{1}{1-u}$, and from (2) the responder's SAD is

$$t_D^*(z) = \begin{cases} G^{-1}\big(1 - z(1 - G(T))\big) & \text{if } z \le \frac{1}{1-G(T)} \\ 0 & \text{if } z > \frac{1}{1-G(T)} \end{cases}. \tag{4}$$

Let $H(z)$ and $h(z)$ be the cdf and the probability density function (hereafter, pdf) for z. We assume that the distribution of z has an increasing generalized failure rate, i.e., the ratio $\frac{zh(z)}{1-H(z)}$ is increasing.[8] First, we establish the optimal deadline to be set by the proposer.

Proposition 3. *Suppose the responder's type z, $z > 1$, has distribution with pdf $h(z)$, cdf $H(z)$, and that it has an increasing generalized failure rate. Given a realization of z, the responder's SAD is given by (4). Define z^* as*[9]

$$z^* = min\left\{ z : \frac{zh(z)}{1 - H(z)} \ge 1 \right\}. \tag{5}$$

Then the optimal deadline to be set by the proposer is $t_D^*(z^*)$, given by (4).

[8] Lariviere (2006) provides alternative characterizations of distributions with increasing generalized failure rates.

[9] Denote the lower bound on the support of $h(z)$ by $\underline{z}$, $\underline{z} \ge 1$. Then z^* equals $\underline{z}$ if $\underline{z}h(\underline{z}) \ge 1$, and it equals the solution of the equation $\frac{zh(z)}{1-H(z)} = 1$ otherwise.

Proof. The probability that the responder will accept the offer with deadline t_D is

$$Pr(Accept) = Pr\big(t_D^*(z) \le t_D\big)\big(1 - G(t_D)\big),$$

where $Pr(t_D^*(z) \le t_D)$ is the probability that the responder's SAD is shorter than t_D, and $1 - G(t_D)$ is the probability that a better alternative does not come before deadline t_D. Let $z_D = \frac{1-G(t_D)}{1-G(T)}$. Note that, by (4), $t_D^*(z_D) = t_D$. Then $1 - G(t_D) = (1 - G(T))z_D$, $Pr(t_D^*(z) \le t_D) = Pr(z \ge z_D) = 1 - H(z_D)$, and

$$Pr(Accept) = \big(1 - H(z_D)\big)\big(1 - G(T)\big)z_D.$$

Differentiating with respect to z_D yields

$$-h(z_D)\big(1 - G(T)\big)z_D + \big(1 - H(z_D)\big)\big(1 - G(T)\big) = \big(1 - G(T)\big)\big(1 - H(z_D)\big)\left(1 - \frac{z_D h(z_D)}{1 - H(z_D)}\right).$$

Since the generalized failure rate $\frac{z_D h(z_D)}{1-H(z_D)}$ is increasing, this derivative is positive (negative) for $z_D < z^*$ ($z_D > z^*$), where z^* is defined by (5), and thus $Pr(Accept)$ is maximized at $z_D = z^*$. This concludes the proof. □

If the responder's type z is known to the proposer, a higher value of z implies a shorter SAD (as follows from (4)), and thus a shorter optimal deadline to be set by the proposer. The same result holds in the case of uncertainty, if we compare two distributions of z where one dominates the other in terms of the failure rate.

Proposition 4. *Suppose that* $\frac{h_1(z)}{1-H_1(z)} \le \frac{h_2(z)}{1-H_2(z)}$ *for all* z*, and both* $\frac{zh_1(z)}{1-H_1(z)}$ *and* $\frac{zh_2(z)}{1-H_2(z)}$ *are increasing. The optimal deadline set by the proposer when the responder's type is drawn from the distribution with pdf* $h_1(z)$ *is less than or equal to the optimal deadline when the responder's type is drawn from the distribution with pdf* $h_2(z)$*.*

Proof. Denote by z_i^*, $i = 1, 2$, the solution of (5) for the distribution with pdf $h_i(z)$. From $\frac{h_1(z)}{1-H_1(z)} \le \frac{h_2(z)}{1-H_2(z)}$, increasing generalized failure rates for $h_1(z)$ and $h_2(z)$, and (5), it follows that $z_1^* \ge z_2^*$, and therefore $t_D^*(z_1^*) \le t_D^*(z_2^*)$. □

Remark 1. *Dominance in terms of the failure rate is stronger than first-order stochastic dominance (Shaked and Shanthikumar, 1994). First-order stochastic dominance would not be sufficient for Proposition 4 to hold. Consider the following example, based on Proposition 2. Set* $G(t) = \frac{2}{3}\frac{t}{T}$, *with* $T = 10$. *Let* $h_1(z)$ *be the probability mass function (pmf) such that* z *is either* 2.6 *or* 2.8 *with equal probabilities, and let pmf* $h_2(z)$ *be such that* z *is either* 1.2 *or* 2.8 *with equal probabilities.*

By (4), when the responder's type is drawn from the distribution with pmf h_1, *the SAD is either* 1 *or* 2 *with equal probabilities; when the responder's type is drawn from the distribution with pmf* h_2, *the SAD is either* 1 *or* 9 *with equal probabilities. By Proposition 2 (equation (3)), the optimal deadline is* 2 *for the first distribution and* 1 *for the second. Thus, even though* h_1 *dominates* h_2 *in the sense of first-order stochastic dominance, the optimal deadline under the distribution with pmf* h_1 *is longer.*

Proposition 4 shows that as the offer gets, on average, more attractive to the responder, the optimal deadline gets shorter. Next we consider the impact of uncertainty about the responder's type on the optimal deadline set by the proposer.

