CALIFORNIA STATE UNIVERSITY
FULLERTON
MIHAYLO COLLEGE
OF BUSINESS AND ECONOMICS

Department of Economics
Working Paper Series

2018/001

Keeping an Eye Out: Behavior in the Presence of Riskless Ambiguity

Nick Huntington-Klein, CSU Fullerton

January, 2018

CALIFORNIA STATE UNIVERSITY, FULLERTON
Mihaylo College of Business and Economics, *Department of Economics*
800 N. State College Blvd., Fullerton, CA 92834-6848 / T 657-278-2228 / F 657-278-3097

KEEPING AN EYE OUT: BEHAVIOR IN THE PRESENCE OF RISKLESS AMBIGUITY*

Nick Huntington-Klein

Department of Economics

California State University, Fullerton

nhuntington-klein@fullerton.edu

January 5, 2018

Abstract

Past research on choice under ambiguity - decisions made when the probability of each outcome is unknown - has typically focused on scenarios in which ambiguity is presented alongside risks with known probabilities. Understanding the response to ambiguity is then difficult to distinguish from the response to risk, confounding our understanding of how subjects behave in conditions of ambiguity. In this paper I take advantage of the fact that all decisions are necessarily ambiguous to some extent to develop a scenario in which ambiguity is present but is otherwise risk-free. Respondents participate in a task that is very similar to a multi-armed bandit but with no random variation in payouts. I show that respondents anticipate ambiguity

*The author was supported by the Jr/Sr Faculty Grant from the Mihaylo College of Business and Economics. I thank Kristin Kleinjans, Skylar Olsen, Vernon Smith, and session participants at the Institute for Research in Experimental Economics for their comments and suggestions, and Jeffrey Rosenblum for research assistance. The experiment described in this paper was approved by the CSU Fullerton IRB. The author declares that he has no relevant or material financial interests that relate to the research described in this paper.

in payouts without prompting, continue to do so for long periods of time, respond differently to gains and losses in this context, and that this tendency is independent of risk preference. I offer a basic, clean framework for creating a riskless ambiguous environment for subjects, and argue that these exploratory results make the case for further study of ambiguity in riskless environments.

1 Introduction

Modern treatments of choice under uncertainty offer a distinction between two kinds of uncertainty: risk, which refers to situations in which the probability that a given outcome will occur is known or can be inferred, and ambiguity, which refers to situations in which the probability of a given outcome is unknown. Ambiguity can come from many sources, including missing information, information of doubtful quality or weight (Camerer and Weber, 1992), conflicting expertise (Cabantous et al., 2011), or imprecise language (Li, 2017). Subjective expected utility theory, as laid out by Savage (1954), cannot account for these scenarios. Following Ellsberg (1961), there have been numerous approaches to expanding decision theory to account for ambiguity. These include max-min decision making as well as more flexible approaches often incorporating Choquet integrals in utility calculation (e.g. Zhang, 2002; Klibanoff et al., 2005).

In addition to theory, there is also work examining the empirical response to ambiguity, which often centers around the finding that people are averse to ambiguity. However, these approaches almost exclusively consider situations in which both risk and uncertainty are present, and known probabilities are compared to unknown probabilities. Such experiments follow from the classic Ellsberg (1961) experiment, in which a random ball will be selected from an urn containing 30 red balls and 60 balls that are either black or yellow. The selected ball determines whether the subject wins or loses. In one setting, subjects choose to win on either red or black. They tend to prefer red, which wins with a known probability of 1/3, to black, which wins with an unknown probability between 0 and 2/3. In the other setting, subjects choose to win on either "red or yellow" or "black or yellow." They tend to prefer "black or yellow," which wins with a known probability of 2/3, to "red or yellow," which wins with an unknown probability between 1/3 and 1.

While these experiments are valuable, they are limited in some ways. Since their focus is on a choice between an option with more ambiguity and an option with less, they can report preference for ambiguity but not how respondents may behave when ambiguity is unavoidable. However, some studies do examine behavior under unavoidable ambiguity. For example, there is analysis of "multi-armed bandit" problems (Bergemann and Välimäki, 2008). In the multi-armed bandit, subjects must repeatedly choose between multiple options with ambiguous but learnable random payout distributions. These reveal how people search for information in initially ambiguous contexts, and they generally do not experimentally vary the amount of ambiguity (Anderson, 2012). Studies of learning-by-doing address related behavior (Engle-Warnick and Laszlo, 2017).

However, multi-armed bandit experiments, as with Ellsberg-like experiments, are subject to a second limitation of empirical studies of ambiguity. These studies, and indeed all studies to my knowledge measuring ambiguity preference, examine behavior in the presence of ambiguity alongside risk. When analyzing actual decision-making behavior in these situations, any inference we can make about human response to ambiguity will be confounded by the presence of risk and the way that people respond to it. If the full model of how respondents evaluate risk is not known, then any given observed behavior in an ambiguity-and-risk experiment cannot be cleanly identified as being a response to ambiguity, since it may instead be a response to risk, or the interaction between risk and ambiguity.

To provide an example of how this confounding could affect results, consider a basic Ellsberg problem where a subject chooses between betting on red, which has a 1/3 probability of winning, against betting on black, which has an ambiguous probability of winning somewhere between 0 and 2/3. Disproportionately selecting red is taken as evidence of ambiguity aversion. Such a conclusion assumes that this is because the ambiguous probability of black is disliked and so undervalued. However, without additional information how the subject responds to risk, an equally valid conclusion is that the known risk of choosing red is improperly judged and for some reason overvalued. Additional information about the response to risk can clear up this problem, but until the response to risk is completely understood, the confounding issue persists to some degree.

In this paper, I construct an exceedingly simple choice task in which outcomes can easily be predicted perfectly after a short learning period. Subjects using any extrapolative method would be able to perfectly forecast future choice-contingent payoffs, and so it can be said that no risk is present. My approach takes advantage of the fact that there is ambiguity inherent in *any* decision, even one without risk. So by presenting a decision with risk stripped away and without presenting additional information, it allows the study to focus on the response to ambiguity. In this context, where there is no risk, ambiguity refers to the fact that it is possible that future payoffs will deviate from past payoff patterns without warning or any previously observed change. Subjects may anticipate these future changes for any number of reasons. They may suspect that the true payoff distribution is constant but contains rare and unforeseen events, or they may suspect the researcher is trying to trick them and will change the payoffs (Taylor and Shepperd, 1996). Either mental model is in effect treating these future payoff changes as ambiguous in nature.

The decision-theoretic problem faced by respondents is mechanically equivalent to a multi-armed bandit with no risk in payoffs. It is also metaphorically similar to the many real-life situations in which outcomes can be predicted accurately most of the time, and any deviations from the predictable status quo are ambiguous. For example, a consumer with access to two grocery stores may determine through experience that store A always has lower prices and better service than store B, and shop exclusively at A. However, having chosen to always go to store A, there is some unknowable probability that B may lower their prices or improve their service, which the customer would miss out on. The only reason the customer would have to visit B would be to check if any of these changes have occurred, even though they would be unable to estimate an a priori probability that such a change had indeed occurred.

