

Submitted to *Management Science*

Beyond Averages: How Do Customers Respond to Wait Time Distributions?

Evgeny Kagan

Carey Business School, Johns Hopkins University, ekagan@jhu.edu, <https://carey.jhu.edu/faculty/faculty-directory/evgeny-kagan-phd>

Kyle Hyndman

Naveen Jindal School of Management, University of Texas at Dallas, KyleB.Hyndman@utdallas.edu, <https://www.kyle-hyndman.com>

Andrew M. Davis

Samuel Curtis Johnson Graduate School of Management, SC Johnson College of Business, Cornell University, adavis@cornell.edu, <https://andrew-davis.weebly.com/>

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and are not intended to be a true representation of the article's final published form. Use of this template to distribute papers in print or online or to submit papers to another non-INFORM publication is prohibited.

Abstract. Many service offerings combine two attributes: price and waiting time. Models of customer choice in this setting usually rely on utility maximization based on expected waiting time and monetary value of the service. We conduct a series of incentive-compatible human-subjects experiments to examine whether customers also consider distributional information about the wait, beyond just the expected waiting time. In the experiments, participants made choices between two time-money bundles: a short wait with a small monetary payoff and a longer, potentially stochastic wait with a larger monetary payoff. We experimentally varied the length, distribution, and information available about the waiting times. Our results show that focusing on average waiting times alone oversimplifies behavior. Decision-makers showed aversion to all four moments of waiting time distributions, with kurtosis having a particularly strong effect. Further, the value of receiving probabilistic information about a wait depended on the complexity of the distribution. For simple binary distributions, ambiguous waits (where the probabilities of duration outcomes are unknown) were evaluated similarly compared to uncertain waits (where probabilities are known). However, for more complex distributions, sharing probabilistic information increased service valuations. Finally, when given the choice, decision-makers were most interested in and responsive to right-tail information about wait durations. Together, these findings provide new behavioral micro-foundations for modeling customer decisions in service systems and offer important implications for service design in practice.

Key words: Service Design, Behavioral Queueing, Experiments

History: July 18, 2024

1. Introduction

Waiting is something we do daily, whether standing in line at a coffee shop, awaiting a driver for pickup at the airport, or sitting in a doctor's waiting room. While some settings, such as government services, offer

no choice but to wait, others show significant variation across service providers. In banking, certain credit card companies advertise direct and immediate access to customer support, while others require lengthy waits to speak to an agent. In healthcare, immediate access to a physician often comes at a high cost, while less expensive options typically involve longer waits. Often, the same provider may offer a range of waiting formats with different price points. For example, ride-sharing services offer customers a choice between expedient yet expensive, and slower, more affordable rides. Food delivery apps and e-commerce platforms charge extra for priority delivery, often with time guarantees. Theme parks allow guests to purchase fast passes to bypass lines. The common theme across all these scenarios is that the customer must evaluate the monetary cost of faster service against the time cost of a longer wait, where the duration of the wait is often not known ex-ante and the waiting time distribution is not fully specified. In this paper, we examine this decision using a series of incentivized human-subject experiments.

1.1. Background

Customer choices that involve waiting are at the center of many queueing models (Naor 1969, Hassin and Haviv 2003, Hassin 2016). In Naor’s classic model, customers observe the length of the line and join the service system if the value R provided by the service exceeds the cost of waiting, which consists of the price p (e.g., admission fee) and the unit waiting cost c , applied to the duration of the wait t . The value of t is not known ex-ante, but customers can compute its expected value ($\mathbb{E}[t]$) based on the length of the line currently observed and the service rate at which customers are processed. In a more general setting where a queue length may be unobservable, customers can still compute expected waiting times based on prior experience or based on the information shared by the service provider (Edelson and Hilderbrand 1975). In either case, customer utility is typically expressed as:

$$U(R, p, c, t) = R - p - c\mathbb{E}[t], \quad (1.1)$$

The utility function in eq. (1.1) serves as a building block for numerous models in the queueing literature (Hassin and Haviv 2003, Hassin 2016). Some recent models have also started incorporating more realistic assumptions about customer behavior, by considering how customers make inferences about R (Veeraraghavan and Debo 2009, Debo and Veeraraghavan 2014, Kremer and Debo 2016), and how they respond to features of the waiting time distribution beyond just the average (e.g., Afèche et al. 2013, Yang et al. 2018, Debo et al. 2023). These models incorporate risk preferences, loss aversion, and other cognitive and psychological constructs that have been shown to explain behavior in other (non-time related) domains (Kagel and Roth 2020). Despite these advancements in modeling customer behavior, the empirical evidence that would justify deviating from the conventional model in eq. (1.1) remain quite limited (see Allon and Kremer

2018, Ingolfsson et al. 2023, for an overview of the relevant behavioral work).¹ We seek to address this gap by developing new behavioral micro-foundations for models of service processes that involve waiting. This includes examining the response to different types of waiting time information in realistic environments, where the wait may be uncertain, where the waiting time distributions may have different shapes, and where the information about the wait may be incomplete.

1.2. Experimental Approach and Preview of Results

To answer our research questions, we conduct a series of incentive-compatible online human-subject experiments that, in total, include 1045 participants. Participants in our experiments make repeated binary choices between two service options: Option A, which consists of a smaller monetary payoff with a short, deterministic wait, and Option B which generates a larger monetary payoff but requires a longer, potentially stochastic wait. These options are chosen to mirror customer choice between two service providers with different service offerings, or a single service provider offering a choice between a more expensive, expedited/priority service with a fixed or guaranteed service time (Option A), and a less expensive service option with an uncertain service time (Option B).

Formally, service option $j \in \{A, B\}$ is defined by a tuple $\{m^j, t^j\}$. Let $m^j \equiv R - p^j$ represent the net monetary value of the service. For example, in the ride-sharing context where the customer chooses between an expedited and a regular pickup option, R is the intrinsic value of a ride, p^A (resp. p^B) is the price of an expedited (resp. regular) pickup option, and m^A (resp. m^B) is the net monetary value of each option. We will focus on settings where $m^A \leq m^B$ (i.e., $p^A \geq p^B$), and where t^A is a constant, while t^B follows some non-degenerate probability distribution, $f(t^B)$ with positive support. We will focus on settings where $t^A \leq \mathbb{E}[t^B]$. That is, Option A corresponds to a lower payoff but involves only a short wait with a guaranteed duration, while Option B corresponds to a higher expected payoff but involves a longer, potentially stochastic wait.

To fix ideas, consider a utility function that is defined as an additive response to the mean, variance, skewness and kurtosis of a waiting time distribution.² In Appendix B.5 we discuss the rationale for using this modeling approach in our setting.³ Denote the mean, variance, skewness and kurtosis of the waiting

¹ Recent experimental literature has focused on strategic interactions between customers in queues (Ülkü et al. 2020, 2022, Estrada Rodriguez et al. 2023, Ibrahim et al. 2024, Frazelle and Katok 2024), rather than on our focal question of how customers process waiting time information. Additionally, several studies have considered buy-now-or-wait decisions and their implications for (promotion) pricing (Özer and Zheng 2012, Osadchiy and Bendoly 2015, Özer and Zheng 2016, Baucells et al. 2017). These studies focus on risk preferences with respect to the availability of the product (as opposed to the uncertainty in the time spent waiting), and consider much longer time intervals (typically, multiple weeks). In contrast, we consider shorter time intervals (minutes or hours) that are more reflective of classic service settings, such as ride-hailing, food delivery, or amusement parks.

² With a slight abuse of terminology, we will refer to these as the first four moments of the distribution. To be sure, the variance is centered by the mean, while the skewness and kurtosis are also centered and, additionally, standardized.

³ Mean-variance models have been used in queueing analysis of customer queue-joining decisions; see, for example, Wang and Zhang (2018). Moment-based decision models are also common in many settings that involve decisions with regards to money, for example, the classic mean-variance analysis of portfolio selection Markowitz (1952, 2014), or the three-moment (mean, variance, skewness) CAPM model (Kraus and Litzenberger 1976, Friend and Westerfield 1980, Harvey and Siddique 2000).

time distribution by $\mathbb{E}[\cdot]$, $\mathbb{V}[\cdot]$, $\mathbb{S}[\cdot]$, $\mathbb{K}[\cdot]$, respectively, and the sensitivity to these moments by β, γ, δ and κ , respectively. Note that since the wait in option A is deterministic, $\mathbb{E}[t^A] = t^A$, and $\mathbb{V}[t^A] = \mathbb{S}[t^A] = \mathbb{K}[t^A] = 0$. If we assume that the customer has no outside options and must choose between option A or B, then the choice of the service option comes down to the comparison of the following utilities:

$$U^A(m^A, t^A) = m^A - \beta t^A \quad (1.2)$$

$$U^B(m^B, t^B) = m^B - \beta \mathbb{E}[t^B] - \gamma \mathbb{V}[t^B] - \delta \mathbb{S}[t^B] - \kappa \mathbb{K}[t^B] \quad (1.3)$$

Our experiments examine the above choice under a variety of conditions, as we vary features of $f(t^B)$. Further, recognizing the practical reality of the choice where the customer may not have access to full distributional information, we examine several scenarios with incomplete information about $f(t^B)$.

The sample sizes, pre-registration links, research questions, and key findings of each of our studies are summarized in Table 1. In Study 1 (§2), we focus on a relatively simple environment where $f(t^B)$ is a two-point distribution (either a short or a long waiting time). We present participants with a series of binary choices (multiple price lists, Holt and Laury 2002), between Option A and Option B as we vary m^B , $\mathbb{E}[t^B]$ and $\mathbb{V}[t^B]$. After making all their decisions, participants experience one of their decisions in real-time. For that decision, they are required to either wait t^A minutes prior to receiving m^A dollars, or wait t^B minutes prior to receiving m^B dollars, depending on their reported preference. Having to experience the (time and money) consequences of their choices provides participants with an incentive to report their preferences truthfully. There are two between-subjects treatments. In the first treatment, t^B is uncertain and participants receive full information about $f(t^B)$, i.e., are informed about the support values and the probabilities of the short/long duration outcomes. In the second treatment, t^B is not merely uncertain, but *ambiguous*, i.e., the probabilities are not known. By examining behavior within each treatment, we can evaluate how participants incorporate the first two moments, $\mathbb{E}[t^B]$ and $\mathbb{V}[t^B]$ into their decisions. By comparing behavior *across* treatments we can identify whether ambiguity about the waiting times further affects their decisions.

The results of Study 1 align with the conventional wisdom that the cost of waiting increases with the expected duration and uncertainty of a wait. However, we find that small amounts of uncertainty do not affect valuations. For example, a wait of four or six minutes (with equal probabilities) is valued similarly to a wait of five minutes with certainty. However, a wait of one or nine minutes (with equal probabilities) is valued significantly less highly than a wait of five minutes with certainty. Further, on average, we observe no significant differences in the evaluations of ambiguous waits (where the probability of the short/long duration outcome is unknown) compared to uncertain waits (where probabilities are known).

Study 2 (§3) focuses on examining whether the first two moments (mean and variance) fully characterize the response to a waiting time distribution, or whether other distributional features also affect behavior. To do this, we examine (in three between-subject treatments) the response to three distributions that hold

mean and variance constant but have different distributional shapes. Two of the treatments are symmetric: in one of these treatments, $f(t^B)$ is a two-point distribution (as in Study 1), while in the second treatment, $f(t^B)$ is uniform. In the third treatment, $f(t^B)$ is exponential and therefore has a right skew. Across all three treatments, we hold constant the mean and variance of the distributions, thus allowing us to isolate the effects of the distributional shape (beyond the first two moments) on decisions.

The results of Study 2 add several important insights. First, we find that the cost of waiting may depend on the waiting time distribution, even when the mean and variance are held constant. Two-point waits are significantly preferred to continuous waits (both uniform and exponential). For the two continuous distributions (uniform and exponential), the specific distributional shape does not have a significant effect on decisions, on average. However, we find that the response to variance depends on the distributional form: while Uniform and Bernoulli waits become significantly less desirable with increased variance, the evaluation of exponential waits does not change with variance. Thus, the long right tail of the exponential distribution appears to have a distinct effect on behavior that is not fully captured by the variance measure.

Study 3 (§4) presents a more systematic evaluation of the effects of each of the four moments on decisions. In particular, Study 3 introduces more flexible (discrete) waiting time distributions, in which we can more precisely manipulate skewness and kurtosis and examine how people respond to each moment. Further, Study 3 consists of two parts: Study 3A and Study 3B. In Study 3A we focus on the full information

Table 1 Pre-registration Links, Research Questions, and Key Results

	Pre-registration Link	Research Questions	Key Results
Study 1 (§2) $N = 206$	https://aspredicted.org/8ZN_XDQ	For a two-point distribution, are uncertain waits valued less than certain waits? Are ambiguous waits valued less than uncertain waits?	Uncertain waits are valued less highly than certain waits. However, small amounts of uncertainty do not affect valuations. No statistically significant differences between ambiguous and uncertain waits.
Study 2 (§3) $N = 316$	https://aspredicted.org/C32_J22	Is the response to waiting time information affected by the distributional form of the wait?	Two-point waits (short/long) are more desirable than continuous waits, but there are no statistically significant differences between valuations of uniform and exponential waits.
Study 3A (§4.1) $N = 219$	https://aspredicted.org/DGX_TC5	What is the response to the skewness and kurtosis of the waiting time distribution?	There is a significant negative response to both skewness and kurtosis. Among the two, the response to kurtosis is substantially larger.
Study 3B (§4.2) $N = 304$	https://aspredicted.org/Y27_X27	Does the behavior from Study 3A replicate under partial information? What type of waiting time information are people most interested in?	Responses to moments under partial information are nearly identical to responses under full information. Further, waits with full or partial distributional information are preferred to waits with no such information (Min/Max only). People are most interested in receiving right-tail information about a wait.

setting. In Study 3B, we extend our analysis to settings with incomplete information, where the decision-maker receives either no information, or only partial information about the shape of the distribution. To do this, we ask participants about the type of waiting time information they are most interested in receiving (left tail/midrange/right tail), and present them with that partial information.

Study 3 replicates the findings from Studies 1 and 2 and adds several new results. First, in addition to mean and variance, we find that decision-makers respond negatively to right-skewed and “fat-tailed” (high-kurtosis) distributions. This behavior is robustly observed both with complete and with partial information about the waiting time distribution. Further, the response to kurtosis is substantially stronger than the response to skewness. Indeed, a simple mean+kurtosis model fits the data better than the more standard mean+variance model and is only slightly worse than the full four-parameter model. Second, different from the “null result” of no ambiguity aversion in Study 1, we find significant ambiguity aversion in Study 3; that is, in the more complex setting of Study 3 (relative to Study 1), providing distributional information significantly decreases the cost of waiting. Finally, we find that the majority of participants strongly prefer to receive right-tail information (as opposed to central tendency or left-tail information) of a wait and respond more strongly to that information.

In §5 we present an integrated discussion of our results and discuss their implications for service design. To supplement this discussion, in §5.3 we present a numerical case study, based on a discrete event simulation of a single-server system with multiple customer classes that differ in their service times. We compare the sojourn time distributions under two priority rules: the First Come First Serve rule (FCFS), and the Shortest Processing Time (SPT) rule. We show that, under the utility function estimated from our experimental data, the classic result that the SPT rule minimizes average waiting costs (Smith et al. 1956, Pinedo 2012) no longer holds. The case study highlights the importance of considering higher-order moments in the design of service systems.