Let $H_0(\varepsilon)$ be the cdf of a distribution with an increasing generalized failure rate, standard deviation one, and the lower bound on the support of ε being $-\delta$, $\delta > 0$. The responder's type is $z = \mu + \sigma\varepsilon$, where $\mu > 1 + \sigma\delta$. Then $z > 1$, and the distribution of z has cdf $H(z|\mu,\sigma) = H_0\left(\frac{z-\mu}{\sigma}\right)$ and an increasing failure rate. The standard deviation of distribution $H(z|\mu,\sigma)$ equals σ, and increasing (decreasing) σ corresponds to stretching (contracting) around μ. In that sense σ serves as a measure of uncertainty about z. We now consider how the optimal deadline changes with σ.

Observe that when $\sigma = 0$, the distribution of z is degenerate with the mass point at μ, and then the optimal deadline equals $t_D^*(\mu)$. The next proposition shows that if σ is not large, the optimal deadline is no shorter than $t_D^*(\mu)$, and if σ is large enough, the optimal deadline is no longer than $t_D^*(\mu)$.

Proposition 5. *Denote by* $z^*(\mu,\sigma)$ *the solution to (5) for a distribution with cdf* $H(z|\mu,\sigma) = H_0\left(\frac{z-\mu}{\sigma}\right)$. *The optimal deadline is then* $t_D^*(z^*(\mu,\sigma))$. *If* $\sigma < \frac{\mu h_0(0)}{1-H_0(0)}$, *then* $t_D^*(z^*(\mu,\sigma)) \geq t_D^*(\mu)$. *If* $\sigma > \frac{\mu h_0(0)}{1-H_0(0)}$, *then* $t_D^*(z^*(\mu,\sigma)) \leq t_D^*(\mu)$. *Inequalities are strict if* $t_D^*(\mu) > 0$.

Proof. Observe that $h(z|\mu,\sigma) = \frac{1}{\sigma} h_0\left(\frac{z-\mu}{\sigma}\right)$. By (5),

$$z^*(\mu,\sigma) = min\left\{ z : \frac{\frac{1}{\sigma} z h_0\left(\frac{z-\mu}{\sigma}\right)}{1 - H_0\left(\frac{z-\mu}{\sigma}\right)} \geq 1 \right\}. \tag{6}$$

First, we show that if $\sigma < \frac{\mu h_0(0)}{1-H_0(0)}$, then $z^*(\mu,\sigma) < \mu$, and therefore $t_D^*(z^*(\mu,\sigma)) \geq t_D^*(\mu)$. Suppose that, contrary to the claim, $\sigma < \frac{\mu h_0(0)}{1-H_0(0)}$ but $z^*(\mu,\sigma) \geq \mu$. By (6) and increasing generalized failure rate, it has to be that $\frac{\frac{1}{\sigma} z^*(\mu,\sigma) h_0\left(\frac{z^*(\mu,\sigma)-\mu}{\sigma}\right)}{1-H_0\left(\frac{z^*(\mu,\sigma)-\mu}{\sigma}\right)} = 1$. But, for $z^*(\mu,\sigma) \geq \mu$,

$$\frac{\frac{1}{\sigma} z^*(\mu,\sigma) h_0\left(\frac{z^*(\mu,\sigma)-\mu}{\sigma}\right)}{1 - H_0\left(\frac{z^*(\mu,\sigma)-\mu}{\sigma}\right)} \geq \frac{\frac{1}{\sigma}\mu h_0(0)}{1 - H_0(0)},$$

which is greater than one for $\sigma < \frac{\mu h_0(0)}{1-H_0(0)}$. Therefore, $z^*(\mu,\sigma) < \mu$. Similarly, suppose that $\sigma > \frac{\mu h_0(0)}{1-H_0(0)}$ but $z^*(\mu,\sigma) \leq \mu$. Then

$$\frac{\frac{1}{\sigma} z^*(\mu,\sigma) h_0\left(\frac{z^*(\mu,\sigma)-\mu}{\sigma}\right)}{1 - H_0\left(\frac{z^*(\mu,\sigma)-\mu}{\sigma}\right)} \leq \frac{\frac{1}{\sigma}\mu h_0(0)}{1 - H_0(0)},$$

which is less than one for $\sigma > \frac{\mu h_0(0)}{1-H_0(0)}$, thus contradicting (6). Therefore, $z^*(\mu,\sigma) > \mu$.

Note that inequalities for $z^*(\mu,\sigma)$ are strict, and since $t_D^*(z)$, given by (4), is strictly decreasing for $z < \frac{1}{1-G(T)}$, $z^*(\mu,\sigma) < (>) \mu$ imply $t_D^*(z^*(\mu,\sigma)) > (<) t_D^*(\mu)$ if $\mu < \frac{1}{1-G(T)}$, i.e., if $t_D^*(\mu) > 0$. □

Proposition 5 agrees with Proposition 2. Small uncertainty (i.e., small σ) in the setting with continuous types corresponds to a small difference between two types. In that case, the optimal deadline is longer compared to the case of no uncertainty. Similarly, large uncertainty (i.e., large σ) in the setting with continuous types corresponds to a large difference between two types. Then the optimal deadline is shorter, compared to the case of no uncertainty, provided that $G(T)$, the probability of a better alternative to come, is sufficiently large.[10]