In this riskless ambiguous scenario, any deviation from the easily determined optimal choice can be interpreted as anticipation of an unforeseen and unpredictable change in payoffs, i.e. ambiguity from the point of view of the respondent. The probability of such a change, whatever reason it is expected to occur for, is necessarily unknowable, and so this anticipatory behavior is evidence of respondents being aware of the ambiguity of payoffs. I am able to measure the extent to which subjects act in anticipation of ambiguous events. I am also able to measure whether the payoff structure affects the level of anticipation. I simulate situations with improving and declining payoffs, as well as situations in which payoffs suddenly change for the better and for the worse.

The intent of the study is not to test any particular model of choice under ambiguity or learning, but rather to provide a basic description of responses to ambiguity that are separate from responses to risk. This is intended to remove the confound between risk and ambiguity, improving the understanding of how subjects respond to both.

I find that subjects repeatedly make selections that are obviously inferior unless they anticipate, despite never having observed change, that payoffs will change in the future. The rate at which subjects make optimal choices rises as they repeatedly observe no change in payoff (or payoff trend), but stabilizes at around 80% rather than at 100%, which is inconsistent with Bayesian learning and probability matching. The rate of optimal choice-making is not statistically related to basic measures of risk aversion. Subjects respond to shocks by increasing their tendency to explore inferior options, moreso when that shock is unpleasant than when it is good. Subjects also explore inferior options more often when conditions are slowly and predictably deteriorating than when they are improving.

2 Separating Risk and Ambiguity

The distinction in economic decision-making between risk and ambiguity extends back to foundational works by Knight Knight (1921) and Keynes Keynes (1936). Here understanding risk as relating to uncertain outcomes for which probabilities are known, and ambiguity as relating to uncertain outcomes for which probabilities are unknown, there have been a number of different decision-theoretic approaches to understanding optimal choice under ambiguity, and ambiguity-sensitive preferences (e.g., Schmeidler, 1989; Chateauneuf, 1994; Epstein and Schneider, 2007). Etner et al. (2012) provides a review of some of these approaches. A long literature has emerged to use data to test the differences between these models (see Kothiyal et al., 2014, for a partial review), although doing so requires assumptions about how subjects interpret the experiment itself (Shmaya and Yariv, 2016).

This theoretical work, as well as descriptive empirical work, is often inspired by the framework set up by Ellsberg (1961). Work following Ellsberg examines how decisions are made in circumstances where a choice is made between an option with more ambiguity and an option with less (Etner et al., 2012), and takes place in experimental settings where risk and ambiguity are both present. These tasks are well-suited to valuing preference for risk over ambiguity, and in testing models of decision theory that account for both risk and ambiguity.

Alongside Ellsberg-inspired tests of ambiguity is the extensive literature on the multi-armed bandit (Bergemann and Välimäki, 2008). In a multi-armed bandit setting, subjects repeatedly choose between different "arms." Each choice results in an immediate payoff, drawn from a payoff distribution specific to that arm. Subjects do not initially know those payoff distributions (payoffs are ambiguous), and must learn about them by trying each arm (over the course of the task, payoff distributions are learned and choices become risky rather than ambiguous). There is a tradeoff between spending time exploring the payoff distributions of the different arms, and exploiting their knowledge to repeatedly select the best arm.

The multi-armed bandit frames ambiguity in a learning context, which happens to have a number of useful economic analogues. The multi-armed bandit has been used in theoretical analysis to study market pricing with unknown demand functions (Rothschild, 1974), competitive research and development (Keller et al., 2005), and job search where employees do not know their own productivity (Jovanovic, 1979), among other things.

However, in empirical application in or outside of the lab, both Ellsberg-type problems as well as the multi-armed bandit propose situations in which some risks are known and others are not. Further, they assume that ambiguity can be reduced and eliminated through learning. It may be more realistic to view at least some portion of ambiguity as consisting of events for which probability cannot be judged because they have never been observed.

This limitation matters. First of all, a framing of ambiguity which assumes that ambiguous probabilities are learnable is somewhat in conflict with the conception of ambiguity (called uncertainty) by Knight (1921), who would argue that firms are able to profit by taking on ambiguity partially because the risks cannot be learned before a decision is made.

Second, from an empirical standpoint, if risk and uncertainty are always tested jointly, then it will be difficult to identify the response to ambiguity alone. If response to risk is at some level unpredictable or follows an unknown model, then any analysis of ambiguity will be confounded. Empirical measurement of ambiguity preferences is not unaware of these problems, and some new measures allow risk preferences to cancel out of ambiguity considerations (Baillon et al., 2016), although this assumes that subjects are capable of allowing mathematically equivalent risk to be treated equally and thus cancel out.

This confounding issue is not insurmountable, and both the Ellsberg-problem and multi-armed bandit literature have provided many careful and useful results about human behavior. But the issue of interpretation does suggest that it may be useful to also examine ambiguity in a setting where it is not alongside known or learnable risks, but instead represents the potential for change in a status quo of certainty. Given that behavior under ambiguity is likely to differ depending on the source and context of that ambiguity (Cabantous et al., 2011), this leaves a wide range of behavior related to ambiguity that is essentially unstudied. There is a gap in the literature related to the study of ambiguity in a riskless setting. In this study I pursue an experiment to fill that gap

3 Experimental Method and Data

In the experiment, subjects are taken to a website with two buttons: A and B, which they choose between. They have $t \in \{1, 2, ..., 40\}$ opportunities to select one button or the other. Each time they select a button, they are rewarded with a certain amount of money, determined exactly by t and a payoff structure. Respondents are not told how much they would have earned by selecting the other button. The use of 40 button pushes is intended to give the respondent time to determine which button offers a superior payout, to capitalize on that payout by repeatedly selecting that button, and to have a chance to return to the lower-payout button to check.

There are six payoff structures. In each, there is a *Better* button, which is randomly determined to be A or B, and a *Worse* button, which always offers a payout that is a fixed amount lower than the *Better* button. The payoff structure the respondent observes is determined randomly at the beginning of the experiment. The six payoff structures are

- (Control) The *Better* button pays \$.10 on each push, and the *Worse* button pays \$.05.
- (Low Alternate Pay) The *Better* button pays \$.10 on each push, and the *Worse* button pays \$.02.
- (Slow Growth) The *Better* button pays \$.0525 in the first round, and pay increases linearly by \$.0025 each round until reaching \$.15 in the 40th round. The *Worse* button always pays \$.05 less than the *Better* button.
- (Slow Decline) The *Better* button pays \$.15 in the first round, and pay decreases linearly by \$.0025 each round until reaching \$.0525 in the 40th round. The *Worse* button always pays \$.05 less than the *Better* button.
- (Good Surprise) The *Better* button pays \$.06 for the first 20 rounds, and \$.14 for the last 20. The *Worse* button always pays \$.05 less than the *Better* button.
- (Bad Surprise) The *Better* button pays \$.14 for the first 20 rounds, and \$.06 for the last 20. The *Worse* button always pays \$.05 less than the *Better* button.