1.3. Contributions

Our results provide new micro-foundations for modeling customer decision-making in service systems. We show that, across multiple studies and parametrizations, people not only respond negatively to increases in the mean and variance, but also to skewness and kurtosis. These findings extend prior experimental research which has focused primarily on the first two moments of the waiting time distribution (Leclerc et al. 1995, Kroll and Vogt 2008, Abdellaoui and Kemel 2014, Festjens et al. 2015, Flicker and Hannigan 2022). Failing to account for higher moment effects may inaccurately represent customer behavior and may have practical implications for service design, such as underestimating the negative impact of highly skewed, or fat-tailed waiting time distributions. Methodologically, we develop a novel incentive-compatible approach to measure preferences over time-money bundles. Much of the extant work in this area has been based on hypothetical-choice (unincentivized) experiments with complete information and simple, Bernoulli distributions (Leclerc

et al. 1995, Kroll and Vogt 2008, Festjens et al. 2015). Our measurements rely on decisions made in a more realistic environment which accommodates settings with more complex distributions and incomplete waiting time information, and where decisions have both time and money consequences.

2. Study 1: Uncertainty and Ambiguity

Study 1 focuses on the effects of uncertainty and ambiguity when evaluating simple, two-point (short/long) waiting time distributions. The experimental design, including sample size, treatment conditions, exclusion criteria, hypotheses, and analysis was pre-registered and can be accessed at https://aspredicted.org/8ZN_XDQ.

2.1. Experiment Design and Hypotheses

Study 1 examines two questions: (i) Whether uncertain waits are valued less than certain waits and (ii) whether ambiguous waits are valued less than uncertain waits. All experiments were programmed in oTree (Chen et al. 2016) and conducted online via Prolific.⁴

2.1.1. Protocol The experiment consisted of six multiple price lists (Holt and Laury 2002). We will refer to each list as a “decision set.” Screenshots of the decision sets are shown in Figure 1. Each decision set consisted of 21 binary choice scenarios. The waiting scenarios are presented in Table 2. In each scenario, Option A was: “Receive \$1.00 and wait 1 minute”, while Option B was “Receive \$ X and wait Y minutes”, where X and Y depended on decision set and the scenario within a decision set. In particular, X was varied within a decision set, while Y was varied across decision sets. In each of the six decision sets, X was varied between \$5.00 and \$1.00 in \$0.20 increments.⁵ Depending on the decision set and treatment, Y could either be deterministic, uncertain (with known probabilities), or ambiguous (with unknown probabilities). In each non-deterministic scenario, we examined a shorter and a longer spread (4/6 and 1/9 minutes for the 5-minute average and 9/11 and 4/16 minutes for the 10-minute average). Decision problems were blocked according to whether they were deterministic or not, and the sequence in which decision sets were presented to participants was randomized both by and within blocks.

The progression of the experiment was as follows. First, participants read through the instructions for the first randomized block (either deterministic or non-deterministic), and completed several comprehension checks. After that, they submitted decisions for each decision problem within the block. Upon completion of the first block, they were informed of a change in the Option B waiting scenarios and made decisions for

⁴ The online setting offers a suitable testing environment for understanding people’s preferences for different service offerings. Previous research has shown that online workers exhibit preferences across multiple domains (risk and time preferences, consumption behaviors) that are representative of the general population (Paolacci and Chandler 2014, Palan and Schitter 2018, Snowberg and Yariv 2021).

⁵ Within a decision set, participants were only allowed at most one switch point from Option B to Option A. This is done to impose strict monotonicity in revealed preferences and is common practice in experiments using multiple price list designs (Gonzalez and Wu 1999, Andersen et al. 2006).

Figure 1 Study 1: Screenshot for Uncertain Treatment

Option A			Option B		
Amount	Wait		Amount	Wait	
Receive \$1.00 bonus	Wait 1 min.	<input type="radio"/> <input type="radio"/>	Receive \$5.00 bonus	Wait 4 or 6 min. with equal chance	
Receive \$1.00 bonus	Wait 1 min.	<input type="radio"/> <input type="radio"/>	Receive \$4.80 bonus	Wait 4 or 6 min. with equal chance	
Receive \$1.00 bonus	Wait 1 min.	<input type="radio"/> <input type="radio"/>	Receive \$4.60 bonus	Wait 4 or 6 min. with equal chance	
•	•		•	•	
•	•		•	•	
•	•		•	•	
Receive \$1.00 bonus	Wait 1 min.	<input type="radio"/> <input type="radio"/>	Receive \$1.40 bonus	Wait 4 or 6 min. with equal chance	
Receive \$1.00 bonus	Wait 1 min.	<input type="radio"/> <input type="radio"/>	Receive \$1.20 bonus	Wait 4 or 6 min. with equal chance	
Receive \$1.00 bonus	Wait 1 min.	<input type="radio"/> <input type="radio"/>	Receive \$1.00 bonus	Wait 4 or 6 min. with equal chance	

Note: The screenshot shows the case of the Uncertain Treatment. In the Ambiguous Treatment, everything was identical except that the word “equal” was replaced with “unknown” for the Option B wait scenarios.

Table 2 Study 1: Treatments and Waiting Scenarios

Scenario (varied within-subject)	Wait Times in Option B	Treatment (varied between-subject)	
		Uncertain	Ambiguous
1	5 min.	with prob. 1	with prob. 1
2	10 min.	with prob. 1	with prob. 1
3	4 min. or 6 min.	with prob. 0.5	with unknown prob.
4	1 min. or 9 min.	with prob. 0.5	with unknown prob.
5	9 min. or 11 min.	with prob. 0.5	with unknown prob.
6	4 min. or 16 min.	with prob. 0.5	with unknown prob.

Note: The sequence of scenarios was determined at random for each participant. First, the sequence of blocks (deterministic, i.e., scenarios 1-2, or non-deterministic, i.e., scenarios 3-6) was determined at random. Second, the sequence of scenarios within each block was determined at random.

the remaining problems in this block. After completing all decision problems, one decision set was selected at random, and one decision within that set was selected at random to count for payment, with the wait from that decision being implemented in real time. Depending on their choice, participants either had to wait for one minute (Option A) or Y minutes, where Y depended on the waiting scenario (Option B) for the selected problem.⁶

2.1.2. Treatments We conducted two between-subjects treatments: an Uncertain treatment, in which participants were given the probabilities for the short and long waiting time in Option B, as well as an Ambiguous treatment, in which they were not. We used 50% probabilities in the Uncertain treatment. To implement ambiguity in the experiment, we used the Stecher et al. (2011) method. Participants were told:

⁶ A total of three participants dropped out during the wait. They did not receive payment and are not included in our data.

As you can see below, waiting times may be uncertain. Suppose you choose option B. To determine the duration we will use a computerized coin whose likelihood of landing on heads or tails is somewhere between 0 and 100%. You can think about it just as a regular coin, but instead of having 50% probability of falling on heads or tails, that probability is X . In fact, even we (the experimenters) do not know X – it will be determined by a random algorithm after you make all your choices.

The advantage of this method is that it explains the uncertainty-generating process with clarity and transparency and minimizes the perception of informational asymmetries between the experimenter and the subjects.⁷

2.1.3. Sample, Exclusions and Incentives All experiments were conducted on weekdays between 10 am and 6 pm Eastern Time. A total of 246 Prolific workers were recruited (average age: 36.5, 54% female). Only US-based workers with an approval rating of at least 99% were eligible. A total of 40 participants were excluded based on pre-registered comprehension and attention checks, resulting in a sample size of 206.⁸ Recruitment was stopped once the target sample of 100 participants per treatment was reached. As noted earlier, participants were incentivized to report their preferences truthfully by experiencing one of their choices at the end of the experiment. During the wait every 30 seconds the time stopped and they were prompted to enter a (randomly chosen) key on the keyboard (See Appendix A for a screenshot of the wait). This was done to ensure that participants were actively waiting (rather than leaving their computers during the wait). After the wait participants were redirected to the exit survey and payment. In addition to their earnings from the experiment, all participants received a \$3.00 show-up fee.

2.1.4. Hypotheses While risk preferences for money have been studied extensively in the experimental literature (Holt and Laury 2002, Eckel and Grossman 2008, Harrison and Cox 2008, Kagel and Roth 2020), research on decision-making under uncertainty in the time domain remains quite limited. Some existing research invokes Prospect theory (Kahneman and Tversky 1979, Tversky and Kahneman 1992) and suggests that since time expenditures are losses, people may be risk-seeking for time (Kroll and Vogt 2008, Abdellaoui and Kemel 2014). More frequently, however, research finds risk-averse behaviors in the time domain (Leclerc et al. 1995, Festjens et al. 2015, Flicker and Hannigan 2022, Kagan et al. 2022).

H1.1: The attractiveness of an uncertain wait decreases in the variance of the wait.

Our next hypothesis concerns people's attitudes towards not just uncertain but ambiguous prospects. Situations where the information about a prospective wait is unknown or incomplete are abundant in practice and are the norm rather than the exception. For example, many food delivery platforms present customers

⁷ Such asymmetries have been shown to confound choices, see Abdellaoui et al. (2011), Stecher et al. (2011).

⁸ The pre-registered exclusion criteria were as follows. First, we removed participants from the sample who made three or more errors on the quiz. Second, we removed participants who selected a dominated option (i.e., choose Option B when the monetary payment was the same, but the waiting time was longer than for Option A, see the bottom decision in Fig. 1).

with a range of possible delivery times, rather than a detailed description of possible outcomes with attendant probabilities. Similarly, ride-sharing services often provide a window of time for when a driver will arrive, instead of an exact pickup time. The customer is typically not informed about the likelihood of the outcomes within that window.

There is no research that we are aware of that examines ambiguous waits. However, research in decision theory suggests that ambiguous prospects involving money are valued less than uncertain prospects (Ellsberg 1961, Einhorn and Hogarth 1986, Camerer and Weber 1992). A common explanation for this behavior is the “comparative ignorance hypothesis” (Fox and Tversky 1995), namely that the feeling of being less confident (or more ignorant) about the data-generating process of the distribution causes the aversive response. In our experiments, all choices involve a clear, deterministic time-money bundle in Option A and a more vague bundle in Option B. Thus, if the findings on ambiguity aversion for money extend to the time domain, we should expect that, all else being equal, people prefer uncertain prospects to ambiguous prospects. Therefore, we hypothesize the following:

H1.2: An ambiguous wait is less attractive than an uncertain wait.


2.2. Results

To test our hypotheses we will use “switching amounts” as our dependent variable. A switching amount is the dollar value of Option B for which the participant switches from Option B (larger payment and a longer, potentially uncertain, wait) to Option A (small dollar payment and a short wait). Recall from Figure 1 that the dollar amounts become progressively smaller as one goes down the list of scenarios: from \$5.00 in decision 1 in the top row to \$1.00 in decision 21 in the bottom row. All else being equal, a larger switching amount therefore indicates a larger price that the subject is willing to pay to avoid the waiting experience in Option B.⁹

2.2.1. Summary Statistics Average switching amounts for each treatment and condition are shown in Figure 2. First, as waiting durations increase, so does the amount of required compensation. Average switching amounts in the first five scenarios (in the left half of Figure 2) are below the last five (right half of the figure). While we did not formulate hypotheses for this behavior, it is not surprising and simply confirms that longer waiting times are perceived as more painful. Second, for each average waiting time, there is very little difference between the deterministic wait and the low variability wait (5 min. vs. 4 or 6 min. and 10 min vs. 9 or 11 min.). However, the switching amount is higher in the scenarios with a greater spread between the short and long waiting times (4 or 6 min. vs. 1 or 9 min. and 9 or 11 min. vs. 4 or 16 min.). Last, comparing the uncertain and ambiguous scenarios, the differences are quite small for all four pairwise comparisons.

⁹ As noted in §2.1.1, we do not allow multiple switching between Option A and Option B.

Figure 2 Study 1: Average Switching Points by Treatment and Type of Wait



2.2.2. Hypothesis Tests To test our hypotheses, we estimated a linear random-effects model where the dependent variable was the switching amount. As explanatory variables we include the ‘average wait,’ the ‘standard deviation’ of the wait, as well as indicator variables for whether the waiting scenario was ambiguous.¹⁰

Table 3 reports the regression coefficients. First, unsurprisingly, there is a significant effect of increasing average waiting time in Option B (denoted by *averageWait* in Table 3): participants are willing to trade each additional minute for approximately 11.5 cents ($p \ll 0.001$). Second, the effects of increasing uncertainty (*sdWait*) do not appear to be linear. Consistent with Figure 2, behavior is quite similar between the deterministic (*sdWait*=0) and the low-variance (*sdWait*=1) settings. However, a larger variance (conditions with *sdWait* = 4 and *sdWait* = 6) prompts participants to demand a greater monetary compensation (both coefficients significant at $p \ll 0.001$). Thus, H1.1 is supported, but only for nontrivial amounts of uncertainty. Turning to H1.2, we see that the coefficient on *Ambiguous Wait* is both quantitatively small and not significantly different from zero, which indicates that, on average, subjects do not view uncertain waits differently from ambiguous waits (in both specifications, $p > 0.2$). That is, we find no support for H1.2. Finally, the indicator for *Male* has a marginally significant effect, with males having slightly lower switching thresholds. While we did not formulate gender-specific hypotheses, we will briefly discuss gender effects in more detail in the next section.

¹⁰ Although in the Ambiguous treatment, the ‘average’ and ‘standard deviation’ of the waits are not technically defined we use them as a shorthand to describe the different waiting scenarios. Further, they would be the average and standard deviation, respectively, if subjects had a uniform belief over the support.

Table 3 Study 1: Regression Results

Dependent variable: Switching Amount (USD)		
Omitted Category: Uncertain Wait, averageWait=5, sdWait=1.	(1)	(2)
averageWait	0.115*** (0.007)	0.115*** (0.007)
sdWait		
0 (Deterministic)	0.008 (0.041)	0.009 (0.041)
4	0.174*** (0.035)	0.174*** (0.035)
6	0.204*** (0.041)	0.204*** (0.041)
Ambiguous Wait	0.064 (0.056)	0.068 (0.056)
Age		0.004 (0.006)
Male		−0.252* (0.134)
Education		−0.002 (0.077)
Income		−0.048 (0.050)
Constant	1.276*** (0.066)	1.311*** (0.251)
R^2	0.080	0.101
No. Observations	1236	1236
No. Participants	206	206

Notes: Random effects regression coefficients are reported (Standard errors in parentheses). *, ** and *** denotes significance at the 10%, 5% and 1% level, respectively.

2.2.3. Additional Analysis Although we did not find any evidence of ambiguity aversion on average, further (non-preregistered) analysis suggests that there are important differences by gender. First, we find that female participants demand more monetary compensation on average - for example, for deterministic waits, the gender gap ranges between \$0.19 and \$0.32. Indeed, the average effect of gender is marginally significant ($p = 0.06$, see Table 3). Second, women demand higher compensation under ambiguity than under risk, and vice versa for men. Additional regression analysis confirms these results with men being ambiguity-seeking ($p = 0.016$), women being slightly ambiguity-averse ($p = 0.061$), and ambiguity attitudes between men and women being significantly different from each other ($p = 0.003$). Detailed analysis and graphs are in Appendix B.1.