Example 1. *To illustrate our results from this section and from Proposition 5 in particular, let $H(z|\mu,\sigma)$ be uniform on $\left[\mu - \sqrt{3}\sigma, \mu + \sqrt{3}\sigma\right]$, so that μ and σ are the mean and standard deviation*

[10] In the case of two types, the condition for the shorter deadline to be optimal is $G(T) > \pi_2$. For continuous types, the range of μ and σ is such that $\mu - \sigma\delta > 1$. For the optimal deadline to be shorter, μ should be large enough so that $\mu - \sigma\delta > 1$ for $\sigma > \frac{\mu h_0(0)}{1-H_0(0)}$, and $G(T)$ should be large enough so that $\mu < \frac{1}{1-G(T)}$.

of z. *To ensure* $z > 1$ *we assume* $\mu - \sqrt{3}\sigma > 1$, *or* $\sigma \le \frac{\mu-1}{\sqrt{3}}$. *Then* $H_0(z)$ *is uniform on* $\left[-\sqrt{3}, \sqrt{3}\right]$, *with* $\delta = \sqrt{3}$. *By (5),* $z^* = \max\left(\frac{\mu+\sqrt{3}\sigma}{2}, \mu - \sqrt{3}\sigma\right)$, *i.e.,*

$$z^*(\mu,\sigma) = \begin{cases} \mu - \sqrt{3}\sigma & \text{if } \sigma < \frac{\mu}{3\sqrt{3}} \\ \frac{\mu+\sqrt{3}\sigma}{2} & \text{if } \frac{\mu}{3\sqrt{3}} \le \sigma \end{cases}.$$

Assume $\mu < \frac{1}{1-G(T)}$. *By Proposition 3, the optimal deadline to be set by the proposer is*

$$t_D^*(z^*(\mu,\sigma)) = \begin{cases} G^{-1}\big(1 - (\mu - \sqrt{3}\sigma)(1 - G(T))\big) & \text{if } \sigma < \frac{\mu}{3\sqrt{3}} \\ G^{-1}\big(1 - \frac{\mu+\sqrt{3}\sigma}{2}(1 - G(T))\big) & \text{if } \frac{\mu}{3\sqrt{3}} \le \sigma \le \frac{1}{\sqrt{3}}\left(\frac{2}{1-G(T)} - \mu\right) \\ 0 & \text{if } \frac{1}{\sqrt{3}}\left(\frac{2}{1-G(T)} - \mu\right) \le \sigma \end{cases}.$$

Consistent with Proposition 5, the optimal deadline is increasing in σ *until* σ *reaches* $\frac{\mu}{3\sqrt{3}}$, *and then it decreases with* σ. *For* $\sigma > (<) \frac{\mu}{\sqrt{3}}$ *the optimal deadline is shorter (longer) than the one with no uncertainty, i.e., than* $t_D^*(\mu)$.

3.3. Generalizations

The main limitation of our model of the responder's search process, specified by Assumption 1, is that we assume that all alternatives from search have the same utility u_H. A more standard approach is to assume that any alternative that comes is drawn from some distribution. In this section we discuss how our results would generalize or change when we take that into account.

In a standard finite-horizon search model the reservation value of search is decreasing over time (see Van Den Berg, 1990). Therefore, the responder's strategy is still characterized by a shortest acceptable deadline. Proposition 1 extends to this setting, though the expression for the SAD would be more complicated.

We now turn to the analysis of the proposer's strategy. If the responder's SAD is known, is it the optimal deadline for the proposer to set? Setting a deadline shorter than the responder's SAD is not optimal, since in this case the offer will be rejected for sure. Setting a deadline t longer than the responder's SAD is not optimal if the probability that the responder will accept the proposer's offer is decreasing in t, for t greater than the SAD. That seems intuitive, since increasing the deadline from t_1 to t_2 gives the responder a chance to accept an alternative resulting from his search process between t_1 and t_2. For the standard search process with i.i.d. draws that intuition is correct, as the following proposition shows.

Proposition 6. *Suppose the responder follows the standard search process (see, e.g., Gilbert and Mosteller, 1966) with T sequential i.i.d. draws, and has an offer from the responder that expires at time t, $t \geq SAD$. Then the probability that the responder will accept the proposer's offer is decreasing in t.*

Proof. Consider two deadlines, t_1 and t_2, with SAD $\leq t_1 < t_2$. When the deadline equals t_1, the responder faces the standard search problem with t_1 draws (since he will stop searching at t_1 and accept the proposer's offer). When deadline is t_2, the responder faces the standard search problem with t_2 draws. If we assume that, when deadline is t_2, he just ignores the first $t_2 - t_1$ draws, the situation is identical to the one with deadline t_1. Since there is positive probability that the responder will get an alternative above the reservation value in the first $t_2 - t_1$ draws, the probability that the responder will accept the proposer's offer with deadline t_2 is lower, compared to the case where the deadline is t_1. (On a side, note that when deadline is t_2, the reservation value at time t is the same as the reservation value at time $t - t_2 + t_1$ in the game with deadline t_1.) □

Proposition 6 is formulated for the search process with finite number of i.i.d. draws. The proof is similar for the case where the search process is formulated in continuous time, with alternatives arriving sequentially according to a Poisson process (Cowan and Zabczyk, 1978). However, the assumption of i.i.d. draws is crucial here. As Example 2 illustrates, when the search process involves non-i.i.d. draws, it might be the case that the proposer would prefer to set the deadline longer than the responder's SAD.