These straightforward and predictable payoff structures each allow the subjects the ability to learn easily which button is *Better*. With the exception of the Good Surprise and Bad Surprise conditions, payoffs never deviate from what can be learned and forecasted easily from only a few button pushes. In these non-Surprise conditions, there is no risk in payoffs. Payoffs do not follow a known or knowable probability distribution aside from full certainty. Once the structure is observed, the only reward-seeking reason for not fully exploiting the *Better* button is the possibility that the payoffs are ambiguous.

The purpose of the Low Alternate Pay condition is to test for the influence of search costs on the propensity to check the *Worse* button. With a lower payout on the *Worse* button, the relative cost of checking is higher, and we might expect respondents to be less willing to see whether its payoff has changed.

The Slow Growth and Slow Decline conditions allow for a test of whether changing conditions affect the tendency to anticipate ambiguity. The presence of changes, even predictable changes, in payouts may alter the anticipation of ambiguous payouts. Differences between behavior under Slow Growth and Slow Decline also allow the results to highlight the differing influence of loss and gain on behavior. Analogously, one can consider whether someone is more or less likely to research alternative investment options when their own stock portfolio is consistently increasing in value as opposed to when it is consistently dropping.

The Good Surprise and Bad Surprise conditions are not strictly riskless since payouts are shown to be at least somewhat unpredictable after round 20, and so must be thought of differently. However, these experimental conditions allow me to observe respondent behavior in response to a shock, and whether that shock is good or bad.

During the experiment, respondents have access to information about the total amount of money they have earned up to that point, the number of button pushes they have remaining, how much money they earned on each previous button push, and which button earned them the money. Following the experiment, I ask demographic questions and elicit risk aversion. Demographic questions include gender, age group, and education level.

Following Falk et al. (2016), risk tolerance is elicited in two ways. First, respondents are given a "staircase" task giving them hypothetical choices between a coin flip that results in a $300 payout with 50% chance and $0 with 50% chance against a safe payout at different levels from $0 to $300 in intervals of $15. Second, respondents are asked "How do you see yourself: are you a person who is generally willing to take risks, or do you try to avoid taking risks?" and told to rank themselves from 0 (avoids risks) to 10 (takes risks). Risk preferences are non-incentivized and are elicited after respondents have completed the button-pushing task. These measures should be interpreted keeping in mind that the subjects may be altering their responses to be consistent with the actual choices they have just made.

The experiment was coded on a personal website, and 605 respondents from predominantly English-speaking countries were recruited on the crowdsourcing site microWorkers.[1,2] As a result, the data includes demographic and risk aversion information on 605 respondents, and information on 24,200 button pushes made by those respondents.

Demographic information is shown in Table 1, and risk profile information is in Figures 1 and 2. The sample is fairly gender-balanced, although it skews young relative to the general population. Educational attainment rates are similar to the general population.

The sample sees themselves as moderately willing to take risks, with most responses gathered in the 5-7 range. 12.6% of the sample was unable to accurately complete the "staircase" risk preference task. These respondents gave inconsistent responses such as preferring the sure money if offered $75 but preferring the coin flip if offered a sure payment of $90. This could represent the relative difficulty of the task, or inattention.[3] Omitting inconsistent responses, respondents are again middlingly risk-averse. Nearly 85% of the sample accepts the sure money at values lower than the expected value of the coin flip ($150). About 14% are very risk averse, shunning the

[1] A sample size of 600, or 100 per treatment, was determined using a power calculation seeking to provide more than 80% power to test the difference in the number of *Worse* button presses per respondent. The allocated budget allowed for the additional five respondents to get a total of 605.

[2] microWorkers can be thought of as similar to its competitor Amazon Mechanical Turk, which has been shown to be high-quality and usable for behavioral studies (Paolacci et al., 2010; Buhrmester et al., 2011; Goodman et al., 2013). microWorkers provides data of similar quality to Mechanical Turk (Peer et al., 2015). Due to limitations in allowable payout structures on microWorkers, respondents were guaranteed a minimum payout of $1.40, even if they earned less in the task. Only 15/605 respondents were at or below this threshold.

[3] Results throughout the paper are similar if these inconsistent responders are omitted from all analyses.

Table 1: Sample Characteristics

Variable	Mean	Variable	Mean
Gender:		Education	
Female	.471	No HS Degree	.033
Male	.521	HS Degree/GED	.154
Other/No Response	.008	Some College	.281
Age:		One-year Cert.	.042
18-25	.318	Associate's Degree	.100
26-35	.364	Bachelor's Degree	.268
36-45	.198	Postgraduate	.122
46-55	.083	Time Spent (secs)	205.934
56+	.037	Total Earned	$3.370
Completed Risk Pref.	.874	N	605