3. Study 2: Role of Distributional Form

We have so far examined how uncertainty and ambiguity affect the choice among time-money bundles in a relatively simple decision environment. In many real-life service settings, however, an uncertain wait may involve a *range* of potential waiting times, rather than the binary short/long waiting times examined in Study 1. Further, within that range, a wait may be right-skewed - likely short but with a long tail of potential outliers - or more uniformly distributed along the range. Study 2 examines how people make decisions when facing these alternative distributional forms. The experimental design, including sample size, treatment conditions, exclusion criteria, hypotheses, and analysis was pre-registered (https://aspredicted.org/C32_J22).

3.1. Experiment Design and Hypotheses

The general structure of Study 2 was analogous to Study 1, but with some key differences.

3.1.1. Treatments We administered three between-subjects treatments:

- Binary treatment: Participants faced either a short or long waiting time in Option B, with equal probabilities (analogous to the Uncertain treatment of Study 1).
- Uniform treatment: The waiting time in Option B was uniformly distributed over an interval.
- Exponential treatment: The waiting time in Option B followed an exponential distribution.

Table 4 summarizes the key details of these treatments. Unlike Study 1, we did not randomize the order of the two blocks (deterministic \rightarrow non-deterministic, or vice versa). Instead, all participants first completed the deterministic scenarios (Scenarios 1-2 in Table 4), followed by the non-deterministic ones (Scenarios 3-6 in Table 4), with the individual scenarios randomized within each block. This change was made because (a) we did not observe any order effects in Study 1, and (b) we are no longer interested in participant behavior in deterministic scenarios; rather, we used deterministic scenarios to help participants get familiarized with the task, before proceeding to the more complex, uncertain scenarios.

Table 4 Study 2: Treatments and Waiting Scenarios

Scenario (varied within-subject)	Treatment (varied between-subject): Distribution used for Option B			Mean and Standard deviation
	Binary	Uniform	Exponential	
1	deterministic	deterministic	deterministic	$\mu = 5, \sigma = 0$
2	deterministic	deterministic	deterministic	$\mu = 10, \sigma = 0$
3	4 or 6, with prob. 0.5	$U[3.27, 6.73]$	$\exp(\lambda = 1, \text{mean shift} = 4)$	$\mu = 5, \sigma = 1$
4	3 or 7, with prob. 0.5	$U[1.54, 8.56]$	$\exp(\lambda = 1/2, \text{mean shift} = 3)$	$\mu = 5, \sigma = 2$
5	9 or 11, with prob. 0.5	$U[8.27, 11.73]$	$\exp(\lambda = 1, \text{mean shift} = 9)$	$\mu = 10, \sigma = 1$
6	5 or 15, with prob. 0.5	$U[1.34, 18.66]$	$\exp(\lambda = 1/5, \text{mean shift} = 5)$	$\mu = 10, \sigma = 5$

Note: The sequence of scenarios within each block was determined at random for each participant. All parameters are in minutes.

We carefully parameterized the distributions to hold the mean and variance constant within each scenario. As in Study 1, we examined different mean-variance combinations. For example, in Scenario 3, all three treatments presented a choice between Option A (\$1 and 1-minute wait) and Option B (larger payment, 5-minute mean wait, 1-minute standard deviation, with distribution depending on treatment). In Scenario 4, we kept the mean the same as in Scenario 3 but increased the standard deviation to 2. In Scenarios 5-6 we examined durations with a higher mean. In all treatments, participants were informed about the probability distribution they were facing in each scenario. To help participants develop an intuitive understanding of each probability distribution, we used simplified density plots with quartiles marked on top of the density (see Figure A1). Each distribution was accompanied by explanations and examples, and participants' understanding was tested in a comprehension quiz.

3.1.2. Sample, Exclusions and Incentives As before, all experiments were conducted on weekdays between 10 am and 6 pm Eastern Time. A total of 372 Prolific workers were recruited (average age: 38.5, 58% female). Only US-based workers with an approval rating of at least 99% were eligible. A total of 56 participants were excluded based on pre-registered comprehension and attention checks, resulting in a sample size of 316. Recruitment was stopped once the target sample of 100 participants per treatment was reached. As in Study 1, participants were incentivized to report their preferences truthfully by experiencing one of their choices (both in terms of time and money) in real-time at the end of the experiment. As before, all participants received a \$3.00 show-up fee.

3.1.3. Hypotheses The literature on preferences for different waiting time distributions is quite limited. Prior studies have primarily focused on binary (short/long) waiting time distributions (Leclerc et al. 1995, Abdellaoui and Kemel 2014, Kroll and Vogt 2008, Festjens et al. 2015, Kagan et al. 2022). Similarly, research on queue-joining behavior does not present participants with an explicit time distribution but rather allows them to infer duration from queue length or configuration (Kremer and Debo 2016, Buell 2021). Also related is the literature on the response to delay information (Soman 2001, Althenayyan et al. 2022, Aksin et al. 2022, Yu et al. 2022, Ansari et al. 2022, Debo et al. 2023). Different from us, these studies focus on events and behaviors that occur *while* waiting, rather than on ex-ante choices among different service offerings.

While two of our treatments present participants with a symmetric waiting time distribution, the third one (Exponential) has a right skew. Several studies examining risk preferences for *money* find that people may prefer right-skewed distributions, even when means and variances are held constant (Brünnner et al. 2011, Wu et al. 2011) – a behavior that has been attributed to factors like over-optimism or overweighting of small probabilities in the right tail (Åstebro et al. 2015). However, unlike money, for prospects involving time, a longer right tail (e.g., in an exponential distribution) would correspond to a longer wait, potentially leading to aversive preferences. It is therefore not clear how people would respond to a right skew in our setting. Given the relative scarcity of research in this area, we do not form a strong directional hypothesis and instead hypothesize that people’s time preferences will be primarily driven by the mean and variance of the time distributions, rather than their specific distributional forms.


H2: Holding constant the mean and variance of waiting times, the attractiveness of a wait is the same for binary, uniform, and exponential waiting time distributions.

3.2. Results

3.2.1. Summary Statistics The average switching amounts are shown in Figure 3. Consistent with Study 1, we see that switching amounts are higher for longer average waits. Further, holding the average wait constant, switching amounts increase with the standard deviation of the wait. However, there also appear to

be some differences between the distributions. As in Study 1, we observe only minimal differences between the deterministic and the $\sigma = 1$ scenarios when the distribution of the wait is binary (\$1.79 vs. \$1.83 and \$2.26 vs. \$2.30). However, this is no longer true for the remaining two distributions. Even a small amount of variability appears to increase switching points relative to the deterministic scenarios (\$1.79 vs. \$1.91, \$1.79 vs \$2.17, \$2.26 vs. \$2.50 and \$2.26 vs. \$2.67). Indeed, the switching amount for both uniform and exponential waits is always higher than for binary waits. Finally, participants appear to be quite insensitive to standard deviation when dealing with exponential waits (\$2.17 vs. \$2.15 and \$2.67 vs. \$2.63).

Figure 3 Study 2: Average Switching Points by Treatment and Type of Wait



3.2.2. Hypothesis Tests Table 5 shows the tests of whether switching amounts differ based on the features of the waiting time distribution. It reports the results of random effects regressions where the dependent variable is the switching amount, and explanatory variables are the mean and standard deviation of the wait in Option B (`averageWait` and `sdWait`), as well as the indicator variables for the probability distribution (`Uniform Wait` and `Exponential Wait`). To test our main pre-registered hypothesis, H2, we conduct the joint test that the coefficients on `Binary Wait`, `Exponential Wait`, and `Uniform Wait` are all equal. The results of these tests are reported in the last row of the table. As can be seen, H2 is rejected – behavior is not the same for all waiting distributions, with the gap being significantly different from zero at $p < 0.05$, both with and without controls.

3.2.3. Additional Analysis Recall that a higher switching amount (dependent variable in Table 5) indicates that participants demand greater monetary compensation for a given wait. Thus, the positive coefficients on both treatment dummies (`Uniform wait` and `Exponential wait`) in Table 5 suggest that

the binary distribution is preferred to both the uniform and exponential distributions. However, only the difference between the binary and uniform distribution is statistically significant in pairwise tests ($p = 0.015$), while the difference between the binary and exponential distribution is not ($p = 0.186$). Further inspection of Figure 3 suggests that the latter (null) result is driven mainly by the insensitivity to variance in choices that involve the exponential distribution. Particularly for the $\mu = 10$ conditions, binary waits are evaluated quite similarly to exponential waits. These results suggest that, consistent with the literature on risk aversion for money (Wu et al. 2011, Br  nner et al. 2011,   stebro et al. 2015), a right skew in a distribution is a salient factor that affects behavior. However, different from this literature, we find that the preference (for or against such distributions) may depend on the average length of the wait. For shorter waits, skewness appears to lower the attractiveness of a wait; however, for longer wait it appears to have no effect.

Further regression analysis confirms that the response to variance is indeed significantly different across treatments. In particular, if we replicate the analysis in column (1) of Table 5 with the addition of interaction terms between each `sdWait` dummy and each distribution dummy (`Uniform Wait` and `Exponential Wait`), we find that the variance effect is statistically significant at $p < 0.05$ for three out of four cases for the binary and uniform case, but is not statistically significant in the exponential case ($p = 0.616$ and $p = 0.671$). Further, after controlling for these interactions, each main effect of the distribution (`Uniform Wait` and `Exponential Wait`) is statistically significant ($p = 0.025$ and $p = 0.005$), suggesting that the binary wait is significantly preferred to both continuous distributions in our data. Further, the difference between the `Uniform Wait` and `Exponential Wait` is not statistically significant ($p = 0.488$), suggesting that there is no strong preference among the two continuous distributions after controlling for their means and variances.

4. Study 3: Higher Moments and Incomplete Information

We have thus far focused on the first two moments (mean and variance) and examined settings with complete as well with ambiguous information (Study 1), as well as different distributional forms (Study 2). In Study 3, we extend our investigation as follows:

- (i) Study 2 suggests that the shape of the waiting time distribution (beyond just the first two moments) may affect decisions. However, the three distributions in Study 2 have different ranges: the exponential distribution has a longer right tail (relative to the other two), while the uniform distribution has a lower minimum and a higher maximum value relative to the binary distribution. Therefore, we cannot tell with the available data whether the choices are driven by the differences in the range or by the differences in the shape of a distribution. To identify the responses to the distributional shape, in Study 3A we will present decision-makers with distributions that have the same range of values, but that vary in their moments.

Table 5 Study 2: Regression Results

Dependent variable: Switching Amount (USD)		
Omitted Category: Binary Wait, averageWait=5, sdWait=1.	(1)	(2)
averageWait	0.098*** (0.005)	0.098*** (0.005)
sdWait		
0 (deterministic)	−0.118*** (0.038)	−0.119*** (0.038)
2	0.046** (0.020)	0.046** (0.020)
5	0.193*** (0.032)	0.193*** (0.032)
Uniform Wait	0.153** (0.062)	0.151** (0.062)
Exponential Wait	0.091 (0.068)	0.090 (0.068)
Age		0.006 (0.004)
Male		−0.080 (0.104)
Education		0.067 (0.061)
Income		0.006 (0.033)
Constant	1.409*** (0.057)	1.076*** (0.206)
R^2	0.087	0.098
No. Observations	1896	1896
No. Participants	316	316
Binary Wait = Uniform Wait = Exponential Wait	$p = 0.043$	$p = 0.045$

Notes: Random effects regression coefficients are reported (Standard errors in parentheses). *, ** and *** denotes significance at the 10%, 5% and 1% level, respectively. The bottom row reports the p -value of the joint test that all three uncertain wait coefficients are equal (pre-registered hypothesis).

(ii) Study 3 will also further our analysis of settings with incomplete information. In Study 1 we have seen that, in the case of the symmetric, binary distribution, ambiguous waits were evaluated similarly to uncertain waits. In Study 3B we extend the comparison of different information regimes to include more realistic scenarios with more complex distributional shapes. In particular, we will examine scenarios where the range (min, max), but not the distribution, is known. In addition, we will examine scenarios where only part of the distribution (for example, left tail) is revealed. Such scenarios are common in many service settings where a customer may have an uninformed prior belief about the duration of a wait or may receive partial information about the range, average, typical, or maximum wait.

4.1. Study 3A: Complete Information

In Study 3A we focus on complete information scenarios (incomplete information scenarios will be examined in Study 3B). The study protocol was similar to Studies 1 and 2, with several key differences that we describe here. As before, the experimental design, including sample size, treatment conditions, exclusion criteria, hypotheses, and analysis was pre-registered (https://aspredicted.org/Y27_X27).

4.1.1. Protocol and Treatments As before, all participants completed six decision sets (with each decision set corresponding to a multiple price list). Similarly, Option A was always a short wait ($t^A = 1$

Table 6 Studies 3A and 3B: Treatments and Waiting Scenarios (Option B)

		Low Range Treatment					High Range Treatment				
Shape		Possible wait times	$\mathbb{E}[t^B]$	$\mathbb{V}[t^B]$	$\mathbb{S}[t^B]$	$\mathbb{K}[t^B]$	Possible wait times	$\mathbb{E}[t^B]$	$\mathbb{V}[t^B]$	$\mathbb{S}[t^B]$	$\mathbb{K}[t^B]$
L1	Heavy left tail	$\{1, 2, \dots, 6\}$ min	2.00	1.40	1.54	5.15	$\{1, 2, \dots, 15\}$ min	4.25	9.19	1.43	4.95
L2	Heavy left tail	$\{1, 2, \dots, 6\}$ min	2.20	2.36	1.44	3.83	$\{1, 2, \dots, 15\}$ min	4.75	15.19	1.38	3.78
L3	Heavy left tail	$\{1, 2, \dots, 6\}$ min	2.50	2.05	0.82	2.80	$\{1, 2, \dots, 15\}$ min	5.50	13.25	0.78	2.81
L4	Heavy left tail	$\{1, 2, \dots, 6\}$ min	2.90	3.49	0.56	1.72	$\{1, 2, \dots, 15\}$ min	6.50	22.25	0.54	1.77
M1	Symmetric	$\{1, 2, \dots, 6\}$ min	3.50	1.05	0.00	4.05	$\{1, 2, \dots, 15\}$ min	8.00	7.00	0.00	3.91
M2	Symmetric	$\{1, 2, \dots, 6\}$ min	3.50	1.85	0.00	2.59	$\{1, 2, \dots, 15\}$ min	8.00	12.00	0.00	2.62
M3	Symmetric	$\{1, 2, \dots, 6\}$ min	3.50	2.92	0.00	1.73	$\{1, 2, \dots, 15\}$ min	8.00	18.67	0.00	1.79
M4	Symmetric	$\{1, 2, \dots, 6\}$ min	3.50	3.45	0.00	1.48	$\{1, 2, \dots, 15\}$ min	8.00	22.00	0.00	1.54
R1	Heavy right tail	$\{1, 2, \dots, 6\}$ min	4.10	3.49	-0.56	1.72	$\{1, 2, \dots, 15\}$ min	9.50	22.25	-0.54	1.77
R2	Heavy right tail	$\{1, 2, \dots, 6\}$ min	4.50	2.05	-0.82	2.80	$\{1, 2, \dots, 15\}$ min	10.50	13.25	-0.78	2.81
R3	Heavy right tail	$\{1, 2, \dots, 6\}$ min	4.80	2.36	-1.44	3.83	$\{1, 2, \dots, 15\}$ min	11.25	15.19	-1.38	3.78
R4	Heavy right tail	$\{1, 2, \dots, 6\}$ min	5.00	1.40	-1.54	5.15	$\{1, 2, \dots, 15\}$ min	11.75	9.19	-1.43	4.95

minute) accompanied by a small payment ($m^A = \$1$), while Option B was a longer, potentially uncertain wait, accompanied by a larger payment ($m^B \in [\$1, \$5]$). At the beginning of the experiment, participants completed two preliminary rounds in which they were presented with a simplified version of the task: in the first round, the wait in Option B was deterministic, and in the second round they were only shown the range, but not the shape of the distribution. In the remaining four rounds, participants saw a random subset drawn from 12 possible distributional shapes (i.e., a total of 24 distributions across the two treatments) summarized in Table 6. In each of the distributions, t^B was distributed according to a discrete probability distribution with support values $\{1, 2, \dots, 6\}$ (resp.: $\{1, 2, \dots, 15\}$) minutes in the Low Range (resp.: High Range) treatment. In each round of the experiment, one of the twelve distributions shown in Table 6 was drawn at random and was used for Option B. Participants were informed about the distribution at the beginning of each round and were reminded of it on the choice screen.