Example 2. *Suppose the (risk neutral) responder's search model is the following: He expects to see alternatives in both periods* 1 *and* 2*. The alternative in period 1 will pay* \$2 *with probability* 0.9 *and* \$0 *with probability* 0.1*. The alternative in period 2 will pay* \$10 *with probability* 0.1 *and* \$0 *with probability* 0.9*. The value of the proposer's offer is* \$1.5*. Should the proposer set deadline* 1 *or* 2*? The shortest acceptable deadline in this situation is* 1*: At the end of period* 1*, the responder will accept the proposer's offer, since the expected value of continuing search (i.e., the reservation value) is* $\$10 \times 0.1 < \$1.5$*. When deadline* $t_D = 1$*, the proposer's offer will be accepted with probability* 0.1

— probability that the responder did not receive an alternative of \$2 *in the first period. Can the proposer do better by setting the longer deadline? When deadline* $t_D = 2$*, the reservation value in the first period becomes* $\$10 \times 0.1 + \$1.5 \times 0.9 = \$2.35$*, which is greater than* \$2*, the highest possible value of the alternative in the first period. Therefore, the responder will reject the first-period alternative, and now the proposer's offer will be accepted with probability* 0.9 *— the probability that the responder did not receive an alternative of* \$10 *in the second period.*

The intuition behind this example is the following: By increasing the deadline, there are two counteracting effects. One is that it gives the responder more time to see the alternatives coming from the search process, which is bad for the proposer. The second is that it increases the responder's reservation value before the deadline is reached, which is good for the proposer. As Proposition 6 shows, the first effect dominates in the case of i.i.d. draws. In Example 2, the second effect dominates.

To conclude, our result that if the responder's SAD is known, the proposer should set a deadline equal to the SAD generalizes to a standard search process with i.i.d. draws of alternatives. However, if the alternatives do not have the same distribution (as in Example 2), this dictate might not hold.

We now consider a situation in which the responder is of one of two types i, $i = 1, 2$, appearing with probabilities π_1 and $\pi_2 = 1 - \pi_1$, respectively. For type i, let t_i denote its shortest acceptable deadline and let $Pr_i(t)$, $t \geq t_i$, denote the probability that the responder will accept the proposer's offer with deadline t. We assume that $Pr_i(t)$ is decreasing. (By Proposition 6, this is the case for a standard search process.) As Proposition 7 below shows, Proposition 2 essentially extends.

Proposition 7. *1) If the SADs* t_1 *and* t_2 *are close enough, the proposer should set the longer deadline. More specifically, if* $Pr_1(t_1) - Pr_1(t_2) < \frac{\pi_2 Pr_2(t_2)}{\pi_1}$*, then the optimal deadline is* t_2*.*

2) If $\frac{\pi_2 Pr_2(t_2)}{\pi_1} > 1$*, the proposer should set the longer deadline.*

3) If the SADs t_1 *and* t_2 *are far enough apart, so that* $Pr_1(t_1) - Pr_1(t_2) > \frac{\pi_2 Pr_2(t_2)}{\pi_1}$*, the proposer should set the shorter deadline.*

Proof. When $t_D = t_1$, the probability that the responder accepts the offer equals $\pi_1 Pr_1(t_1)$, i.e., the probability that the responder is of type 1 times the probability that the responder will still be available at time t_1. When $t_D = t_2$, the probability that the responder accepts the offer equals $\pi_1 Pr_1(t_2) + \pi_2 Pr_2(t_2)$. Then setting the deadline to t_1 is better than setting the deadline to t_2 if and only if

$$\pi_1 Pr_1(t_1) < \pi_1 Pr_1(t_2) + \pi_2 Pr_2(t_2) \Leftrightarrow$$
$$Pr_1(t_1) - Pr_1(t_2) < \frac{\pi_2 Pr_2(t_2)}{\pi_1}.$$

□

The first claim of Proposition 7 states that if the uncertainty surrounding the responder's SAD is small (in the sense that t_1 and t_2 are close to each other), then the proposer should set the longer deadline. The second states the condition under which setting the longer deadline is always optimal. The third is inverse to the first claim, and states when setting the shorter deadline is optimal. These statements agree with Proposition 2.

To summarize, the theoretical analysis suggests that the responder's strategy is to accept the proposer's offer if the deadline is longer than the responder's shortest acceptable deadline (SAD) and to reject the offer if the deadline is shorter than the SAD. If the responder's SAD is known to the proposer, the SAD itself is the optimal deadline for the proposer to set. If the proposer is uncertain about the responder's SAD, she should set a longer deadline unless the degree of uncertainty is very large (Proposition 2 and Proposition 5). These conclusions also hold in more general search settings (Proposition 6 and Proposition 7). We now turn to the experimental (descriptive) analysis of this game.

4. A Behavioral Experiment

4.1. Overview and Method

We examined decision behavior of both responders and proposers in the ultimatum deadline game. In total, 142 subjects participated in the study. Sixty-nine were assigned randomly to the role of responder and 73 to the role of proposer. The experimental subjects were recruited from the MBA participants at INSEAD. They were paid based on their performance; they did not participate to fulfill a course requirement. In addition to a fixed S$5 showup fee, subjects in the responder role earned an average of S$21 for the 15 minute experiment, while those in the proposer role earned an average of S$8.