Figure 1: Self-Stated Risk Preference

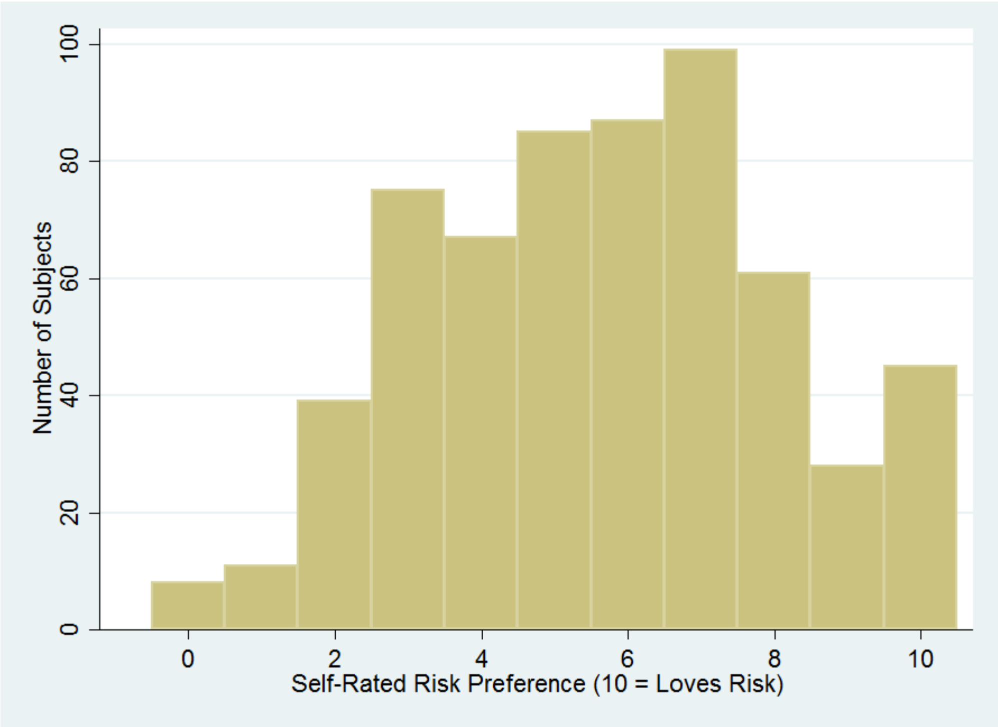

Figure 2: Risk Preference Elicited Comparing Sure Payment to \$0/\$300 Coin Flip

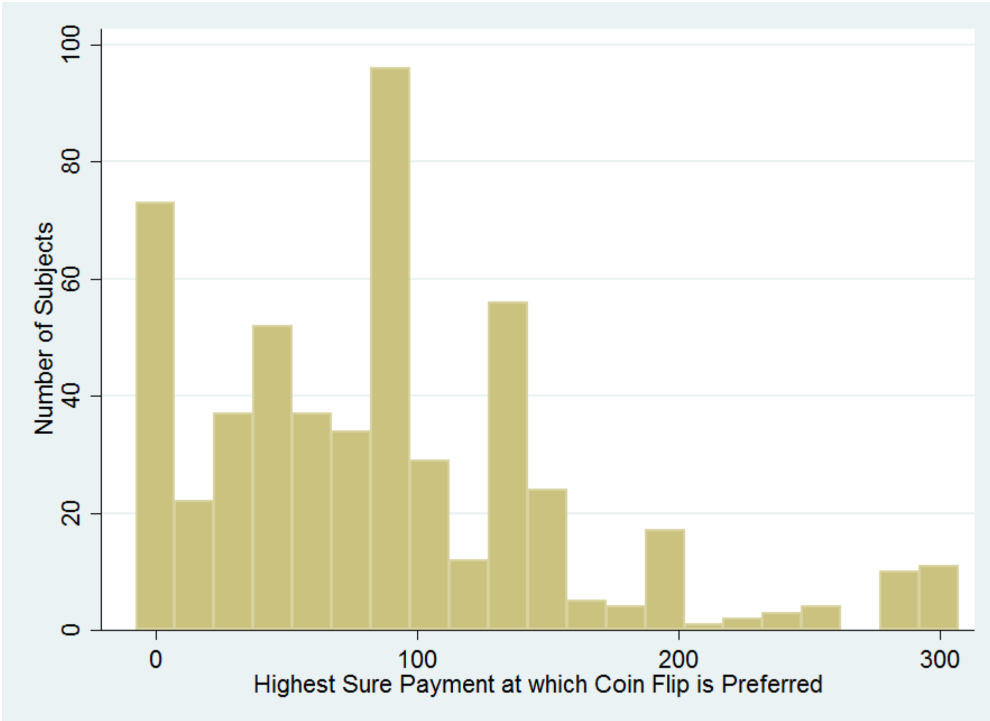

coin flip if even \$15 is offered for sure. The two measures of risk preference have a statistically significant correlation of .406.

These demographic data suggest that the sample is, if not perfectly representative of a wider population, at least relatively well-matched on the small number of variables recorded. To the extent that ambiguity-anticipating behavior varies with demographics, we might expect that average behavior here is not too far off from that in the broader population, supporting external validity.

4 Results

The basic questions under examination in this study are: (1) do subjects select the *Worse* button after the payoff structure should be known?; (2) how does the rate of *Worse*-button pushing change with learning?; and (3) how does the structure of the payout affect the rate of *Worse*-button pushing?

The answer to the first question is an emphatic yes. To look at this, I omit respondents in the "Surprise" experimental conditions, since those payout structures are not perfectly riskless. Allowing the first ten button pushes as a training period, only 52 respondents out of 394 remaining respondents pushed the same button for the 11th through 40th button pushes, indicating that most respondents opted to test an inferior option again to see if, with no indication or evidence, it had changed to a superior option.

Omitting the 8 respondents who pressed the same button 40 times (including the training period) and so never saw the payouts of both buttons, on average respondents checked the *Worse* button 8.62 times in the last 30 button pushes. The anticipation of ambiguously-determined changes in payouts was not only present but fairly common. Figure 3 shows the distribution of *Worse*-button presses. 13.2% always pressed the *Better* button, which is the ex-post optimal approach, and what would be predicted by someone choosing rationally under the assumption of no ambiguity in payouts. 59.9% of the sample checked the *Worse* button more than three times. The median number of times checking the *Worse* button is 6.

The second question, about how learning affects behavior in this context, is a little harder to address. One approach is to look at the average rate at which subjects choose the *Better* button over time. Figure 5 shows the proportion of button presses that select the *Better* button over time, both for all riskless experimental conditions and for the Normal condition separately. Both tell a similar story. There is an initial period, over the first 13 pushes or so, when learning primarily occurs.

Then, there follows a period in which the rate of correct pushes stalls around 70%. This continues for about 20 pushes, roughly half the length of the entire exercise. There is no statistically detectable rise in the rate over this period. Finally, near the end of the available pushes, the rate jumps to about 75% and stays there to the end of the experiment.

What we see from this pattern is that, first, there is clearly a learning period. This is to be expected, given that respondents need to time to determine which button is *Better* and to recognize that they never observe the *Worse* button being superior. Second, we can note that the choice to