The advantage of the random sampling design used in this study is that it allows us to separately identify the response to each of the four moments of the waiting time distribution. The distributions span a wide range from right-skewed distribution with a heavy left tail (Distributions L1-L4 in Table 6), to symmetric (Distributions M1-M4 in Table 6), to left-skewed with a heavy right tail (Distributions R1-R4 in Table 6). A graphical representation of each of the 24 distributions is displayed in Fig. B2 in Appendix B. Importantly, these distributions vary in all four moments (except for the M1-M4 distributions which are symmetric and thus have the same mean and skewness). Evaluating the attractiveness of these distributions will thus allow us to robustly identify the response to each of the four moments.

4.1.2. Sample, Exclusions and Incentives As in our first two studies, all experiments were conducted on weekdays between 10 am and 6 pm Eastern Time. A total of 250 Prolific workers were recruited (average age: 39.1, 51.9% female). Only US-based workers with an approval rating of at least 99% were eligible. A total of 31 participants were excluded based on pre-registered comprehension and attention checks,

resulting in a sample size of 219. Recruitment was stopped once the target sample of 100 participants per treatment was reached.

4.1.3. Hypotheses We are not aware of existing studies examining the response to higher moments when evaluating time durations; however, some testable predictions can be derived from our prior results in Study 1 and Study 2. Specifically, Study 1 suggests that people respond negatively to the first two moments of a distribution. Further, Study 2 shows that binary waits (which have the lowest skewness and kurtosis among the distributions examined in Study 2) are preferred to both exponential and uniform waits, suggesting that people may respond negatively to the third and to the fourth moment of a distribution. Therefore, we hypothesize the following:

H3.1: The attractiveness of a wait decreases in its mean, variance, skewness, and kurtosis.

4.1.4. Results To conserve space, in the main text, we focus on presenting the results of hypothesis tests for H3.1 (summary statistics and more detailed discussion of our testing approach are in Appendix B). To test H3.1, we estimate the utility function described in equations (1.2) and (1.3) using the conditional logit model (McFadden 1974, Ben-Akiva 1974). We use this approach because, unlike Study 1 and Study 2, Study 3 is based on a random sampling design, so that each participant experiences a different set of comparisons when making decisions. The conditional logit approach directly models this choice process by specifying the probability of choosing each alternative as a function of its attributes. The conditional logit model assumes that decision-makers make probabilistic choices based on a comparison of the utilities offered by Alternative 1 (eq. 1.2) and Alternative 2 (eq. 1.3), with the addition of a random error term. If we assume that the random error follows a Type I extreme value distribution (Gumbel distribution), the probability of choosing each alternative can be shown to follow a logistic distribution, allowing us to estimate (1.2)-(1.3) using maximum likelihood estimation. See Appendix B.4 for technical details.

Table 7 presents four nested models. The models progressively incorporate higher moments of $f(t^B)$ (distribution of the wait in Option B), starting with mean only, then adding variance, skewness, and finally kurtosis. To evaluate the statistical significance of each higher-order moment for model fit, we conducted a series of likelihood ratio tests. The p -values for these tests are as follows: comparing the mean/variance model to the mean-only model ($p \ll 0.001$), the mean/variance/skewness model to the mean/variance model ($p = 0.002$), and the full model including kurtosis to the mean/variance/skewness model ($p \ll 0.001$). These results strongly support H3.1.

While each of the three moments (variance, skewness, kurtosis) helps improve model fit, the most substantial improvement comes from including kurtosis. This is evidenced by the marked decrease in log-likelihood from 9259.7 in the mean/variance/skewness model to 9155.3 in the full model (a drop of 104.4 points). In contrast, the improvement in fit from adding skewness to the mean-variance model is only 6.2 points. In addition to the models presented in Table 7 we estimated more parsimonious models that included

Table 7 Study 3A: Estimation Results

	Mean	Mean/Var	Mean/Var/Skew	Mean/Var/Skew/Kur
λ (rationality parameter)	1.421	1.431	1.442	1.518
β (response to mean)	0.184	0.167	0.172	0.150
γ (response to variance)		0.016	0.014	0.010
δ (response to skewness)			0.047	0.033
κ (response to kurtosis)				0.074
Log Likelihood	9294.7	9265.9	9259.7	9155.3
No. of Participants	219	219	219	219
No. of Observations	1095	1095	1095	1095

only two parameters. This analysis shows that even in the absence of variance, kurtosis appears to be a stronger predictor of decisions than skewness (mean + variance model: log-likelihood = 9265.9, mean + skewness model: log-likelihood = 9281.2, mean + kurtosis model: log-likelihood = 9174.62). Taken together, these results suggest that kurtosis plays a significant role in explaining behavior, potentially more so than variance and skewness.

4.2. Study 3B: Incomplete Information

Study 3B is analogous to Study 3A and uses the same set of probability distributions. However, in a subset of decisions, participants were presented with partial, rather than full information about the distributions, allowing us to examine several questions related to information processing when choosing among service options. Detailed instructions are in Appendix A.3. As before, the experimental design, including sample size, treatment conditions, exclusion criteria, hypotheses, and analysis was pre-registered (https://aspredicted.org/DGX_TC5).

4.2.1. Protocol and Treatments The design of Study 3B closely followed that of Study 3A. As in Study 3A, participants first completed two (preliminary) rounds in which they experienced a deterministic scenario and a scenario in which they knew the range, but not the distribution of t^B . Participants were then informed that in the remainder of the experiment, they would receive some information about t^B and were asked to indicate what type of information they were most interested in receiving. A screenshot of the elicitation procedure in the Low Range treatment is shown in Figure 4. As shown in the figure, participants were given the following choices: reveal the probability mass in $\{1, 2\}$ range, in $\{3, 4\}$ range, or in $\{5, 6\}$ range. Analogously, the choices in the High Range treatment were: reveal the probability mass in $\{1, 2, 3, 4, 5\}$ range, in $\{6, 7, 8, 9, 10\}$ range, or in $\{11, 12, 13, 14, 15\}$ range. Each participant indicated their top and second choice only once. Then, in each round of the experiment, their top choice was revealed with 75%, and the second choice was revealed with 50%. Note that if both choices are revealed (which

Figure 4 Study 3B: Screenshot of Elicitation of Information Preferences (Low Range Treatment)

In the previous round, the waiting time in Option B was described by the following table:

Option B	
Possible Waiting Times	1, 2, 3, 4, 5, or 6 Minutes
Likelihood of each time	Unknown

In the remainder of the experiment you will receive partial information about Option B. The table below shows three regions: 1-2 minutes (short waiting time), 3-4 minutes (moderate waiting time) and 5-6 minutes (long waiting time).

Option B			
Possible Waiting Times	1-2 Min. (short wait)	3-4 Min. (moderate wait)	5-6 Min. (long wait)
Likelihood	?%	?%	?%

Below you are asked to indicate the regions whose likelihood (%) you are most interested in learning. After you do this, we will reveal some of your choices. In particular, in each of the remaining rounds we will reveal your top choice with 75% chance and your second choice with 50% chance. In other words, it is not guaranteed that you will receive all the pieces of information that you request, but **it is more likely that you receive the information that you rank higher.**

What information would you like us to reveal?

Top choice

Second choice

✓ -----

- Reveal likelihood (%) of waiting 1-2 minutes (short wait)
- Reveal likelihood (%) of waiting 3-4 minutes (moderate wait)
- Reveal likelihood (%) of waiting 5-6 minutes (long wait)

happens with probability $0.75 \times 0.5 = 0.375$), the probability mass in the third range is complementary to the first two, and is also revealed.¹¹

In sum, Study 3B differed from Study 3A in two ways. First, in a subset of the rounds, participants only saw partial waiting time information for Option B: left tail, midrange, or right tail, where the probability of a range being revealed was correlated with their stated information preference. Second, even when participants received all three pieces of information (left tail, midrange and right tail), the granularity of information was different. In Study 3A participants received probabilities for each possible support value. In Study 3B, they only received probability for each range (e.g., probability that t^B falls into the $\{1, 2\}$ range, the $\{3, 4\}$ range, or the $\{5, 6\}$ range), but not the individual probabilities. Receiving this type of approximate or coarse probabilistic information is reflective of many real-life choices in service settings.

4.2.2. Sample, Exclusions and Incentives As before, all experiments were conducted on weekdays between 10 am and 6 pm Eastern Time. A total of 346 Prolific workers were recruited (average age:

¹¹ The random draw (determining which pieces of information are revealed) was performed independently at the beginning of each period. The probabilistic information revelation mechanism will allow us to control for selection effects when analyzing the responses.

38.6, 56% female). Only US-based workers with an approval rating of at least 99% were eligible. A total of 41 participants were excluded based on pre-registered comprehension and attention checks, resulting in a sample size of 304. Recruitment was stopped once the target sample of 150 participants per treatment was reached.¹²

4.2.3. Hypotheses To examine whether the results of Study 3A hold in a more incomplete information setting, we will first re-examine H3.1, i.e., the effects of higher moments on decisions. This will allow us to test whether the previously identified effects are robust to more realistic decision environments common in practical settings where customers make decisions. Additionally, Study 3B will allow us to examine informational preferences. We are not aware of any literature that directly examines such preferences. Methodologically, the closest to us are Fréchette et al. (2017), who study information acquisition for risky monetary gambles. They find that people are most interested in the left tail of a distribution (which represents the worst-case outcome scenario), though they also find substantial heterogeneity in preferences. Given that money-based preferences do not necessarily extrapolate to decisions that concern time (Leclerc et al. 1995, Soman 2001), we take a more cautious approach and anticipate that heterogeneity in information acquisition preferences will prevail, and people will, on average, be indifferent across the three pieces of information.

H3.2: People are indifferent between receiving probabilistic information about the left tail, midrange, or right tail.

In addition to examining informational preferences, we will test several ancillary hypotheses that help support our prior results in the incomplete information setting. These hypotheses examine the response to tail and midrange information and follow directly from our prior results on how people process information about the four moments of a distribution.

H3.3:

- a) Response to left tail (right tail) probability mass increase is positive (negative).*
- b) Response to midrange probability mass increase is positive.*
- c) Response to tails is stronger than response to midrange.*
- d) Response to right tail is stronger than response to left tail.*
- e) Response to right tail is stronger for long waits.*
- f) Response is stronger to one's top choice than to one's second or third choice.*

¹² We use a slightly larger sample size (relative to Studies 1, 2, and 3A) because of the higher number of different scenarios that participants could experience (12 distributions \times five possible information scenarios).

Table 8 Study 3B: Estimation Results (Both Tails and Midrange Revealed)

	Mean	Mean/Var	Mean/Var/Skew	Mean/Var/Skew/Kur
λ (rationality parameter)	1.296	1.310	1.318	1.367
β (response to mean)	0.193	0.180	0.182	0.173
γ (response to variance)		0.018	0.017	0.008
δ (response to skewness)			0.046	0.038
κ (response to kurtosis)				0.068
Log Likelihood	7194.8	7167.1	-7163.8	-7113.5
No. of Participants	304	304	304	304
No. of Observations	785	785	785	785

4.2.4. Hypothesis Tests We begin by re-testing H3.1 using the data from Study 3B. Recall that a (random) subset of choices in Study 3B presented participants with all three pieces of information (left tail, midrange, right tail). However, different from Study 3A, participants were not provided with the probabilities of each value within the left tail, the midrange, or the right tail.¹³ Structural estimation results are shown in Table 8. Several observations are in order. First, all parameter estimates are quite close to the ones presented in Table 7. That is, the response to each moment is approximately the same with complete probabilistic information vs. with “coarse” information about the tail and midrange probability mass. Second, the rationality parameter is somewhat lower compared to the full information case, suggesting that decisions are somewhat noisier. Third, adding the response to each moment (variance, skewness, kurtosis) significantly improves model fit (LR tests, all $p < 0.01$). However, among the three higher moments, kurtosis continues to be the most important predictor of behavior: Log-likelihood improves by 27.7, 3.3, and 50.3, respectively, as we add each higher moment to the mean-only model.

We continue with H3.2., i.e., examine distributional information preferences (preference for learning about the left tail, midrange, or right tail). Summary statistics are in Appendix B.2. To test H3.2 we conduct a series of Binomial tests examining differences in the preference ordering in each treatment. All pairwise statistics on the preference orderings among the three pieces of information, as well as the corresponding p -values are in Table 9. The test results strongly reject H3.2: across the two treatments, 78% of participants prefer to receive left tail information to midrange ($p \ll 0.001$), and 68% of participants prefer right tail to left tail ($p \ll 0.001$). Further, there is a slight preference for left tail information over midrange, with the comparison being statistically significant for the pooled data ($p = 0.038$). Together, these tests suggest that people have a consistent preference for right-tail information relative to both left-tail and midrange information.

We next test H3.3, i.e., the responses to different types of incomplete information about a waiting time distribution. Detailed statistical comparisons and discussion are in Appendix B.4. To test H3.3a and H3.3b,

¹³ Note that the probability that each participant saw all three pieces of information was not affected by their decision; therefore, the analysis presented here does not suffer from selection effects.

we perform random effects regressions with a single explanatory variable (tail/midrange probability mass) and control for treatment. The results fully support H3.3a and partially support H3.3b. In particular, the response to midrange probability mass increases is not always significantly different from 0, suggesting that people do not respond as strongly to a reduction in variance, as they do to a reduction in means. Third, the response partially support H3.3c and reject H3.3d. That is, tail information is sometimes weighted more than midrange, and there is no asymmetric response to the tails. Finally, the results partially support H3.3e. The response to the right tail is significantly stronger for long waits, but not under all information scenarios. Finally, we do not find significant support for H3.3f; that is, information preferences do not predict the strength of one's response to that information.