The subjects were recruited through an MBA core course. The students were quite diverse, representing a range of nationalities, ages, and backgrounds. In class, they were presented with a brief description of the ultimatum deadline game and the protocols for taking part in the online experiment. Then, each of them received a personalized email invitation to participate in the study. The email contained a link to the online software for playing the game, a notification of whether the student would play the role of proposer or responder, and a personalized password to access the experimental software.

Once logged into the experimental site, each subject was provided with extensive written instructions describing the task and the way in which payoffs would be determined. Subjects were told that proposers and responders would be matched randomly at the conclusion of the experiment, and then paid according to the rules of the game once everyone had participated, which was expected to be — and was — about one week later. To aid intuition, the problem was described as a job search problem.

Subjects in *both* roles were provided with the payoff contingencies. Specifically, they were told that the responder could earn S$40 if he received and accepted a better alternative ($u_H = \text{S\$40}$), S$10 if he accepted the proposer's offer ($u_O = \text{S\$10}$), and that he would earn nothing otherwise ($u_L = \text{S\$0}$). The proposer would earn S$40 if the responder accepted her offer.

The experimental "season" lasted for a period of 10 days ($T = 10$). Subjects in both roles were provided with a table showing the probabilities of the responder receiving a better alternative *on or before* each day t, and another one showing the probabilities of receiving a better one *after* each day t given that he has not received one sooner. These probabilities are displayed together in Table 1.[11] As an example, the probability that a responder will receive a better alternative *on or*

[11] These probabilities correspond to $G(t) = 0.5t/T$.

before day 7 is 0.35, and the probability that he will receive a better alternative *after* day 7 (given no alternative on or before day 7) is 0.23. Given these probabilities, the expected value maximizing SAD for the responders is 7. This is the point at which expected value of waiting for a better alternative dips below (to S$9.20) the value of the sure thing offer from the proposer (S$10).

Table 1 Probabilities for the behavioral experiment.

	t									
	1	2	3	4	5	6	7	8	9	10
Pr(Alt. on or before t)	5%	10%	15%	20%	25%	30 %	35%	40%	45%	50%
Pr(Alt. after t given no alt. yet)	47%	44%	41%	38%	33%	29 %	23%	17%	9%	0%

The subjects in the responder role were asked to provide their shortest acceptable deadline, and those in the proposer role were asked to set their deadline for the responder. These values were input using a slider that could be set to any integer value between 1 and 10 days, inclusive. Once these were input, the subjects were asked the following question: "What strategy did you use for making this decision?" Next, the responders were asked: "What do you think the chances are that [the proposer] is going to give an offer with a deadline shorter than your shortest acceptable deadline?" Similarly, the proposers were asked: "What do you think the chances are that the [responder] will accept your offer?" Responses were input using a slider that could be set between 0% and 100%, inclusive.

Once all the data were collected, for each experimental responder whether and when a better alternative came was resolved precisely according to the probabilities provided to the participants. Each responder was then randomly paired with one proposer, and payoffs to each subject in each pair were determined according to the rules of the game. (A small number of proposers had to be paired with multiple responders because of unequal numbers. In these cases, we simply used the first responder to determine the proposer subject's payoff.)

In short, the game played by the experimental subjects was exactly the ultimatum deadline game described in Section 3, and the game was played under incentive compatible payoffs (i.e., with real money).

4.2. Results

Responder SADs and Proposer Deadlines. The distributions of the responders' shortest acceptable deadlines (left panel) and of the deadlines set by proposers (right panel) are shown in Figure 1. Notice first that there is considerable heterogeneity in the responders' shortest acceptable deadlines.[12] The mean SAD was 6.97 (SD = 2.01), and they ranged from 1 to 10. Recall that the expected value maximizing SAD for the game used in the experiment is 7. There was no significant difference between the mean SAD and the expected value maximizing SAD, $t(68) = 0.12$, $p = 0.90$.

Likewise, we find considerable variability in the deadlines set by the subjects in the proposer role, with a mean deadline of 6.00 (SD = 1.81). Assuming the proposers would have faced expected value maximizing responders, their deadlines would have been too short on average – since the optimal deadline would have been 7 (the same as the EV maximizing responder's SAD), $t(72) = 4.72$, $p < 0.001$. However, we have already seen that this is an inappropriate criterion to evaluate the quality of the proposers' deadlines, since the responders did not behave like EV maximizers. Instead, we must consider the *actual* distribution of the responders SADs in order to decide on the proposers' payoff maximizing strategy. Figure 2 shows the expected payoffs to a proposer as a function of its deadline, given the responders' SADs. The subjects in the proposers role would have maximized their expected payoffs by setting a deadline of 10 (the longest possible deadline). Thus, the experimental proposers set deadlines that were too short (and significantly so, $t(72) = 18.88$, $p < 0.001$). Their payoffs suffered considerably: On average, they earned only 40% of the payoffs they could have earned by setting their deadlines optimally (i.e., to 10).[13]

Probability Estimates. Recall that the responders estimated the probability that the proposer would give an offer with a deadline *shorter* than their SAD. Figure 3 (left panel) shows the mean probability estimate for the responders as a function of their SAD (solid line), and also the actual (or

[12] The subjects' verbal descriptions of their strategies for settings SADs and deadlines revealed heterogeneity in strategic reasoning as well.