Figure 3: Number of Presses of *Worse* Button After First Five Presses

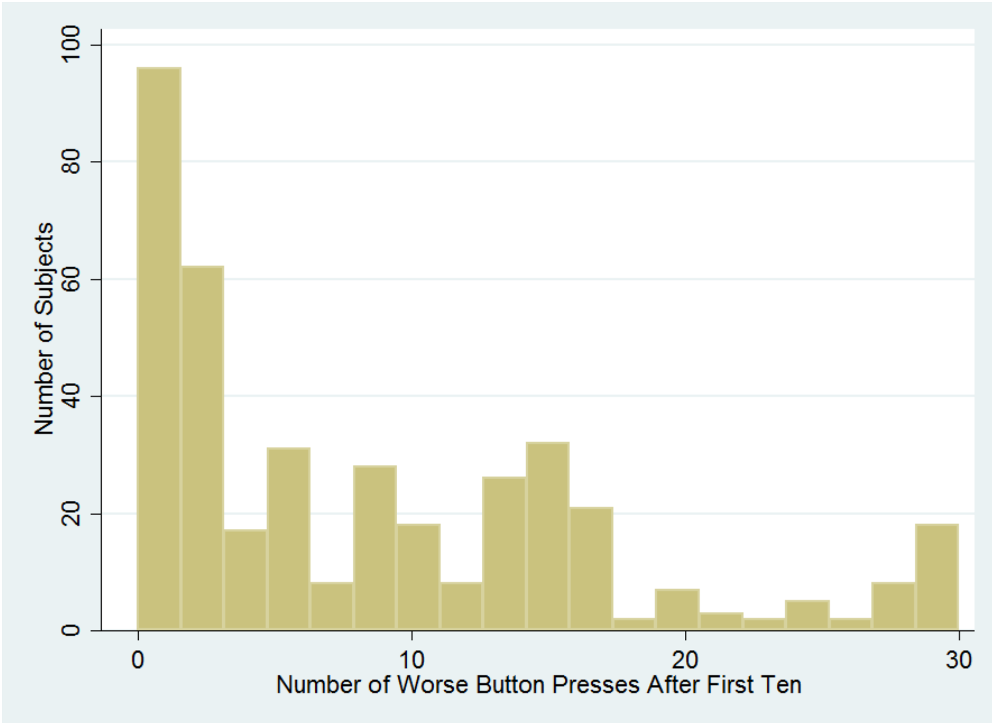

Figure 4: Rate of Pushing the *Better* Button Over Time

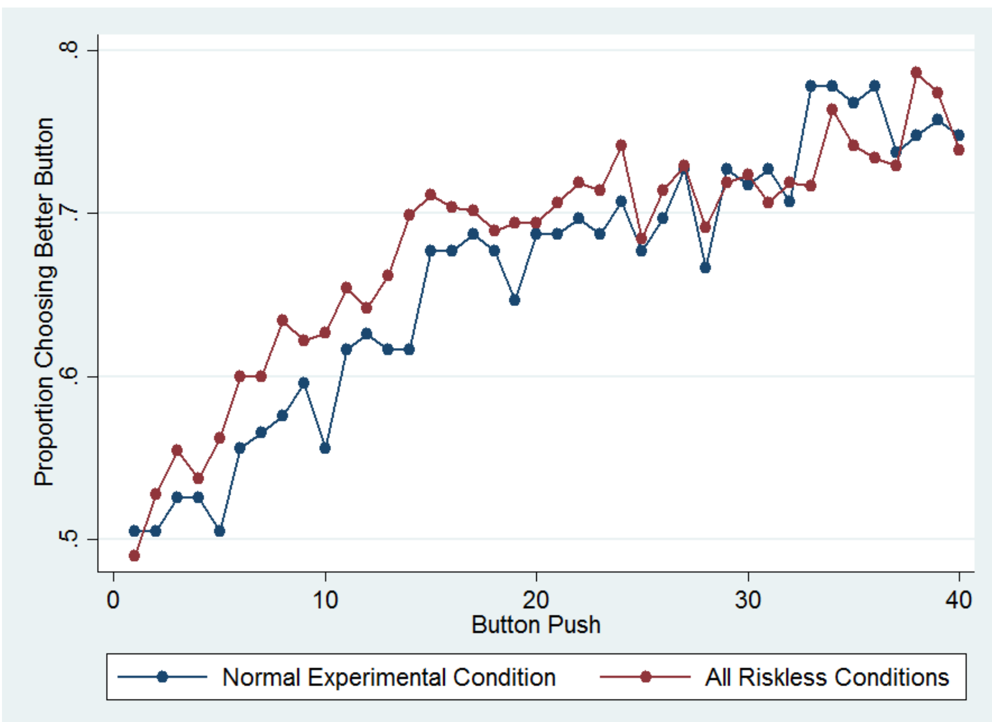

push the *Worse* button is not purely a matter of learning or Bayesian updating about risk. If learning about risk were the only explanation, and ambiguity could be eliminated by learning, we would expect the rate of correct choices to rise past 70% as more information continues to show that the *Better* button remains the *Better* button, and we would expect the rate of *Better* button pushes to continually rise to near 100% rather than stalling.

The observed behavior is inconsistent with Bayesian learning that concerns only risks, and also is inconsistent with the irrational behavior of probability matching. Probability matching is a phenomenon observed in some multi-armed bandit-like settings in which each option is chosen at a rate equal to the probability that it is the optimal choice (Vulkan, 2000). Since there is never any indication that the *Worse* button is optimal, probability matching would suggest that the bad button is never pushed after a learning period, since it has zero probability of being the best.

Generally, we have a period of learning followed by a period after learning. But in this period after learning, suboptimal decisions are still made at a relatively high rate, in a way that is not explained by probability matching. Subjects act in a way consistent with a persistent belief in the ambiguity of payoffs. They check the bad button regularly, which makes sense only if they suspect that the payoff structure will change.

Instead, there appears to be a period of exploration (checking both buttons) followed by a period of partial exploitation (selecting the *Better* button), with sporadic exploration (checking the *Worse* button). Exploration and exploitation comes in strings, where subjects either explore or exploit multiple times in a row. The strings of exploitation are longer, as shown in Figure 5. Measuring the total number of same-button pushes in a row after the first ten pushes, 63.48% of *Worse*-button pushes are in strings of one, a single button push before returning to the *Better* button. But only 29.11% of *Better*-button pushes are in strings of one. Still, exploitation does not last excessively long. The median string length for *Better*-button pushes is only 3.

As time goes on, even though the average rate at which the *Better* button is pushed does not rise much, the length of strings does increase. Regressing string length linearly on the push number for

Figure 5: Length of Strings of Same-Button Pushes

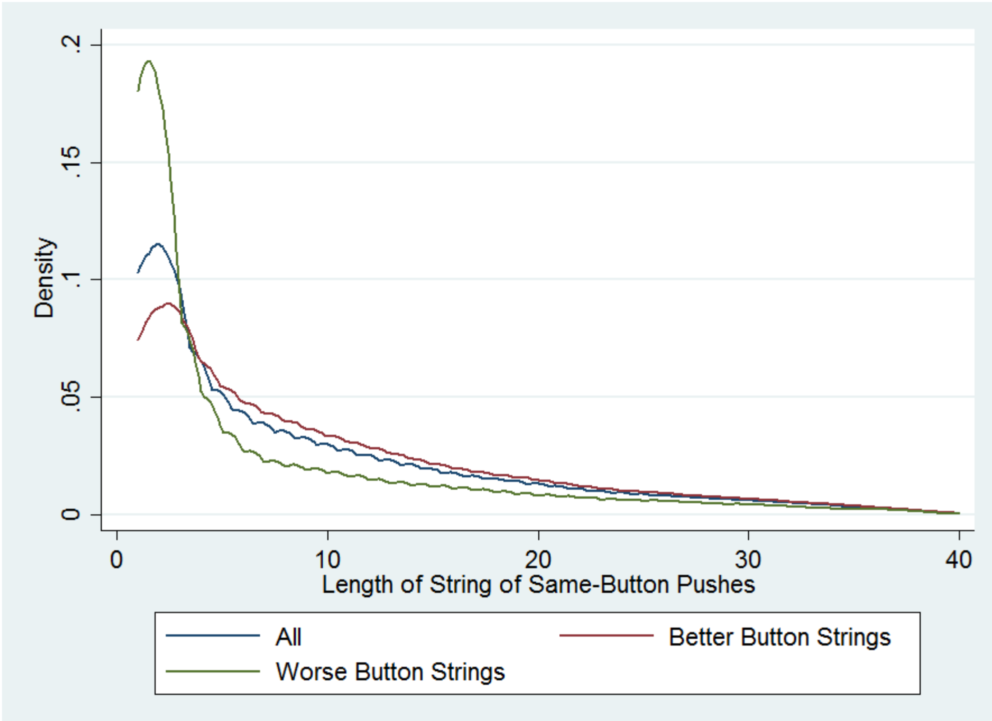

Better-button strings after the first ten pushes shows that a string that ends one button-push later is on average .289 pushes longer. However, this apparent increase in confidence is accompanied by more thorough exploration - the same regression on *Worse*-button pushes shows that a string of *Worse*-button pushes that ends one button-push later is on average .088 pushes longer.

As we would expect, learning impacts subject behavior, and there is some indication that the anticipation of ambiguity dies over time, as the status quo is enforced more strongly. But at least within the length of time given to subjects in this study, *Worse*-button checking never dies out. Section 4.3 discusses the results of a follow-up experiment that lasts for 80 pushes rather than 40, and *Worse*-button checking does not die out in that setting either. Whatever expectations lead the subject to anticipate a change in the status quo should persist in the absence of that change actually having occurred yet. The length of the reigning status quo gives little indication as to when it will end.