4.2.5. Additional Analysis In addition to examining behaviors separately in Studies 3A and 3B, we can further examine the effects of information completeness. To do so, we pool the data from Studies 3A and 3B and examine the response to receiving distributional information relative to scenarios with no such information (Min/Max information only). In particular, we will focus on three scenarios:

- In both, Study 3A and Study 3B, a subset of scenarios presented participants with no distributional information (Min/Max information only).
- In Study 3A, a subset of scenarios presented participants with complete distributional information.
- In Study 3B, a subset of scenarios presented participants with “coarse” distributional information, i.e., probability mass in the left tail, midrange and right tail, but not the individual probabilities.

By pooling the data from Studies 3A and 3B, we can therefore evaluate whether receiving more detailed information about the waiting time distribution affects how people evaluate the wait. The results are in Table 10. We observe that both, coarse and complete distributional information coefficients are significantly negative (both $p \ll 0.01$ in col. 1 and 2), suggesting that receiving probabilistic information indeed has a positive effect on the evaluation of a wait (negative coefficients indicate that participants demand smaller monetary amounts to be compensated for a wait). Note that this result stands in contrast to Study 1 (Figure 2 and Table 3), where probabilistic information was not found to affect decisions. Further, the difference between the coarse information scenarios of Study 3B and the complete information scenarios of Study 3A is minimal and not statistically significant ($p > 0.1$).

Table 9 Study 3B: Information Preferences (Test of H3.2)

Treatment	Right Tail \succ Midrange		Left Tail \succ Midrange		Right Tail \succ Left Tail	
	Frequency	Test = 50%	Frequency	Test = 50%	Frequency	Test = 50%
High Range	73.20%	$\ll 0.001$	54.90%	0.129	63.40%	0.001
Low Range	82.78%	$\ll 0.001$	55.63%	0.096	73.51%	$\ll 0.001$
Pooled	77.96%	$\ll 0.001$	55.26%	0.038	68.42%	$\ll 0.001$

Note: The χ^2 test for differences in the preference ordering between the High vs. Low Range treatment cannot reject that preferences are the same ($p = 0.207$).

In sum, we find that the value of probabilistic information depends on the complexity of the probability distribution. While probabilistic information does not affect decisions for simple binary gambles (as in Study 1), it can have substantial effects when participants face a more complex range of possibilities (as in Study 3) and do not have strong priors about the shape of the distribution.

5. Integrated Discussion and Conclusions

The results of our studies can be used towards developing new, behaviorally grounded models of customer behavior in service systems, as well as towards deriving new insights for service design in practice. In this section, we summarize these results and discuss their implications.

5.1. Summary of Key Results and Discussion

How does the distribution of waiting times affect behavior? At a high level, our findings suggest that mean-only models oversimplify people's evaluations of waiting time. Our findings highlight the importance of considering the variability of waiting times, rather than just the averages. While this result may not be surprising given prior literature (e.g., Leclerc et al. 1995, Kroll and Vogt 2008, Flicker and Hannigan 2022), much of the extant evidence has been based on hypothetical decisions, and none of the existing studies include both time and money components of the service offering, as we have done. Our results show that variability has robust effects on behavior across a wide range of distributions (discrete, continuous, symmetric, asymmetric) and information scenarios (complete/incomplete information). Further, we have shown that the amount of variability matters: while a large amount of uncertainty reduces the value of a wait, small amounts do not significantly affect valuations.

Moving beyond the first two moments, we have shown that the specific distributional form can also matter. In particular, our second study revealed a preference for binary waiting time distributions (either short or

Table 10 Studies 3A + 3B: Value of Information

Dependent variable: Switching Amount (USD)		
Omitted Category: Min/Max only information	(1)	(2)
Coarse information (Study 3B)	−0.114*** (0.040)	−0.162*** (0.051)
Complete information (Study 3A)	−0.121*** (0.034)	−0.182*** (0.052)
Age		0.003 (0.003)
Male		−0.173** (0.084)
Education		0.044 (0.048)
Income		0.032 (0.027)
Constant	2.518*** (0.061)	2.294*** (0.165)
R^2	0.047	0.074
No. Observations	1113	1113
No. Subjects	523	523

Notes: Random effects regression coefficients are reported (Standard errors in parentheses). Asymmetric distributions (L1-L4 and R1-R4) and partial information scenarios are excluded. Round, distributional shape and treatment are controlled for. *, ** and *** denotes significance at the 10%, 5% and 1% level, respectively.

long waits) over continuous distributions, even when the mean and variance are held constant. One of the continuous probability distributions that we examined was the exponential distribution. We were interested in the response to exponentially distributed waiting times given their ubiquitous use in queueing models (Karlin and McGregor 1958, Kleinrock 1976, Whitt 2002) and their representativeness of certain real-world service processes, for example in call centers (Brown et al. 2005). Comparing the response to uniform vs. exponential waits, we found that, on average, the two were evaluated similarly, suggesting that the skewness of the waiting time distribution is not a key factor in explaining behavior. However, for the exponential distribution, people responded primarily to averages, but not to variance. A plausible explanation of this behavior is that for distributions with a long right tail, people react more strongly to maximum wait times, rather than to standard measures of variability, such as variance.

In our third study we examined a richer set of distributions. These distributions varied systematically in each of the first four moments (mean, variance, skewness, kurtosis), allowing us to estimate a series of utility models and better characterize decisions. The main result of this study is that while mean and variance explained a large portion of the variation in the data, accounting for skewness only minimally improved model fit. However, the fourth moment (kurtosis) explained a large portion of behavior. Indeed, the mean-kurtosis model was found to fit the data better than the mean-variance model (and almost as well as the full four-parameter model), suggesting that the thickness of the tails of a waiting time distribution is a key predictor of behavior. Further reinforcing this result, we found that decision-makers were particularly sensitive to and concerned about the maximum possible wait time, and responded more strongly to changes in the right tail than the left tail or central tendency.

What is the role of information completeness? Given that most practical settings do not present customers with full distributional information, we were also interested in behavior under incomplete information where information was shared with decision-makers in more aggregate, coarse form. If our results regarding preferences for higher moments hold in this limited information setting, then we can be more certain that incorporating these preferences into queueing models will produce results that are consistent with real-world decisions and outcomes. We found that preferences in the limited information setting were remarkably similar to the complete information setting. The mean and kurtosis continued to be the most important predictors of behavior, followed by variance and skewness. Further, we found that for simple, binary distributions providing probabilistic information did not increase service valuations. However, for more complex distributions, we found that probabilistic information - even in relatively coarse form - made the service offering more desirable.

5.2. Implications for Service Design

Our findings have implications for managing waiting times in a variety of service settings. The result that small amounts of variation do not affect decisions suggests that service providers can introduce some variability in wait times without having to fear negative consequences. This may be relevant for services that


involve delivery logistics or transportation. For example, in ridesharing, offering a range of wait times (e.g., “Your driver will arrive in 5-7 minutes”) may be more accurate than a single estimate. As long as the range is not too wide, customers may not heavily penalize the uncertainty. Further, the result that kurtosis is a salient metric for customer decisions suggests that minimizing the risk of extremely long delays should be a key priority, even if this means accepting a slightly higher average wait time or increased variance. Such improvements can be achieved by reallocating staffing to the periods where the tail may grow large, rather than to periods with high waiting time averages. Appropriate training and transfer policies may also facilitate the reduction of tail probabilities (See Dada et al. 2023, Hathaway et al. 2023, for related modeling frameworks).

Second, our results can help improve the effectiveness of communicating waiting times to customers. One of our results was that simple binary waits were preferred to continuous distributions, suggesting that customers may prefer distributions that can be simplified into a small number of discrete outcomes, rather than a complex range of possibilities. For example, e-commerce retailers sometimes encourage customers to choose slower, more environmentally sustainable modes of delivery. Such delivery options can be made more attractive when they feature a limited number of possible waiting times. Focusing customer attention on a handful of possible delivery times (e.g., “typical” delivery times) is likely more effective than communicating an interval for possible delivery times. Conversely, presenting customers with a complex range of potential wait times may help nudge them towards expedited delivery options. Among the different types of information, providing maximum or “worst-case” wait information is likely to be the most salient factor driving customer preferences. For service providers who want to maximize the uptake of a service offering, providing right-tail information may therefore not be desirable. More generally, our results provide empirical motivation for models that explicitly account for behavioral factors in modeling the response to waiting time information (see Debo et al. 2023, Guda et al. 2023, for recent examples).

5.3. Numerical Case Study: Priority Rules and Customer Response to Higher Moments

We include a numerical case study that illustrates the broader implications of our findings. The case study is based on a discrete event simulation of a single-server system with two customer types. Each customer type is characterized by its arrival rate $\lambda_i, i = 1, 2$, and by its (deterministic) service time $\mu_i, i = 1, 2$. We focus on scenarios with $\lambda_1 = \lambda_2$ and $\mu_1 < \mu_2$, and examine heavy-traffic systems where utilization ranges between 0.8 and 0.97. To focus on the response to waiting times, we assume that there is no admission fee. To evaluate the disutility of waiting in this service system we compute the sojourn time distribution under different priority rules and evaluate the response to that distribution based on the utility parameters estimated in Table 7 (full specification):

$$U(t_r^s) = -(0.150\mathbb{E}[t_r^s] + 0.010\mathbb{V}[t_r^s] + 0.033\mathbb{S}[t_r^s] + 0.074\mathbb{K}[t_r^s]), \quad (5.1)$$

Figure 5 Case Study: Priority Rules and Distributional Moments

Note: Results are based on a discrete event simulation performed for utilization levels of 0.8, 0.825, 0.85, 0.875, 0.9, 0.925, 0.95, and 0.975 with 100,000,000 customer arrivals for each run. Customer disutility is computed based on eq. 5.1.

where t_r^s is the random variable denoting the sojourn time in the steady state (averaged across both customer types) under priority rule r and $\mathbb{E}[t_r^s]$, $\mathbb{V}[t_r^s]$, $\mathbb{S}[t_r^s]$ and $\mathbb{K}[t_r^s]$ are the first four moments of the distribution characterizing that random variable.

In Figure 5 we compare the moments of the sojourn time distributions, and the customer response (disutility based on eq. 5.1), under two priority rules: the First Come First Serve (FCFS) rule, and the Shortest Processing Time (SPT) rule. The classic result is that the SPT rule minimizes average waiting costs (Smith et al. 1956, Pinedo 2012). Our analysis suggests that this result may not hold when considering higher-order moments. Panel a) examines the case where $\mu_1 = 1$ and $\mu_2 = 9$. In this setting, the classic result appears to hold. The SPT rule is quite effective in reducing average waiting times (up to 48% reduction for high utilization scenarios, relative to FCFS). Despite increasing skewness and kurtosis relative to FCFS, the overall disutility under SPT remains lower. Panel b) presents a scenario with a smaller service time difference between customer classes ($\mu_1 = 3, \mu_2 = 7$). In this case, SPT leads to higher customer disutility. SPT's advantage in reducing means is somewhat less pronounced (relative to the $\mu_1 = 1, \mu_2 = 9$ case), while the remaining three moments are all larger under SPT than under FCFS. Therefore, the composite response to the four moments favors the FCFS rule. Together, these comparisons highlight the importance of considering higher-order moments in the design of service systems.

5.4. Limitations and Future Directions

Our studies were conducted in a controlled, context-free experimental setting. The lack of context and the simplicity of the choices is an important and necessary first step in understanding focal trade-offs; however, this also means that some of the findings may not fully generalize to certain real-world service environments, which may include factors like emotional stakes, sunk costs, social preferences, and other contextual factors. Further, our experiments focused on virtual queue settings where the (remaining) waiting time is known but the queue advancements are not displayed. Thus, our results may not extrapolate to physical queues where the advancement of the queue may be more salient than the passage of time (see Kumar et al. 1997, Althenayyan et al. 2022, for related experiments). Finally, modeling and experimentally examining richer server behaviors (as in Shunko et al. 2018, Rosokha and Wei 2024) and explicitly accounting for the interactions between servers and customers (as in Frazelle and Katok 2024) may be a useful next step.

Our study has focused solely on people's decisions *prior to* joining a service system. We did not examine retrospective or in-process evaluations of a wait (Carmon and Kahneman 1996, Luo et al. 2022) or potential violations of customer expectations, caused by unexpected delays or overly-optimistic reference points (Yang et al. 2018, Yu et al. 2022, Debo et al. 2023). Real-world service experiences often involve dynamic expectations that evolve over time (Das Gupta et al. 2016, Deshmane et al. 2023), which can shape the overall experience in ways not captured by our experiment. Future research could explore how decisions are influenced by the discrepancy between expected and experienced waiting times, and how service providers can manage these evolving expectations through effective communication and service design. Incorporating these richer, more contextual elements of the customer experience is an important direction for extending the insights from our study.

References

- Abdellaoui M, Baillon A, Placido L, Wakker PP (2011) The rich domain of uncertainty: Source functions and their experimental implementation. *American Economic Review* 101(2):695–723.
- Abdellaoui M, Kemel E (2014) Eliciting prospect theory when consequences are measured in time units: “time is not money”. *Management Science* 60(7):1844–1859.
- Afèche P, Baron O, Kerner Y (2013) Pricing time-sensitive services based on realized performance. *Manufacturing & Service Operations Management* 15(3):492–506.
- Aksin OZ, Gencer B, Gunes ED (2022) How observed queue length and service times drive queue behavior in the lab. *Available at SSRN 3387077* .
- Allon G, Kremer M (2018) Behavioral foundations of queueing systems. *The handbook of behavioral operations* 9325:325–366.
- Althenayyan A, Cui S, Ulku S, Yang L (2022) Not all lines are skipped equally: an experimental investigation of line-sitting and express lines. *Georgetown McDonough School of Business Research Paper* (4179751).