[13] This penalty for deviating from optimality is substantially greater than those typically found in similar behavioral experiments, such as those on optimal stopping cited above. Due to the "flat maximum" problem, in other studies, we have found that behavioral biases tend to produce no more than a 10% payoff penalty (e.g., Bearden, Rapoport, and Murphy, 2006). The payoff losses to the proposers in the current study are quite substantial.

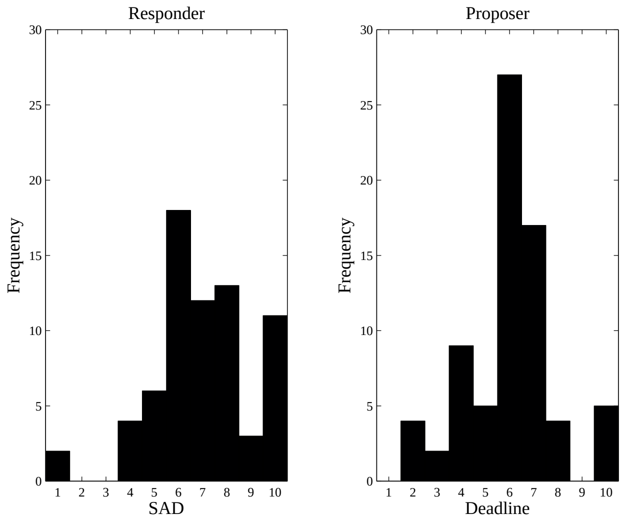

Figure 1 Distributions of SADs set by responders (left panel) and deadlines set by proposers (right panel) from the behavioral experiment.

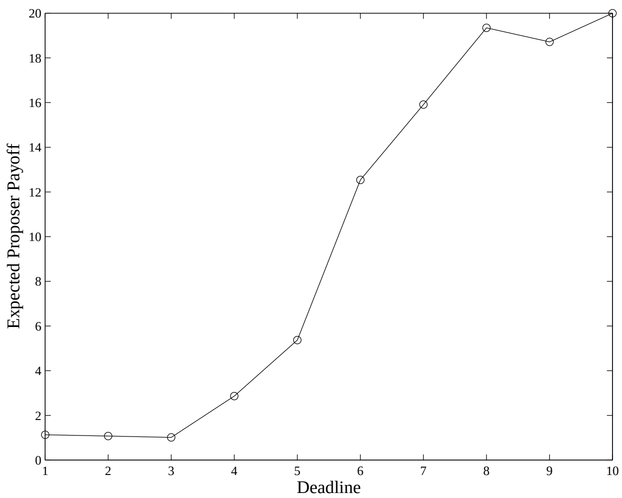

Figure 2 Expected payoff to an experimental proposer as a function of its deadline, given the actual distribution of SADs given by the experimental responders.

true) probabilities (dashed line). Interestingly, responders with shorter SADs (< 7) overestimated the probability that they would reject the proposer's offer, whereas those with longer SADs (≥ 7) underestimated the probability that they would reject the offer. Similarly, the subjects in the proposer role were asked to estimate the probability that their deadlines would be acceptable. From Figure 3, right panel (together with Figure 1, which provides a basis for weighting the estimates for each deadline), we can see that there was a general tendency to overestimate the probability

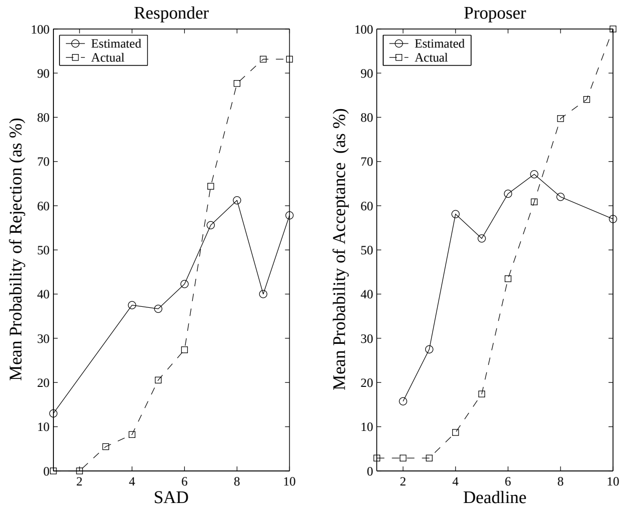

Figure 3 Mean estimated and actual probabilities by responders (left panel) and proposers (right panel).

that the deadlines would be acceptable to the responders.

4.3. Summary of Behavioral Experiment

Our experimental study of the ultimatum deadline game produced several important findings. First, we find that responders' shortest acceptable deadlines are indeed heterogeneous: They do not all set the same one. Thus, the assumption that the responder's type is (or could be) known is violated empirically. Therefore, proposers must take into consideration the variability in SADs in order to set suitable deadlines. Second, the proposers tended to set deadlines that were considerably shorter than they ought to have been in order to maximize their expected payoffs. (Note that the proposers' payoffs are maximized by maximizing the probability that the responders will accept their offers; so risk aversion of proposers, for instance, cannot account for the shorter deadlines.) Based on the probability estimate data, we found that responders tended to expect that proposers would set longer deadlines than they in fact did, and proposers tended to expect that responders would set shorter SADs than they actually did. The net result was a large number of rejected offers. That is, behavioral biases caused many responders to end up without offers and many proposers to fail to get their offers accepted.