This section shows that, overall, subjects continue to check the *Worse* button throughout the task, even though there is no knowable probability that suggests this is an optimal choice. In the next section, I examine whether this tendency differs based on the exact structure of the payouts.

4.1 Differences Between Payout Structures

There are significant differences in behavior between the different payout structures. Compared to Normal, there are .409 more *Better* presses in the Slow Growth condition and 2.244 more *Better* presses in the Slow Decline condition. The 2.244 difference is significant at the 5% level, and the Slow Growth and Slow Decline levels are different at the 10% level. Each of these cross-type comparisons has roughly 100 observations in each type.

The Surprise conditions see more *Better* presses overall than Normal, 3.198 for Good Surprise and 3.151 for Bad Surprise, both significantly different from Normal at the 1% level.

There is no significant difference in the number of *Worse* presses between the Normal and Low Alternate Pay conditions.

Figure 6 shows the rate of selecting the *Better* button across experimental conditions. The figure demonstrates that the Surprise difference appears to be due to unexpectedly quicker learning rates for the Surprise conditions, which may be random given that the *Better*-pushing rate is higher in both conditions from the outset.

Figure 6 allows for a more detailed comparison of the different experimental conditions. It first allows us to look at the effect of the Surprise - good or bad - on exploration. For both the Good Surprise and the Bad Surprise conditions, the rate of *Better* pushes drops significantly after the surprise at the 20th push. However, the drop is much larger for the Bad Surprise than the Good Surprise, and comes more quickly.

In both cases, evidence of change inspires exploration of the *Worse* button, but there is asymmetry in the responses. Consistent with loss aversion, there is a bigger and more immediate response to the Bad Surprise. Breaking the response down to a finer grain and running an interrupted time series design with a break after the 21st push and quadratic button-push trends, there is an insignif-

Figure 6: Rate of Pushing the *Better* Button for All Conditions (Five-Push Average)

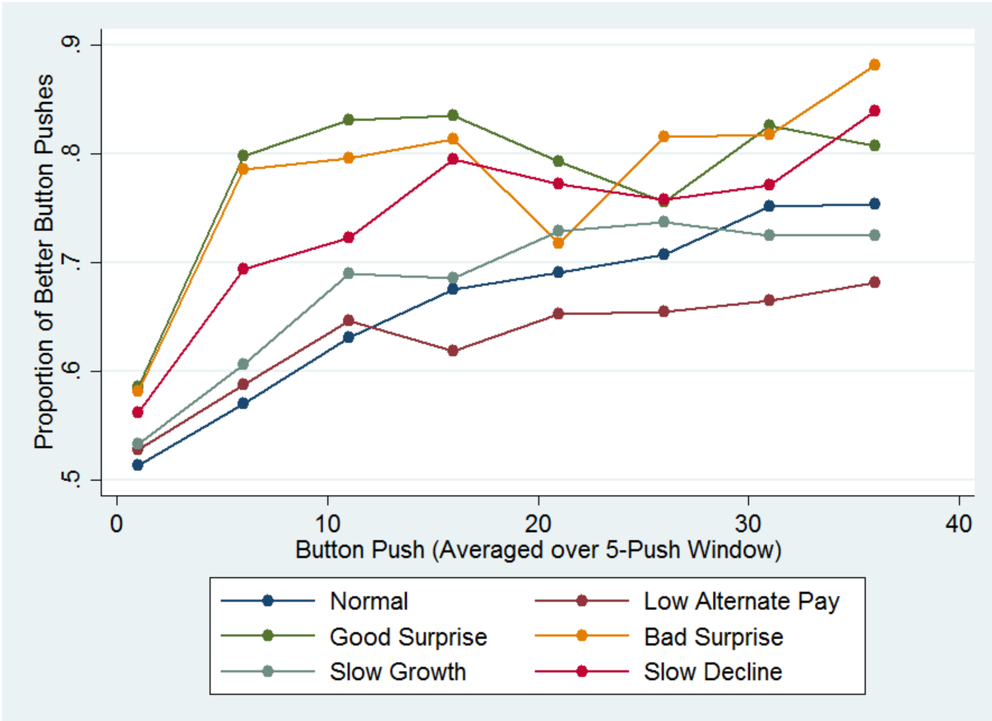

icant decrease of .006 in the rate of selecting the *Better* button for the Good Surprise condition, and a drop of .152, significant at the 1% level, in the Bad Surprise condition.

There is also asymmetry between the Slow Growth and Slow Decline conditions. At all points, there is more exploration in the Slow Growth condition than in the Slow Decline condition. Subjects are less willing to accept the cost of checking the *Worse* button when payoffs are declining.

The observation that behavior varies in response to payoff structure and the injection of surprise also supports the notion that checking the *Worse* button follows some intention and is not just random action or boredom. It makes intuitive sense that subjects would prepare and explore at different rates under improving, declining, or constant conditions.

Table 2: Relationship Between Risk Preference and Number of Wrong Button Presses

Variable	(1)	(2)
Self-stated Risk Preference	.274	
	(.176)	
Highest Certain Payment Preferred		-.006
to Coin Flip		(.006)
Constant	7.049	8.764
	(1.081)	(.744)
N	402	351

4.2 Exploratory Behavior and Risk Preferences

In this experiment, there is no risk present in payoffs. In calculating the optimal series of decisions, a researcher would not consult a measure of risk preference. However, it is worthwhile to see if risk preference is related to subject behavior in the task. A strong relationship between risk preference and task behavior could be interpreted as subjects perceiving the task to actually be risky, or that risk preference and behavior under ambiguity are correlated within subjects. A lack of a relationship between risk preference and task behavior would suggest that response to risk and ambiguity are indeed two separable processes, and further that studies finding positive correlation between risk preference and ambiguity preference (as in **?**) may be affected by the fact that ambiguity preference is often measured using a risk-and-ambiguity task .

To examine this, in Table 2 I regress the number of times the Wrong button is pressed in the riskless scenarios after the 10-press learning period on the self-stated risk preference as well as the highest certain payment for which the $0/$300 coin flip is preferred to the certain payment.[4] In both cases, a higher number indicates a higher preference for risk.

The point estimate on the self-stated risk preference is positive, indicating a difference of 2.5 button pushes comparing the least and most risk-averse. However, the relationship between risk preference and exploratory behavior is statistically insignificant, and the alternate measure of risk preference suggests no relationship whatsoever.

[4]This analysis is only performed for the portion of the sample that consistently completed the task.

These results should be taken with some skepticism, given that the risk preference measures are not incentivized, and that this is a null result estimated with a modest sample size. However, the lack of a relationship here is suggestive that risk preference and behavior under ambiguity can be thought of as two separate things. Further, given that this result is not consistent with prior work finding a relationship between risk aversion and behavior under ambiguity, it may be worthwhile to develop measures of ambiguity aversion that do not rely on a risk-and-ambiguity scenario to test the robustness of that relationship.

4.3 Extending the Length of the Experiment

The experiment described above found that subjects pressed the *Better* button at an increasing rate as they learned about their options, but that the rate of pressing the *Better* button tended to level out around 70% and largely stays there until the end of the experiment, increasing by the end to about 75%. A natural follow-up question is whether this occurs only because the time frame of the experiment was too short for subjects to reach 100% or get close.