- Andersen S, Harrison GW, Lau MI, Rutström EE (2006) Elicitation using multiple price list formats. *Experimental Economics* 9:383–405.
- Ansari S, Debo L, Ibanez M, Iravani S, Malik M, et al. (2022) Under-promising and over-delivering to improve patient satisfaction at emergency departments: Evidence from a field experiment providing wait information. *Tuck School of Business Working Paper* (4135705).
- Åstebro T, Mata J, Santos-Pinto L (2015) Skewness seeking: risk loving, optimism or overweighting of small probabilities? *Theory and Decision* 78:189–208.
- Baucells M, Osadchiy N, Ovchinnikov A (2017) Behavioral anomalies in consumer wait-or-buy decisions and their implications for markdown management. *Operations Research* 65(2):357–378.
- Ben-Akiva M (1974) Structure of passenger travel demand models (doctoral dissertation, massachusetts institute of technology) .
- Brown L, Gans N, Mandelbaum A, Sakov A, Shen H, Zeltyn S, Zhao L (2005) Statistical analysis of a telephone call center: A queueing-science perspective. *Journal of the American statistical association* 100(469):36–50.
- Brünnner T, Levínský R, Qiu J (2011) Preferences for skewness: evidence from a binary choice experiment. *The European Journal of Finance* 17(7):525–538.
- Buell RW (2021) Last-place aversion in queues. *Management Science* 67(3):1430–1452.
- Camerer C, Weber M (1992) Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of risk and uncertainty* 5:325–370.
- Carmon Z, Kahneman D (1996) The experienced utility of queuing: real time affect and retrospective evaluations of simulated queues. *Duke University: Durham, NC, USA* .
- Chen DL, Schonger M, Wickens C (2016) oTree – An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance* 9:88–97.
- Dada M, Hathaway B, Kagan E (2023) The omnichannel service desk: Live agents, chatbots or both? *Chatbots or Both* .
- Das Gupta A, Karmarkar US, Roels G (2016) The design of experiential services with acclimation and memory decay: Optimal sequence and duration. *Management Science* 62(5):1278–1296.
- Debo L, Shumsky RA, Ansari S, Iravani S, Liu Z (2023) *Wait Time Information Design* (SSRN).
- Debo L, Veeraraghavan S (2014) Equilibrium in queues under unknown service times and service value. *Operations Research* 62(1):38–57.
- Deshmane A, Martinez-de Albeniz V, Roels G (2023) Intertemporal spillovers in consumer experiences: Empirical evidence and service design implications. *Available at SSRN 4507191* .
- Eckel CC, Grossman PJ (2008) Men, women and risk aversion: Experimental evidence. *Handbook of Experimental Economics Results* 1:1061–1073.

- Edelson NM, Hilderbrand DK (1975) Congestion tolls for poisson queuing processes. *Econometrica: Journal of the Econometric Society* 81–92.
- Einhorn HJ, Hogarth RM (1986) Decision making under ambiguity. *Journal of business* S225–S250.
- Ellsberg D (1961) Risk, ambiguity, and the savage axioms. *The quarterly journal of economics* 75(4):643–669.
- Estrada Rodriguez A, Ibrahim R, Kremer M (2023) Persuasive communication in social service operations. Available at SSRN 4672269 .
- Festjens A, Bruyneel S, Diecidue E, Dewitte S (2015) Time-based versus money-based decision making under risk: An experimental investigation. *Journal of Economic Psychology* 50:52–72.
- Flicker B, Hannigan C (2022) On people’s utility over wait fundamentals and information.
- Fox CR, Tversky A (1995) Ambiguity aversion and comparative ignorance. *The quarterly journal of economics* 110(3):585–603.
- Frazelle AE, Katok E (2024) Paid priority in service systems: Theory and experiments. *Manufacturing & Service Operations Management* 26(2):775–795.
- Fréchette GR, Schotter A, Trevino I (2017) Personality, information acquisition and choice under uncertainty: An experimental study. *Economic Inquiry* 55(3):1468–1488.
- Friend I, Westerfield R (1980) Co-skewness and capital asset pricing. *The Journal of Finance* 35(4):897–913.
- Gonzalez R, Wu G (1999) On the shape of the probability weighting function. *Cognitive psychology* 38(1):129–166.
- Guda H, Dawande M, Janakiraman G (2023) The economics of process transparency. *Production and Operations Management* 32(6):1812–1829.
- Harrison GW, Cox JC (2008) *Risk aversion in experiments* (Emerald Group Publishing).
- Harvey CR, Siddique A (2000) Conditional skewness in asset pricing tests. *The Journal of finance* 55(3):1263–1295.
- Hassin R (2016) *Rational queueing* (CRC press).
- Hassin R, Haviv M (2003) *To queue or not to queue: Equilibrium behavior in queueing systems*, volume 59 (Springer Science & Business Media).
- Hathaway BA, Kagan E, Dada M (2023) The gatekeeper’s dilemma: “when should i transfer this customer?”. *Operations Research* 71(3):843–859.
- Holt CA, Laury SK (2002) Risk aversion and incentive effects. *American Economic Review* 92(5):1644–1655.
- Ibrahim R, Estrada Rodriguez A, Zhan D (2024) On customer (dis) honesty in priority queues: The role of lying aversion. *Management Science* .
- Ingolfsson A, Mandelbaum A, Schultz K, Yom-Tov GB (2023) Preface to the special issue on behavioral queueing science: The need for a multidisciplinary approach. *Operations Research* 71(3):791–797.
- Kagan E, Dada M, Hathaway B (2022) Ai chatbots in customer service: Adoption hurdles and simple remedies. *Johns Hopkins Carey Business School Research Paper* (23-03).

- Kagel JH, Roth AE (2020) *The handbook of experimental economics, volume 2* (Princeton university press).
- Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):363–391.
- Karlin S, McGregor J (1958) Many server queueing processes with poisson input and exponential service times. .
- Kleinrock L (1976) *Queueing systems, volume II: Computer applications* (Wiley).
- Kraus A, Litzenberger RH (1976) Skewness preference and the valuation of risk assets. *The Journal of finance* 31(4):1085–1100.
- Kremer M, Debo L (2016) Inferring quality from wait time. *Management Science* 62(10):3023–3038.
- Kroll EB, Vogt B (2008) Loss aversion for time: an experimental investigation of time preferences. *Working Paper Series* .
- Kumar P, Kalwani MU, Dada M (1997) The impact of waiting time guarantees on customers' waiting experiences. *Marketing science* 16(4):295–314.
- Leclerc F, Schmitt BH, Dube L (1995) Waiting time and decision making: Is time like money? *Journal of consumer research* 22(1):110–119.
- Luce RD (1959) *Individual Choice Behavior: A Theoretical Analysis* (New York: Wiley).
- Luo J, Valdés L, Linardi S (2022) Experienced and prospective wait in queues: A behavioral investigation. *Available at SSRN 4169028* .
- Markowitz H (1952) The utility of wealth. *Journal of political Economy* 60(2):151–158.
- Markowitz H (2014) Mean–variance approximations to expected utility. *European Journal of Operational Research* 234(2):346–355.
- McFadden D (1974) The measurement of urban travel demand. *Journal of public economics* 3(4):303–328.
- Mitton T, Vorkink K (2007) Equilibrium underdiversification and the preference for skewness. *Review of Financial Studies* 20(4):1255–1288.
- Nakamura Y (2015) *Journal of Economic Theory* 160:536–556.
- Naor P (1969) The regulation of queue size by levying tolls. *Econometrica: journal of the Econometric Society* 15–24.
- Osadchiy N, Bendoly E (2015) Are consumers really strategic? implications from an experimental study. *Implications from an Experimental Study (October 2015)* .
- Özer Ö, Zheng Y (2012) Behavioral issues in pricing management .
- Özer Ö, Zheng Y (2016) Markdown or everyday low price? the role of behavioral motives. *Management Science* 62(2):326–346.
- Palan S, Schitter C (2018) Prolific. ac—a subject pool for online experiments. *Journal of Behavioral and Experimental Finance* 17:22–27.

- Paolacci G, Chandler J (2014) Inside the turk: Understanding mechanical turk as a participant pool. *Current directions in psychological science* 23(3):184–188.
- Pinedo ML (2012) *Scheduling*, volume 29 (Springer).
- Rosokha Y, Wei C (2024) Cooperation in queueing systems. *Management Science* .
- Shunko M, Niederhoff J, Rosokha Y (2018) Humans are not machines: The behavioral impact of queueing design on service time. *Management Science* 64(1):453–473.
- Smith WE, et al. (1956) Various optimizers for single-stage production. *Naval Research Logistics Quarterly* 3(1-2):59–66.
- Snowberg E, Yariv L (2021) Testing the waters: Behavior across participant pools. *American Economic Review* 111(2):687–719.
- Soman D (2001) The mental accounting of sunk time costs: Why time is not like money. *Journal of behavioral decision making* 14(3):169–185.
- Stecher J, Shields T, Dickhaut J (2011) Generating ambiguity in the laboratory. *Management Science* 57(4):705–712.
- Tversky A, Kahneman D (1992) Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty* 5:297–323.
- Ülkü S, Hydock C, Cui S (2020) Making the wait worthwhile: Experiments on the effect of queueing on consumption. *Management Science* 66(3):1149–1171.
- Ülkü S, Hydock C, Cui S (2022) Social queues (cues): Impact of others’ waiting in line on one’s service time. *Management Science* 68(11):7958–7976.
- Veeraraghavan S, Debo L (2009) Joining longer queues: Information externalities in queue choice. *Manufacturing & Service Operations Management* 11(4):543–562.
- Wang J, Zhang ZG (2018) Strategic joining in an m/m/1 queue with risk-sensitive customers. *Journal of the Operational Research Society* 69(8):1197–1214.
- Whitt W (2002) *Stochastic-process limits: an introduction to stochastic-process limits and their application to queues* (Springer).
- Wu CC, Bossaerts P, Knutson B (2011) The affective impact of financial skewness on neural activity and choice. *PloS one* 6(2):e16838.
- Yang L, Guo P, Wang Y (2018) Service pricing with loss-averse customers. *Operations research* 66(3):761–777.
- Yu Q, Zhang Y, Zhou YP (2022) Delay information in virtual queues: A large-scale field experiment on a major ride-sharing platform. *Management Science* 68(8):5745–5757.

Appendix A: Instructions

A.1. Study 1 Instructions

This section includes the stimuli (Instructions, screenshots and protocols) for Study 1. In this set of instructions, a subject first completed the non-deterministic block, and then the deterministic block. The sequence (deterministic → non-deterministic or vice versa) was randomized for each subject. Differences between the two treatments (Uncertain and Ambiguous treatments) are described in square brackets. Instructions for Study 2 and 3 are analogous and are described at the end of this Appendix.

Welcome to this decision-making study. The study takes approximately 15 minutes to complete. Please pay close attention to the instructions. To ensure that you understand the instructions we will ask you several questions as we go along. You will only be allowed to proceed to the task if you can answer those questions. If you complete this experiment you will receive a participation payment of \$3. In addition, you will receive a bonus payment. Bonus payments range between \$1 and \$5 and depend on your decisions and on chance.

In the following, you'll face 6 decision problems. Each decision problem consists of 21 choices between two options: "Option A" and "Option B". We now describe how each decision problem works.

As noted, with each decision problem, there will be a list of decisions that you will need to make. Each decision is a paired choice between "Option A" and "Option B". Both options consist of amounts of money and lengths of time you must wait to receive the given amount of money. "Option A" is a relatively small amount of money with a short wait. "Option B" is a relatively larger amount of money, typically with a longer wait.

An example of the decisions you will make is in the table below. This example will not be played for real money or time - it is just for you to get used to the interface.

[Screenshot of Uncertain treatment:]

Option A		Option B	
Amount	Wait	Amount	Wait
Receive \$1.00	Wait 1 min.	Receive \$4.00	Wait 2 min. or 8 min. with equal chance
Receive \$1.00	Wait 1 min.	Receive \$3.50	Wait 2 min. or 8 min. with equal chance

[Screenshot of Ambiguous treatment:]

Option A		Option B	
Amount	Wait	Amount	Wait
Receive \$1.00	Wait 1 min.	Receive \$4.00	Wait 2 min. or 8 min. with unknown chance
Receive \$1.00	Wait 1 min.	Receive \$3.50	Wait 2 min. or 8 min. with unknown chance

Your task is to choose which option, A or B, that you prefer for each decision presented to you. Once you have completed all of the 6 decision problems, you will be taken to a new screen in order to determine your payoff. Please click "Next" now to continue with the instructions.

As noted earlier, you will make several choices between option A and option B. After you have made all choices for all decision problems, we will first randomly select (with equal chance) one of the decision problems. For that decision problem, we will use the option that you have chosen, A or B, to determine your payoff. For example, suppose that the row in the table below is selected for payment. Suppose also that you selected Option B for this decision. In this case, after clicking "Next", you will be taken to a waiting screen, where you will be asked to wait for either 2 minutes or 8 minutes (with equal chance). If you successfully wait for the required length of time, you will be paid \$4.00. If you do not successfully wait for the required length of time, then you will not receive any payment.

[Subject sees an example where Option B is chosen]

On the other hand, suppose that the row in the table below is selected for payment. Suppose also that you selected Option A for this decision. In this case, you will be asked to wait 1 minute, after which you will receive \$1.00.

[Subject sees an example where Option A is chosen]

Note: Because each decision problem and each question within a decision problem is equally likely to be chosen for payment, you have no incentive to lie on any question. If you lied on a question, and if that question is chosen for payment, then you would end up with the option you like less.

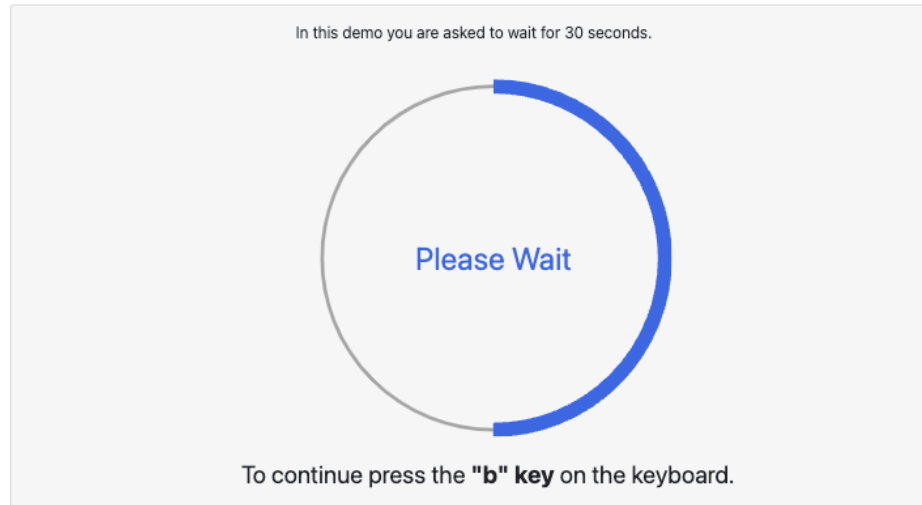
[Subject completes a comprehension quiz]

As noted earlier, you will make several choices between Option A and Option B. After you have made all choices for all decision problems, we will randomly select (with equal chance) one of the decision problems and implement your choice for that problem. If the choice that has been selected involves waiting, you will be presented with a wait. During the wait we will use attention checks. The following page demonstrates what you would experience in the waiting stage.

[The following paragraph was only shown in the Ambiguous treatment.]

As you can see below, waiting times may be uncertain. Suppose you choose option B. To determine the duration we will use a computerized coin whose likelihood of landing on heads or tails is somewhere between 0 and 100%. You can think about it just as a regular coin, but instead of having 50% probability of falling on heads or tails, that probability is X . In fact, even we (the experimenters) do not know X – it will be determined by a random algorithm after you make all your choices.

[Screenshot of Waiting Demo]



You are now ready to begin making choices. Remember that one of your choices will be selected for real payment/wait at the end of the study.

[Subject completes the first block of questions (Non-deterministic for this subject).]

In the following decision problems you will be asked to make similar choices. However, rather than having a shorter and a longer waiting times, Option B will now involve a no uncertainty. That is, the waiting time for both options will be known ahead of time. Remember that one of your choices may be selected for real payment/wait at the end of the study.


[Subject completes the second block of questions (Deterministic for this subject).]

[One of the decisions made by the subject is drawn at random. The subject experiences the time outcome of the selected option (Option A or Option B, depending on the subject's prior choice) for that decision. After completing the wait, the subject proceeds to payment and exit survey.]