5. Conclusions

Most contract offers, including employment offers, have expiration dates. At one extreme is a take-it-or-leave-it (or immediately exploding) offer where the deadline can be thought of as zero (on the spot). At the other extreme is an offer that is valid for quite a long time, until the responder makes up his mind (for instance, such offers are quite common for big-shot hires, both in academia and industry). In many other instances, the deadline would fall somewhere in between. Steve Ballmer's three-week deadline for Yahoo! is one example. A one-week deadline for a McKinsey offer to a new MBA is another.

When setting the deadline, the proposer would like to minimize the chances that the offer will be rejected. The responder may reject the offer for two reasons: First, he can already have a better alternative. Second, he might decide that it is better to reject the proposer's offer and hope for a better one to come later. The proposer's challenge is to balance these two factors. Giving a take-it-or-leave-it offer will minimize the probability that the responder will find a better alternative before the deadline, but unless the responder values the proposer's offer very highly, he will reject it in hopes of finding something better. On the other hand, setting a very long deadline gives the responder enough time to explore other alternatives and thus will minimize the probability that the responder will reject the offer hoping for something better. However, doing so increases the probability that the responder will find a better alternative before the proposer's deadline.

In our model the responder is searching for other alternatives following a finite-horizon search process. This horizon might be determined by market conditions (e.g., the job market clears by March) or might arise from the responder's own commitments (to find a job before January, to sell a house within one month). At the deadline, the responder will accept the proposer's offer only if its value exceeds the expected value of continuing the search. Since the expected value of continuing the search is decreasing in time, the responder's strategy is characterized by a shortest acceptable deadline (SAD): Accept (reject) the proposer's offer if its deadline is longer (shorter) than SAD. This result is formalized in Proposition 1.

This responder's strategy implies that, in general, neither very short (take-it-or-leave-it) nor very long deadlines (e.g., equal to the horizon of search process) are optimal for the proposer. A take-it-or-leave-it offer will be rejected unless the responder is quite pessimistic about his (potential) outside alternatives. An overly long deadline is not optimal because an offer with a shorter deadline (but longer than the SAD) would also be accepted by the responder, but the longer deadline gives the responder more time to find a better alternative. We show that indeed the proposer should set the deadline equal to the responder's SAD, if this SAD is known to the proposer. This conclusion is also valid for a standard search process in which the coming alternatives are i.i.d. (The details are in Proposition 6 and Example 2.)

In most instances, the proposer would not know the responder's SAD exactly, but would have some guesses about the responder's potential alternatives (resulting from his search process) and thus about the responder's SAD. We address this issue in two settings: First, we assume that the responder's SAD can take only one of two values (Propositions 2 and 7). We show that if two possible SADs are not very far from one another (put differently, if the uncertainty about the responder's SAD is not very large), it is optimal to set the deadline equal to the longer SAD. If deadlines are far from each other, we establish the conditions for the shorter or the longer deadline to be optimal. Next, we assume that the responder's SAD is drawn from a continuous distribution (Section 3.2.2). In this setting the conclusions are similar: When uncertainty about the responder's SAD is not very substantial, it is better to set a longer deadline, compared to the case of no uncertainty. Intuitively, this result follows from the fact that (under mild uncertainty) it is better to make sure that the responder does not reject the offer because he hopes to find a better alternative (i.e., because the offer's deadline is shorter than the responder's SAD).

Our theoretical analysis yields two main takeaways: First, we would expect the optimal deadline to be longer than zero (i.e., take-it-or-leave-it offers are not optimal) and shorter than the horizon of the responder's search process (i.e., it is suboptimal to allow the responder to examine all potential alternatives before making a decision). Second, if the responder's SAD is unknown, it is better to set a deadline longer than if it were known. (Formally, Proposition 5 assumes one-to-one relationship between the responder's SAD and type z, and shows that if some noise is added to the responder's type, the optimal deadline is longer than it would be without noise.) We test these prescriptions in the behavioral experiment.

In the behavioral experiment we find that the *average* SAD set by the responders equals the SAD that would be optimal if a responder is maximizing expected payoff. However, there is considerable variability in the responders' SADs. Given the variability, our model suggests that the proposers should set a longer deadline than the average SAD of the responders. That was not the case: The average deadline set by the proposers was *shorter* than the average SAD of the responders. Non-optimal deadlines seem to be driven by mis-calibrated beliefs about how opponents' are likely to behave: In aggregate, proposers underestimated the chances that their deadlines would be rejected.

One reasonable inference to draw from the experimental data is that the proposers underestimated both the mean and the variance of the SAD distribution. Thus it seems that our theoretical results have important prescriptive value: If one is not sure about the responder's reaction, it is safer to set a deadline that is longer than one's (best point) estimate of the responder's SAD. Hedging in the direction of longer deadlines is a safer policy.

The focus of this paper is on advising proposers on setting deadlines. Our model captures the main tradeoff when setting the deadline, namely the balance between the probability that the responder will reject the offer in hope of getting better alternatives later and the probability that the responder will find a better alternative prior to the deadline being reached. There are many directions in which the model can be extended, and many related issues to be explored. For example, the value of the offer can be a decision variable for the proposer. It is not in many job market situations, e.g., in academic job markets for the rookies, because in many instances salary for a particular position is nonnegotiable. Of course, in many settings, e.g., Microsoft's bid for Yahoo!, deciding on the value of the offer would also be important. Another extension could include multiple responders, with potentially different utilities for the proposer, where in addition to the offer deadline, the proposer would also need to decide which responder to approach first.