In this section I describe a small follow-up experiment in which subjects were given 80 button pushes rather than 40. All subjects were in the Control payout setting, where the *Better* button paid $.10 per push, and the *Worse* button paid $.05. The sample size, 39 respondents, was smaller than in the original experiment. The follow-up experiment is otherwise exactly the same as in the original. After dropping the two respondents who chose the same button 80 times in a row and thus did not see the alternate payoff, the rate at which respondents pushed the *Better* button is shown in Figure 7.

The rate of pushing the *Better* button in this follow-up is similar to the original experiment over the first 40 pushes, with a quick learning curve that flattens near 70% rather than 100%. However, the rate of pushing the *Better* button increases after the 40-push mark. The rate does not converge to near 100%, though, and instead stabilizes near 80%. Like in the original experiment, there is no statistically significant relationship between risk preference and performance in the task.

Figure 7: Rate of Pushing the *Better* Button in Extended-Length Experiment

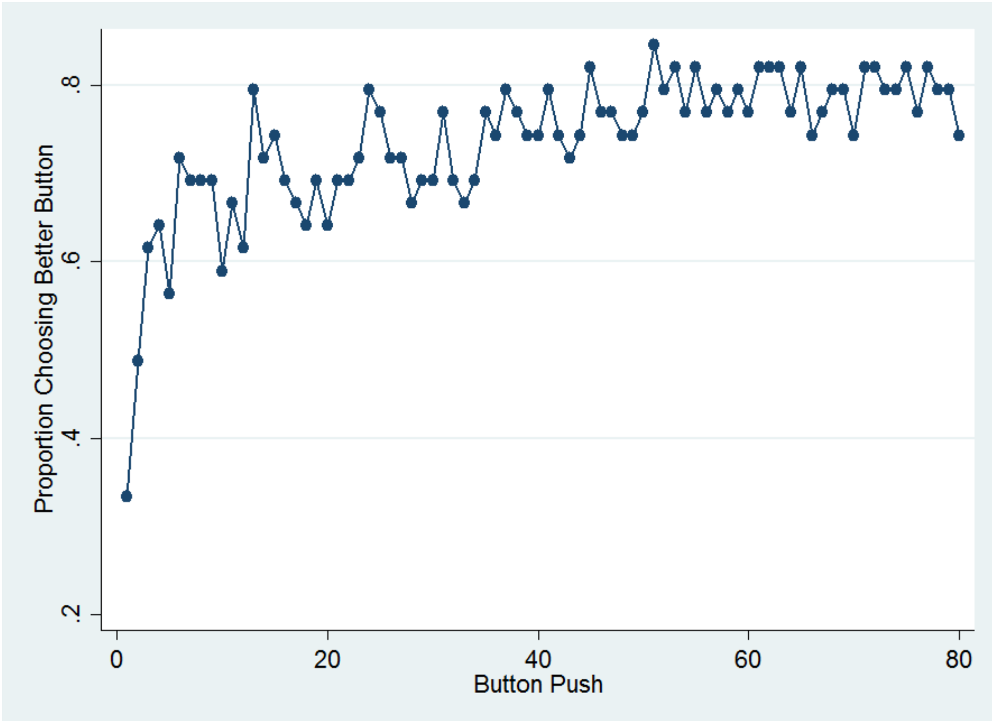

5 Conclusion

In this paper, I presented the results of an experiment in which subjects were told to repeatedly choose between two buttons, A and B, one of which is *Better* and one of which is *Worse*. The difference between the two is fixed through the entire experiment, and each payout stream is easy to predict. This is a task in which optimal choice could be easily predicted by any basic statistical method if there is assumed to be no ambiguity. There is no risk, i.e. no knowable or observable non-certain probability distribution, for any of the payouts in most conditions.

Bayesian learning in the absence of ambiguity would suggest continual convergence towards the *Better* button, eventually reaching something at or near 100%. However, in this experiment, subjects continue to choose the *Worse* button at fairly high rates through to the end of the experiment. Even after a training period, the rate at which the *Better* button is chosen stabilizes around 70%, rather than anywhere near 100%. Doubling the length of the experiment saw the rate rise to only 80%.

This study is intended only to be exploratory, and I do not make any claims to be able to favor a given model of choice under ambiguity. Doing so would require additional strong assumptions about how subjects interpret the task (Shmaya and Yariv, 2016). The behavioral-model implications of these results may be different, for instance, if the source of ambiguity is subjects assuming that the payoff structure will change over time, subjects assuming that the payoff structure is constant but contains rare events, subjects assuming that the researcher may be trying to deceive them, or subjects assuming that the researcher *must* have designed some reason to push the *Worse* button.

Regardless of the source of ambiguity, we do see subjects checking the *Worse* button at constant rates after long periods of no deviation in payoff structure. Regardless of the particular model underlying this behavior, for *Worse*-button selection to be consistent with any intentional attempt to earn a higher payoff, subjects must be anticipating that future payoffs will be different from current payoffs in some way that is not based on extrapolation from previous observations of payoffs. This can be understood as an anticipation of ambiguity in payoffs.

The prevalence of this ambiguity-anticipating behavior in this study has several broader implications. The first is that subjects do not need to be prompted with risk to anticipate ambiguity. This is perhaps not surprising given that, in normal settings, people are generally aware that there are things they don't know, and that things can sometimes change in unpredictable ways, even if there would be no basis on which to estimate a probability of that change. Choice tasks, whether in the laboratory or in the real world, have naturally ambiguous payouts that respondents may be aware of, whether or not that ambiguity is intentional.

Second, response to ambiguity is a concept that exists coherently without risk. When ambiguity and risk are studied together without a fully known model of the response to risk, inference about the response to ambiguity may be confounded by the presence of risk in the task. Although risk preference and ambiguity preference as commonly measured are related, this paper finds no significant link between risk preference and behavior in the choice task. Current measures of ambiguity aversion, performed with tasks that compare known probabilities to unknown probabilities, may be related to risk preference largely because they also depend on subjects evaluating risk. There is room for future work in developing measures of ambiguity preference, and models of behavior under ambiguity, that do not rely on risk. A full model of behavior under uncertainty will benefit from understanding response to risk, response to ambiguity, and the response to the interaction of the two. Currently, response to ambiguity alone is ignored. Baillon et al. (2016) is a step in the right direction in this regard, although those insights may be pushed further towards a measure of ambiguity preference that does not rely at all on risk assessment.

The goal of this paper is to emphasize that, even though ambiguity is studied almost exclusively alongside risk, it is possible and useful to study it alone, and subjects continue to respond to ambiguity in the absence of risk. Further, our current understanding of behavior under risk, and under ambiguity, is confounded. An Ellsberg-type task that compares known probabilities to unknown probabilities can only be used to make inference about response to ambiguity alone if it is already known how the subject will respond to risk. Given that ambiguity appears in many situations that are not risky - formally, every situation in which there is less than full information has some ambiguity - this is not a trivial distinction. There is much more to be explored here.