A.2. Study 2 Instructions

Instructions are analogous to Study 1, with one key change. Subjects see a visual representation of the respective distribution (Binary, Uniform, or Exponential, depending on the treatment) and are trained to develop an intuitive understanding of the risks involved. Quartiles are added for continuous distribution to explain each distribution. After the training subjects answer several additional quiz questions. Additionally, subjects see the graph of the distribution during each decision (parametrized for the respective decision set). The visual representations for the $\mu = 5, \sigma = 2$ scenario are reproduced below (sizes are scaled to text width).

Figure A1 Study 2: Presentation of Probability Distributions. Wait Time in Option B ($\mu = 5, \sigma = 2$ scenario)



A.3. Study 3 Instructions

Instructions are analogous to Study 1 and Study 2, with the following differences. Before completing the first non-deterministic block (decision set 2) subjects saw the following instructions.

In the remaining decision problems you will be asked to make similar choices. However, the waiting time in Option B will now be uncertain. This means, if you choose Option B, and that choice is selected for real payment/wait, you will not know exactly how much time you will need to wait. As before, the amount of money you will be paid is larger if you choose Option B. The Tables below show the waiting times for Option A and Option B will look like. In particular, in Option B you could wait anywhere between 1 and 15 minutes. However, you do not know how likely it is that you have to wait each amount of time. For example, the chance of waiting exactly 5 minutes could be 5%, or it could be 75%. All you know is that the time cannot be less than 1 minute or greater than 15 minutes.

Option A		Option B	
Waiting Time	1 Minute	Possible Waiting Times	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, or 15 Minutes
Likelihood	100%	Likelihood of each time	Unknown

[Subject completes a short comprehension quiz.]

[Subject completes the decision set.]

In the remainder of the experiment, you will receive partial information about Option B. The table below shows three regions: 1-5 minutes (short waiting time), 6-10 minutes (moderate waiting time) and 11-15 minutes (long waiting time).

[Subject sees the following:]

Below you are asked to indicate the regions whose likelihood (%) you are most interested in learning. After you do this, we will reveal some of your choices. In particular, in each of the remaining rounds, we will reveal your top choice with a 75% chance and your second choice with a 50% chance. In other words, it is not guaranteed

Option B			
Possible Waiting Times	1-5 Min. (short wait)	6-10 Min. (moderate wait)	11-15 Min. (long wait)
Likelihood	??%	??%	??%

that you will receive all the pieces of information that you request, but it is more likely that you receive the information that you rank higher. What information would you like us to reveal?

[Subject indicates their top choice and second choice and continues with the instructions for the partial information scenarios (Scenario 3-6).]

[A screenshot of a scenario with partial information being revealed is reproduced below.]

In this round, the following information (highlighted in blue) is revealed.

Option B			
Possible Waiting Times	1-5 Min. (short wait)	6-10 Min. (moderate wait)	11-15 Min. (long wait)
Likelihood	80.0%	??%	??%

[Subject continues to making their decisions for the given distribution.]

[Subject continues to making their decisions for the given distribution.]

Appendix B: Additional Analysis

B.1. Study 1: Analysis of Gender Differences


The gender split for the average switching amount in each condition is shown in Figure B1.¹⁴ First, the graph confirms that female participants demand more monetary compensation on average - for example, for deterministic waits, the gender gap ranges between \$0.19 and \$0.32. Second, women demand higher compensation under ambiguity than under risk, and vice versa for men. A regression analysis confirms these results with men being ambiguity-seeking ($p = 0.016$), women being slightly ambiguity-averse ($p = 0.061$), and ambiguity attitudes between men and women being significantly different from each other ($p = 0.003$).¹⁵ Further exploratory analysis suggests that there are no significant differences across our male and female participants in terms of age, income, or education; nor, do we find significant interaction effects of gender with these variables.

B.2. Study 3 Distributions

Figure B2 shows the probability distributions used in Study 3A and Study 3B. In Study 3A, the distribution was described to participants via a table with all possible waiting time values and their probabilities. In Study 3B participants received information about left tail, right tail, and midrange regions (depending on their choice and luck), but not about the individual probabilities within those regions. Participants were *not* presented with graphs of distributions.

¹⁴ Three participants indicated that their gender was neither male nor female. These participants are not shown in Figure B1.

¹⁵ These comparisons are based on the same regression specification as the one used in Table 3, column (1), with the additional interaction term between treatment and gender.



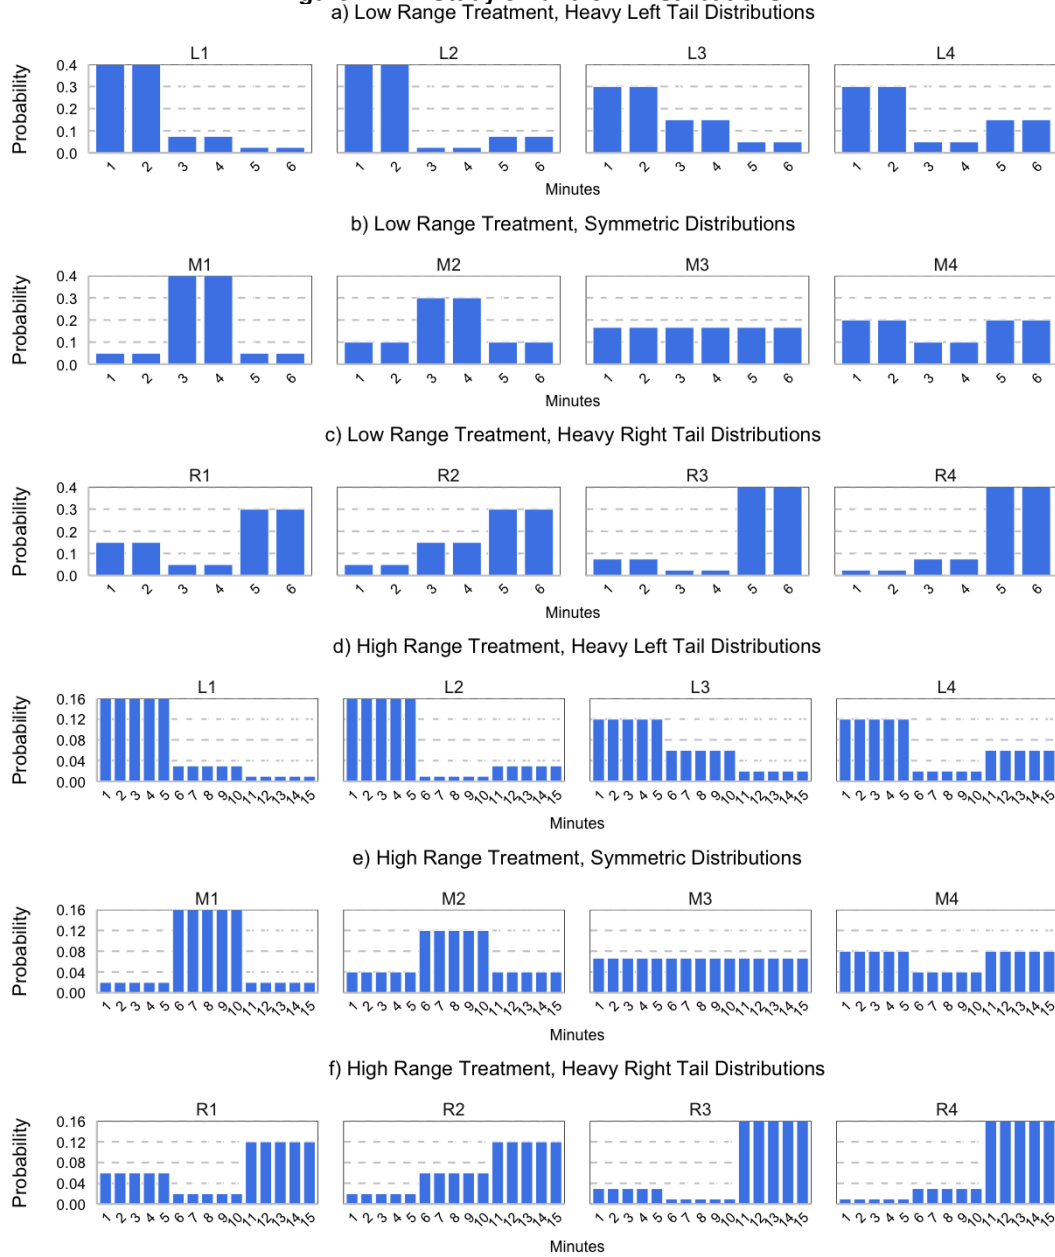
B.3. Study 3A: Summary Statistics

Figure B3 shows the switching amounts for different time distributions. A higher switching point indicates a higher cost of waiting in a given scenario. Recall that in one of the preliminary rounds participants evaluated the waits in the absence of the distributional information. These decisions are shown as horizontal dashed lines. The remaining 24 bars show the switching amounts for each of the 24 distributions presented in Table 6.

Consider first the leftmost four sets of bars. These bars show scenarios L1-L4, i.e., scenarios in which the waiting time distribution is left-tail heavy; that is, lower waiting times (1-2 minutes in the Low Range treatment and 1-5 minutes in the High Range treatment) are more likely than longer times. Unsurprisingly, revealing this information leads to a lower cost of waiting relative to the no-information scenario. Next, consider the middle four sets of bars, M1-M4. These are symmetric distribution scenarios (U-shaped, flat, and bell-shaped, respectively). In these scenarios, revealing waiting time information does not necessarily result in a lower switching amount in the Low Range treatment, but does so in the High Range treatment. Finally, consider the rightmost four sets of bars, R1-R4, which present the response to less favorable (right-tail heavy) distributions. In the High Range treatment, participants respond quite negatively to the revealed information with all four bars extending above the no-information scenario. In contrast, the response in the Low Range treatment appears to be minimal.

B.4. Study 3B: Summary Statistics and Additional Hypothesis Tests


B.4.1. Summary Statistics on Information Preferences Figure B4 presents the relative frequencies of top and second choices by treatment. In both the High Range and Low Range treatments, the top choice for the majority of participants (56.2% and 66.9%, respectively) is the right tail of the waiting time distribution, suggesting a tendency to focus on the upper bound of potential wait times, rather than the central tendency or lower tail. However, the second choice preferences are somewhat different across treatments, with the midrange being the most common second choice (47.1%) in the High Range treatment, while the Low Range treatment shows a more even split, with midrange (40.4%) and left tail (37.1%) being closely tied. This suggests that as the overall range of potential waiting times increases,

Figure B2 Study 3A and 3B: Distributions

there is a slight shift away from the tails and towards the midrange, though the interest in the right tail dominates across both treatment conditions.

Figure B5 shows the switching amounts for different time distributions. Recall that a higher switching point indicates a higher cost of waiting in a given scenario. Panel (a) focuses on scenarios in which either no information (dashed line) or all three pieces of information were revealed. The data show that, unsurprisingly, revealing information that makes a wait appear more desirable results in lower switching amounts. Consider first the leftmost four sets of bars. These bars show scenarios in which the waiting time distribution is left-tail heavy; that is, lower waiting times (1-2 minutes in the Low Range treatment and 1-5 minutes in the High Range treatment) are more likely than longer times. Unsurprisingly, revealing this information leads to a lower cost of waiting relative to the no-information scenarios. Next, consider the

Figure B3 Study 3A: Switching Amounts by Treatment and Type of Wait



middle four sets of bars. These are symmetric distribution scenarios (flat, bell-shaped, and U-shaped, respectively). In these scenarios, revealing waiting time information does not necessarily result in a lower switching amount in the Low Range treatment, but does so in the High Range treatment. Finally, consider the rightmost four sets of bars, which present the response to less favorable (right-tail heavy) distributions. In the High Range treatment, participants respond quite negatively to the revealed information with all four bars extending above the no-information scenario. In contrast, the response in the Low Range treatment appears to be minimal. In sum, participants appear to be incorporating the information into their decisions, with the response being somewhat stronger upon seeing the right tail, and in the High Range treatment.

Next, consider panel (b) of Figure B5. This panel plots the partial information (probability mass in the left tail, midrange, or right-tail) revealed to a participant against the switching amount, along with a LOESS (LOcally Estimated Sums of Squares) fit. Each scatter point corresponds to one decision. The left part of panel (b) shows that participants in both treatments require a smaller monetary payment as the probability mass in the left tail increases. However, the response in the Low Range treatment appears to be substantially more flat and noisy than in the High Range

Figure B4 Study 3B: Information Preferences


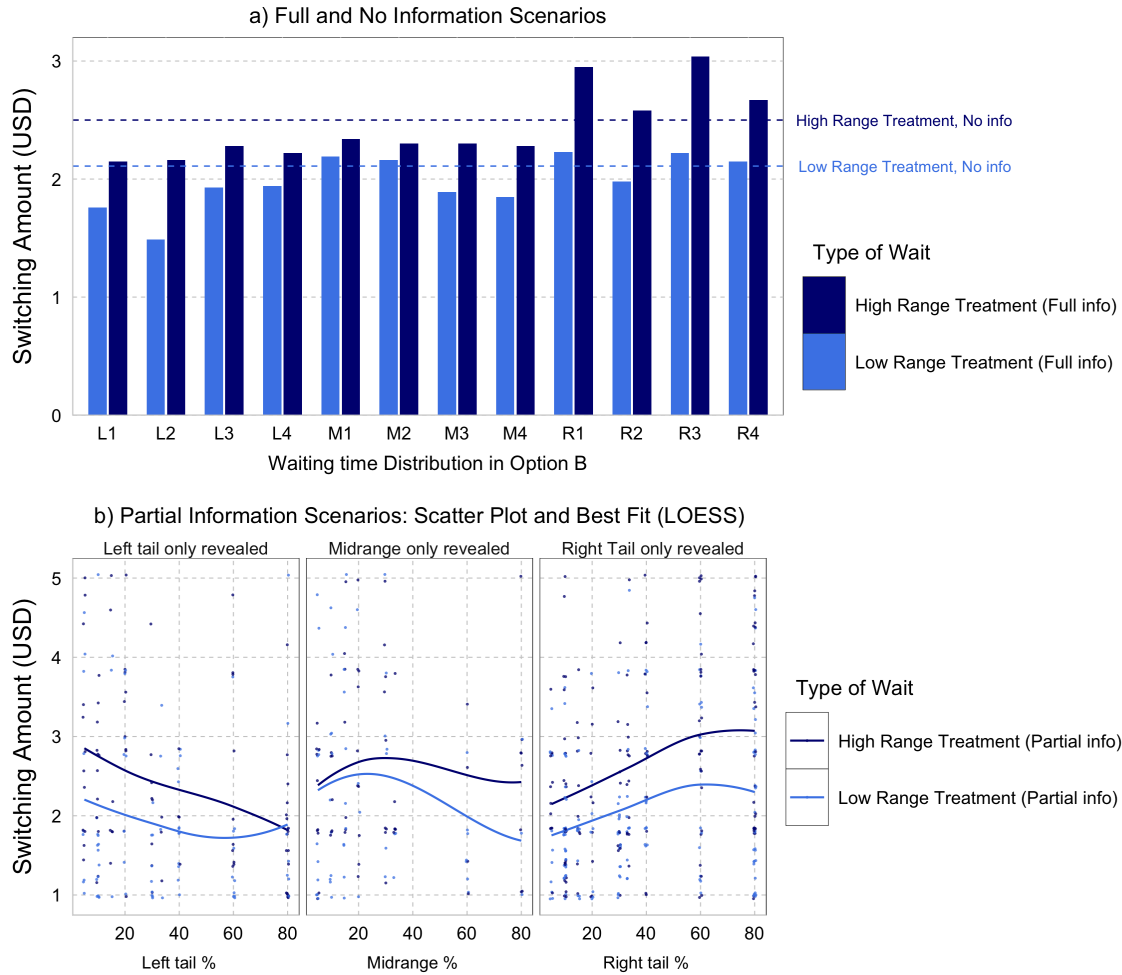



Figure B5 Study 3B: Switching Amounts by Treatment and Type of Wait

treatment. The middle part of panel (b) shows the response to midrange information. The response appears quite noisy, though there is a decreasing trend in the Low Range treatment. The right part of panel (b) shows that, as the right tail probability mass goes up, people demand a greater monetary payment. Consistent with our hypotheses, the response to the right tail appears to be somewhat more defined than the response to the left tail or midrange, particularly for the High Range treatment.