However, we would expect that the tradeoff between longer and shorter deadline established here would be an important factor in those extended settings as well.

It is hard to resist noting that Yahoo! did not respond to Microsoft's ultimatum. Though we are not sure, it is quite plausible that Yahoo! wanted more time to explore its alternatives with Google and AOL. Anticipating that, perhaps Ballmer should have set a longer deadline.

Acknowledgments

We thank Stephen Chick, Horacio Falcao, Paul Kleindorfer, Claudio Mezzetti, and Robert Winkler for helpful comments. This research has been partially funded by the INSEAD Alumni Fund (IAF).

References

Avery, C., C. Jolls, R. A. Posner, A. E. Roth. 2007. The new market for Federal judicial law clerks. *University of Chicago Law Review* **74** 447–486.

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.

Bearden, J. N. 2001. Ultimatum bargaining experiments: The state of the art. Working paper.

Bearden, J. N., R. O. Murphy, A. Rapoport. 2005. A multi-attribute extension of the secretary problem: theory and experiments. *Journal of Mathematical Psychology* **49** (5) 410–422.

Bearden, J. N., A. Rapoport, R. O. Murphy. 2006. Sequential observation and selection with rank-dependent payoffs: an experimental study. *Management Science* **52** (9) 1437–1449.

Cowan, R., J. Zabczyk. 1978. An optimal selection problem associated with the Poisson process. *Theory Probab. Appl.* **23** 584–592.

Cox, J. C., R. L. Oaxaca. 1989. Laboratory experiments with a finite-horizon job-search model. *Journal of Risk and Uncertainty* **2** (3) 301–329.

Freeman, P. R. 1983. The secretary problem with its extensions: a review. *International Statistical Review* **51** (2) 189–206.

Gale, D., L. S. Shapley. 1962. College admissions and the stability of marriage. *American Mathematical Monthly* **69** 9–14.

Gilbert, J. P., F. Mosteller. 1966. Recognizing the maximum of a sequence. *American Statistical Association Journal* 35–73.

Gronau, R. 1971. Information and frictional unemployment. *The American Economic Review* **61** (3) Part 1. 290–301.

Karni, E., A. Schwartz. 1977. Search theory: the case of search with uncertain recall. *Journal of Economic Theory* **16** (1) 38–52.

Lariviere, M. 2006. A Note on Probability Distributions with Increasing Generalized Failure Rates. *Operations Research*, 54(3) 602-604.

Li, H., S. Rosen. 1998. Unraveling in matching markets. *American Economic Review* **88** (3) 371–387.

Li, H., W. Suen. 2004. Self-fulfilling early-contracting rush. *International Economic Review* **45** (1) 301–324.

Lippman, S. A., J. J. McCall. 1976. Job search in a dynamic economy. *Journal of Economic Theory* **12** (3) 365–390.

Mortensen, D. T. 1970. Job search, the duration of unemployment, and the Phillips curve. *The American Economic Review* **60** (5) 847–862.

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

Niederle, M., A. E. Roth. 2007. Making markets thick: designing rules for offers and acceptances. Working Paper.

Presman, E. L., I. M. Sonin. 1972. Equilibrium points in a game related to the best choice problem. *Theory Probab. Appl.* **20** 770–781.

Roth, A. E. 1995. Bargaining experiments. *Handbook of Experimental Economics.* J. Kagel and A. E. Roth, eds. Princeton University Press, Princeton, NJ, 253–348..

Roth, A. E., M. A. O. Sotomayor. 1990. *Two-Sided Matching.* Cambridge University Press, Cambridge, UK and New York.

Roth, A. E., X. Xing. 1994. Jumping the gun: Imperfections and instituitions related to the timing of market transactions. *The American Economic Review* **84** (4) 992–1044.

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** 284–329

Sanfey A. G., J. K. Rillig, J. A. Aronson, L. E. Nystrom, J. D. Cohen. 2003. The neural basis of economic decision-making in the ultimatum game. *Science* **300** 1755–1758.

Seale, D. A., A. Rapoport. 1997. Sequential decision making with relative ranks: an experimental investigation of the "secretary problem". *Organizational Behavior and Human Decision Processes* **69** (3) 221–236.

Seale, D. A., A. Rapoport. 2000. Optimal stopping behavior with relative ranks: the secretary problem with unknown population size. *Journal of Behavioral Deicsion Making* **13** (4) 391–411.

Shaked, M., J.G. Shanthikumar. 1994. *Stochastic Orders and Their Applications.* Academic Press, San Diego.

Stigler, G. J. 1961. The economics of information. *The Journal of Political Economy* **69** (3) 213–225.

Stigler, G. J. 1962. Information in the labor market. *The Journal of Political Economy* **70** (5) Part 2: Investment in Human Beings. 94–105.

Suen, W. 2000. A competitive theory of equilibrium and disequilibrium unraveling in two-sided matching. *RAND Journal of Economics* **31** (3) 101–120.

Van Den Berg, G. J. 1990. Nonstationarity in job search theory. *Review of Economic Studies* **57** 255–277.

Europe Campus
Boulevard de Constance,
77305 Fontainebleau Cedex, France
Tel: +33 (0)1 6072 40 00
Fax: +33 (0)1 60 74 00/01

Asia Campus
1 Ayer Rajah Avenue, Singapore 138676
Tel: +65 67 99 53 88
Fax: +65 67 99 53 99

www.insead.edu

INSEAD

The Business School
for the World