References

Anderson, Christopher M. 2012. Ambiguity Aversion in Multi-Armed Bandit Problems. *Theory and Decision* 72(1):15–33.

Baillon, Aurélien, Zhenxing Huang, Asli Selim, and Peter P Wakker. 2016. Measuring Ambiguity Attitudes for All (Natural) Events.

Bergemann, Dirk, and Juuso Välimäki. 2008. Bandit Problems. In *The new palgrave dictionary of economics*, ed. Steven Durlauf and Larry Blume. Cowles Foundation Discussion Paper, London, England: Macmillan Press.

Buhrmester, Michael, Tracy Kwang, and Samuel D. Gosling. 2011. Amazon's Mechanical Turk: A New Source of Inexpensive, yet High-quality, Data? *Perspectives on Psychological Science* 6(1).

Cabantous, Laure, Denis Hilton, Howard Kunreuther, and Erwann Michel-Kerjan. 2011. Is Imprecise Knowledge Better than Conflicting Expertise? Evidence from Insurers' Decisions in the United States. *Journal of Risk and Uncertainty* 42(3):211–232.

Camerer, Colin F., and Martin Weber. 1992. Recent Developments in Modeling Preferences: Undertainty and Ambiguity. *Journal of Risk and Uncertainty* 5:325–370.

Chateauneuf, Alain. 1994. Modeling Attitudes Towards Uncertainty and Risk Throguh the Use of Choquet Integral. *Annals of Operations Research* 52:3–20.

Eckel, Catherine C., and Philip J. Grossman. 2008. Men, Women, and Risk Aversion: Experimental Evidence. In *Handbook of experimental economics results*, ed. Charles R. Plott and Vernon L. Smith, 1st ed., chap. 113, 1061–1073. Elsevier.

Ellsberg, Daniel. 1961. Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics* 75(4):643–669.

Engle-Warnick, Jim, and Sonia Laszlo. 2017. Learning-by-doing in an Ambiguous Environment. *Journal of Risk and Uncertainty* 55:71–94.

Epstein, Larry G., and Martin Schneider. 2007. Learning under ambiguity. *Review of Economic Studies* 74(4):1275–1303.

Etner, Johanna, Meglena Jeleva, and Jean-Marc Tallon. 2012. Decision Theory Under Ambiguity. *Journal of Economic Surveys* 26(2):234–270.

Falk, Armin, Anke Becker, Thomas Dohmen, David Huffman, and Uwe Sunde. 2016. The Preference Survey Module : A Validated Instrument for Measuring Risk, Time, and Social Preferences.

Goodman, Joseph K., Cynthia E. Cryder, and Amar Cheema. 2013. Data Collection in a Flat World: The Strengths and Weaknesses of Mechanical Turk Samples. *Journal of Behavioral Decision Making* 26(3):213–224.

Jovanovic, Boyan. 1979. Job Matching and the Theory of Turnover. *Journal of Political Economy* 87(5):972–990.

Keller, Godfrey, Sven Rady, and Martin Cripps. 2005. Strategic Experimentation with Exponential Bandits. *Econometrica* 73(1):39–68.

Keynes, John M. 1936. *The General Theory of Employment, Interest and Money*. London: Macmillan.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. A Smooth Model of Decision Making under Ambiguity. *Econometirca* 73(6):1849–1892.

Knight, Frank H. 1921. *Risk, Uncertainty and Profit*. Boston: Hart Schaffner & Marx.

Kothiyal, Amit, Vitalie Spinu, and Peter P. Wakker. 2014. An Experimental test of Prospect Theory for Predicting Choice under Ambiguity. *Journal of Risk and Uncertainty* 48:1–17.

Li, Chen. 2017. Are the Poor Worse at Dealing with Ambiguity? *Journal of Risk and Uncertainty* 54(3):239–268.

Paolacci, Gabriele, Jesse Chandler, and Panagiotis Ipeirotis. 2010. Running Experiments on Amazon Mechanical Turk. *Judgment and Decision Making* 5(5):411–419.

Peer, Eyal, Sonam Samat, Laura Brandimarte, and Alessandro Acquisti. 2015. Beyond the Turk: An Empirical Comparison of Alternative Platforms for Crowdsourcing Online Behavioral Research.

Rothschild, Michael. 1974. A Two-Armed Bandit Theory of Market Pricing. *Journal of Economic Theory* 9:185–202.

Savage, Leonard J. 1954. *The Foundation of Statistics*. New York: Wiley.

Schmeidler, David. 1989. Subjective Probability and Expected Utility without Additivity. *Econometrica* 57(3):571.

Shmaya, Eran, and Leeat Yariv. 2016. Experiments on Decisions Under Uncertainty: A Theoretical Framework. *American Economic Review* 106(7):1775–1801.

Taylor, Kevin M., and James A. Shepperd. 1996. Probing Suspicion Among Participants in Deception Research. *American Psychologist* 51(8):886–887.

Vulkan, Nir. 2000. An Economist's Perspective on Probability Matching. *Journal of Economic Surveys* 14(7):101–118.

Zhang, Jiankang. 2002. Subjective Ambiguity, Expected Utility, and Choquet Expected Utility. *Economic Theroy* 20(1):159–191.

Appendix A Experimental Instrument

Respondents in the sample pool are shown the opportunity to take the survey. Those who opt in are shown a standard disclaimer provided by the institutional review board before beginning. Following this, they are given the instructions:

> On the next page, you will be shown two buttons, A and B.
>
> You may choose to press either A or B. Each button press will earn you money.
>
> You will be given 40 chances to choose between A and B. Each time, you may press whatever button you like. You do not have to choose the same button every time.
>
> The total amount of money you've earned so far will be shown on the screen, as well as the amount you have earned for each button press.
>
> When you are done, any money you've earned beyond the guaranteed $1.40 for completing the task will be given to you as a microWorkers bonus. That $1.40 will not be taken away from you, even if you fail to earn $1.40 here.
>
> DO NOT press the Back or Refresh buttons on your browser, or you may be denied your bonus.
>
> That's it! Go ahead and earn some money.

Following these instructions, they are taken to the experiment page, which includes:

- The phrase "Please choose between A and B. Each button press will earn you money. Scroll down to see full instructions again." and, at the bottom of the page, the instructions reprinted.
- Information on the amount of money earned so far and the number of remaining button pushes.
- The A and B buttons

- Under the A button, a list of each of the times the A button had been pushed, and the money earned each time, and similarly for the B button.

Pushing the A or B button on this page returns them to the page until the button pushes became exhausted, at which point the respondent is taken to a page with demographic and risk aversion questions, as described in the main text.

There are two risk aversion questions asked. One is “How do you see yourself: are you a person who is generally willing to take risks, or do you try to avoid taking risks?” with ratings from 0 to 10 offered.

The other risk aversion question is “Suppose you’ve been given a series of choices. These choices are hypothetical and will not affect your bonus. You may choose to either **accept** a fixed amount of money for certain, or you may choose to **flip a coin**. If you flip the coin, then you win $300 if the coin comes up heads, and $0 if the coin comes up tails.” followed by a series of radio button pairs reading “Given a choice, would you: ○ accept $X, or ○ flip the coin?” with $X starting at $0 and incrementing by $15 until it reaches $300.