B.4.2. Tests of H.3.3 We present the marginal effects for each of the subhypothesis. We begin with H3.3a. Panel a) of Figure B6 focuses on full information scenarios and shows the marginal effects of changes in the left tail probability, midrange, and right tail probability on switching amounts. Panel b) repeats the analysis for partial information scenarios. In both cases, we perform random effects regressions with a single explanatory variable (tail/midrange) and control for treatment. The results can be summarized as follows. First, the results support H3.3a. That is, participants react positively (negatively) to increases in the left (right) tail probability mass. All four coefficients on the tails are significantly different from 0 ($p < 0.01$). Second, the results partially support H3.3b. In both full information (panel a) and partial information (panel b) scenarios, switching amounts decrease as the midrange probability mass goes up.


Figure B6 Study 3B: Tail/Midrange Effects on Switching Amount



However, the effect is not statistically significant in the full information scenario and quite close to the 0.05 cutoff in the partial information scenario ($p = 0.345$ and $p = 0.034$, respectively). Third, the results partially support H3.3c. The response to tails is significantly stronger than the response to midrange in the full information case (both $p \ll 0.001$) but not in the partial information case ($p = 0.751$ and $p = 0.229$). Fourth, the results reject H3.3d. The response to right tail is not significantly stronger (in absolute value terms) than the response to left tail ($p = 0.342$ and $p = 0.262$). Overall, these comparisons suggest that people update their beliefs about the waiting time distribution when provided with tail and midrange information and weigh both tails similarly. However, when they receive detailed information about the waiting time distribution, they tend to focus more on the tails than on the midrange.

Next, we turn to H3.3e., i.e., potential differences in behavior between the Low Range and the High Range treatments. The marginal effects of the right tail probability mass on switching points in each treatment are shown in Figure B7. We have hypothesized that the right tail effect is stronger for the High Range treatment than in the Low Range treatment. Figure B7 suggests that this is true for the full information scenario, but not for the partial information scenario. Indeed, the right tail effect is approximately twice as strong in the High Range treatment than in the Low Range treatment in the full information scenario (the difference between coefficients is significantly different from zero, $p \ll 0.001$). However, the effect is approximately the same across treatments in the partial information scenario ($p = 0.714$). Finally, we test H3.3f in which we hypothesized that people may respond more strongly to the information they ranked higher. Contrary to H3.3f, our tests show that switching points are not significantly related to information preferences (all p -values above 0.2).

Figure B7 Study 3B: Right Tail Effects on Switching Amount by Treatment



B.5. Structural Estimation Details

In the main text (§4) we have examined the utility parameters that characterize the decision-maker's preferences in Study 3. In this section, we expand on this analysis. We discuss the theoretical and econometric foundations of our utility model, include utility model estimation results for Study 1 and Study 2, and show that they are consistent with our main results.

B.5.1. Theoretical and Econometric Foundations A natural starting point for estimating the utility over time and money components is a utility-based approach. For example, one could assume that the utility over money-time pairs (m, t) is given by $u(m, t) = m - \beta v(t)$. We could then, for example, make assumptions on $v(t)$, such as constant relative risk aversion (CRRA) or constant absolute risk aversion (CARA). Given our data, we first show that this approach presents us with problems.

To see this, fix a CRRA utility,

$$v(t) = \frac{t^{1-\gamma} - 1}{1-\gamma} \quad (\text{B.5.1})$$

and consider the following four waiting scenarios, (1) $t = 5$, (2) $t = 10$, (3) $t \in \{3, 7\}$ with equal chance and (4) $t \in \{5, 15\}$ with equal chance, which were decision problems faced by subjects in the Binary Treatment of Study 2. Denote by \bar{m}_i , $i \in \{1, 2, 3, 4\}$ to be the average amount of money required to switch between Options A and B. From Figure 3, we know that $\bar{m}_1 = 1.79$, $\bar{m}_2 = 2.26$, $\bar{m}_3 = 1.90$, and $\bar{m}_4 = 2.58$. The fact that $\bar{m}_4 > \bar{m}_2$ and $\bar{m}_3 > \bar{m}_1$ implies that $\gamma < 0$. That is, holding the average constant, introducing variability in waiting makes a prospective wait less attractive.

Using the CRRA utility, we can find values for $\theta = (\beta, \gamma)$ to exactly match the switching points for scenarios (1) and (3), and also for scenarios (2) and (4).¹⁶ In the former case, we obtain $\theta_{13} \approx (0.0758, -0.879)$, while in the latter case, we obtain $\theta_{24} \approx (0.0252, -1.004)$. So for longer average waits (e.g., 10 minutes versus 5 minutes), subjects appear to be more risk averse in the time dimension, but there is a level effect where they are less sensitive overall. On the other hand, if we tried to find a single set of parameters to best-fit all four scenarios, then we obtain $\theta_{1234} \approx (0.3306, 0.481)$; importantly, we no longer obtain $\gamma < 0$, which is inconsistent with our data.

For this reason, in our estimation model presented in the main text (eq. (1.2)-(1.3) in §1 and estimation results in §4) we use a moment-based model, in which utility over time depends on the average wait, the variance of the wait as well as higher moments of the waiting time distribution. This is similar to the approach taken by Mitton and Vorkink (2007) in the context of studying preferences for skewed lotteries in the money dimension. We apply it here to the time dimension. We also note that the moment-based approach is the dominant approach for modeling investor preferences in the portfolio selection literature in finance (See, for example, the three-moment CAPM model, which adds skewness to the classic mean-variance CAPM, see Kraus and Litzenberger 1976, Friend and Westerfield 1980, Harvey and Siddique 2000).

Following the moment-based approach, the utility of the waiting scenario (m, t) is given by:

$$U(m, t) = m - \beta \mathbb{E}[t] - \gamma \mathbb{V}[t] - \delta \mathbb{S}[t] - \kappa \mathbb{K}[t] \quad (\text{B.5.2})$$

¹⁶ These problems lead to systems of two equations in two unknown variables and can be solved numerically.

where $\mathbb{E}[\cdot]$, $\mathbb{V}[\cdot]$, $\mathbb{S}[\cdot]$ and $\mathbb{K}[\cdot]$ are the mean, variance, skewness and kurtosis operators.¹⁷ We expect the coefficients, $(\beta, \gamma, \delta, \kappa)$ to all be greater than zero – indicating that higher mean, variance, skewness and kurtosis of waiting time distributions lower the desirability for waiting. In Study 1 and the Binary and Uniform treatments of Study 2, skewness is always zero. In the Exponential treatment of Study 2, skewness is always equal to 2 for exponentially distributed waits. In Study 3, for the case of “full” information, skewness may be 0, positive or negative depending on the underlying distribution. For a given model, let Θ denote the vector of parameters to estimate.

To estimate the model parameters, we must first model the comparison between the two waiting scenarios in Option A (wait one minute and receive \$1) and Option B (the waiting scenarios of interest). In this case, the utility of Option A is simply given by $U_A = 1 - \beta$, while the utility for Option B is given by eq. B.5.2. For simplicity, denote this by U_B . Let $\Delta_{AB}(\Theta) := U_A(\Theta) - U_B(\Theta)$ be an index function. Using a distribution function $F(\Delta_{ab}(\Theta))$, this index function is linked to the observed choices. This function maps any real number to a number in the interval $[0, 1]$. The probability that the decision maker chooses Option A over Option B is given by $Pr(A, B; \Theta) = F(\Delta_{ab}(\Theta))$. Luce (1959) shows that if we choose $F(\cdot)$ as the logistic CDF where λ is the inverse standard deviation parameter, then the probability that the decision-maker i chooses Option A over Option B for waiting scenario j is equal to the binary logit such that:

$$Pr^{ij}(A, B; \Theta, \lambda) = \frac{e^{\lambda U_A(\Theta)}}{e^{\lambda U_A(\Theta)} + e^{\lambda U_B(\Theta)}}, \quad (\text{B.5.3})$$

where λ can be interpreted as a rationality parameter. When $\lambda = 0$, the decision maker does not care about utility and simply chooses at random, while as $\lambda \rightarrow \infty$, the decision maker chooses the waiting scenario that gives a higher utility.

Finally, we can then write the likelihood function as:

$$L(\Theta, \lambda) = \prod_{i=1}^N \prod_{j=1}^{100} Pr^{ij}(A, B; \Theta, \lambda)^{\mathbf{1}_{[c_{ij}=A]}} (1 - Pr^{ij}(A, B; \Theta, \lambda))^{\mathbf{1}_{[c_{ij}=B]}}, \quad (\text{B.5.4})$$

where $c_{ij} \in \{A, B\}$ is the choice of the decision maker i for waiting scenario j , and $\mathbf{1}_{[\cdot]}$ is the indicator function equal to 1 if the condition in $[\cdot]$ is satisfied, 0 otherwise.

B.5.2. Estimation Results for Studies 1-3 Our estimation results are presented in Tables B1-B4. For each study, we report the model parameters based on several specifications, starting from the mean-only model, and then progressively adding the higher moments of the waiting distribution for Option B such as variance, skewness, and kurtosis. In Study 1 we cannot separately identify the higher moments and focus on mean and variance. In Study 2 we can identify higher moments; however, because the distributions are not sufficiently different from each other, we can only identify one higher moment. In Study 3, we can identify all four moments.

First, consider the results for Study 1-2 (Tables B1-B2). In all cases, the coefficients are of the expected positive sign. Moreover, the overall fit improves substantially and significantly once we account for both variance, skewness and kurtosis. That is, the likelihood ratio test rejects the null of no improvement in model fit at $p \ll 0.01$ each time we add a further parameter.¹⁸ Although not shown, for Study 2, we estimated each of the three models where we restricted

¹⁷ It is well known that such a formulation implies preferences which are inconsistent with the Expected Utility Theorem. However, see Nakamura (2015) for an axiomatization of mean-variance-type preferences that reconciles this approach with the Expected Utility Theorem.

¹⁸ If we penalize for number of parameters by using the BIC, then for all studies, the largest model is still preferred.

the preference parameters, (β, γ) , to be the same across all three treatments.¹⁹ In all cases, we easily reject the model where preference parameters are the same across all treatments. This is consistent with our descriptive analysis where we showed that the waiting distribution does appear to matter. Furthermore, once accommodating for skewness for the Exponential treatment, we see that the coefficients on both the mean and variance of the waiting distribution drop substantially. Thus, there is a rather large, negative, base effect of the Exponential distribution. It remains to be seen whether this is because of the unboundedness of the distribution or because of its positive skewness, where long waits are possible, if unlikely. Comparing the Binary and Uniform treatments, the negative effect of the average wait is quite similar, while variance has a larger impact on behavior under the uniform distribution. We also see that the rationality parameter is somewhat lower in the Uniform (and Exponential) treatment than in the Binary treatment. This is unsurprising because it is cognitively more challenging to evaluate continuously distributed waits than it is to evaluate waits with a two-point distribution.

Consider next Study 3. Tables B3 and B4 are replications of Tables 7 and 8 in the main text, and are included for completeness and to allow comparisons with the analogous estimation in Tables B1 and B2. As noted in the main text, since the models are nested we can perform Likelihood-Ratio tests. These show that each additional parameter

¹⁹ We do not need to impose this restriction on the skewness coefficient, δ , because skewness is only operative in the Exponential treatment.

Table B1 Estimation Results For Study 1

	Mean	Mean/Var
λ (rationality)	1.396	1.403
β (mean)	0.185	0.174
γ (var.)		0.008
LogLike	-5733.9	-5710.8
Num. Obs.	630	630

Table B2 Estimation Results For Study 2

		Mean	Mean/Var	Mean/Var/Higer
Exponential	λ (rationality)	1.391	1.390	1.512
	β (mean)	0.192	0.191	0.100
	γ (var.)		0.002	0.006
	δ (higher)			0.674
Binary	λ (rationality)	1.676	1.678	1.754
	β (mean)	0.165	0.154	0.104
	γ (var.)		0.013	0.015
	δ (higher)			0.368
Uniform	λ (rationality)	1.340	1.343	1.402
	β (mean)	0.175	0.157	0.108
	γ (var.)		0.020	0.022
	δ (higher)			0.370
LogLike		-16594	-16513	-16349
Num. Obs.		1896	1896	1896

significantly improves the fit, relative to the more constrained version ($p \ll 0.01$). Note, however, that the biggest improvement in both Table B3 and B4 comes from adding kurtosis. In particular, the improvement in log-likelihood from adding kurtosis is approximately 17 times larger than the improvement from adding skewness in Table B3, and 22 times larger in Table B4. Indeed, including kurtosis appears to improve the fit more than including variance if we focus on the changes to Loglikelihoods. Finally, we estimate two-parameter models and find that even in the absence of variance, kurtosis appears to be a stronger predictor of decisions than skewness (mean + skewness model: log-likelihood = -9281.2 , mean + kurtosis model: log-likelihood = 9174.62). Taken together, these analyses suggest that kurtosis of the waiting time distribution is a key predictor of preferences in this setting.

Table B3 Estimation Results for Study 3A

	Mean	Mean/Var	Mean/Var/Skew	Mean/Var/Skew/Kur
λ (rationality)	1.421	1.431	1.442	1.518
β (mean)	0.184	0.167	0.172	0.150
γ (var.)		0.016	0.014	0.010
δ (skew.)			0.047	0.033
κ (kur.)				0.074
LogLike	-9294.7	-9265.9	-9259.7	-9155.3
Num. Obs.	1095	1095	1095	1095

Table B4 Estimation Results for Study 3B (All Three Pieces of Information Revealed)

	Mean	Mean/Var	Mean/Var/Skew	Mean/Var/Skew/Kur
λ (rationality)	1.296	1.310	1.318	1.367
β (mean)	0.193	0.180	0.182	0.173
γ (var.)		0.018	0.017	0.008
δ (skew.)			0.046	0.038
κ (kur.)				0.068
LogLike	-7194.8	-7167.1	-7163.8	-7113.5
Num. Obs.	785	785	785	785