Preference Discovery

Jason Delaney*[1], Sarah Jacobson†[2], and Thorsten Moenig‡[3]

[1]*School of Business, Georgia Gwinnett College, 1000 University Center Lane, Lawrenceville, GA 30043, USA*
[2]*Department of Economics, Williams College, 24 Hopkins Hall Dr., Williamstown, MA 01267, USA*
[3]*Temple University, Fox School of Business, Alter Hall 611, 1801 Liacouras Walk, Philadelphia, PA 19122, USA*

July 2019

Abstract

Is the assumption that people automatically know their own preferences innocuous? We present an experiment studying the limits of preference discovery. If tastes must be learned through experience, preferences for some goods may never be learned because it is costly to try new things, and thus non-learned preferences may cause welfare loss. We conduct an online experiment in which finite-lived participants have an induced utility function over fictitious goods about whose marginal utilities they have initial guesses. Subjects learn most, but not all, of their preferences eventually. Choice reversals occur, but primarily in early rounds. Subjects slow their sampling of new goods over time, supporting our conjecture that incomplete learning can persist. Incomplete learning is more common for goods that are rare, have low initial value guesses, or appear in choice sets alongside goods

*`jdelaney@ggc.edu`
†`sarah.a.jacobson@williams.edu`; Corresponding author
‡`moenig@temple.edu`

that appear attractive. It is also more common for people with lower incomes or shorter lifetimes. More noise in initial value guesses has opposite effects for low-value and high-value goods because it affects the perceived likelihood that the good is worth trying. Over time, subjects develop a pessimistic bias in beliefs about goods' values, since optimistic errors are more likely to be corrected. Overall, our results show that if people need to learn their preferences through consumption experience, that learning process will cause choice reversals, and even when a person has completed sampling the goods she is willing to try, she may continue to lose welfare because of suboptimal choices that arise from non-learned preferences.

Keywords: discovered preferences, preference stability, learning
JEL codes: D81, D83, D01, D03

1 Introduction

> "You do not like them. So you say. Try them! Try them! And you may."
>
> *Green Eggs and Ham*
> Dr. Seuss

Do you know what you like? Neoclassical microeconomic choice theory is grounded in the assumption that people make choices according to a stable ranking that represents their true preferences. However, as proposed in Plott (1996), it is possible that people don't know their tastes until they discover them through consumption experience. When preferences are not fully discovered, people may make choices that don't maximize utility. If this is the case, some standard results of neoclassical microeconomic theory come into question. In this paper, we start from an assumption that preferences must be learned from experience, and, using an experiment simulating preference learning, we explore implications for choice patterns and welfare, focusing on the extensive margin: what is and is not learned.

Consider an encounter with a new food. For example, one of this paper's authors had not eaten celeriac until a few years ago. She *a priori* believed it untasty. When she tried it, she discovered that she likes it. Experience yielded a more accurate assessment of her preferences, and she now enjoys a more efficient level of celeriac consumption. Still, because of her initial misperception, she might have missed out on a lifetime of celeriac appreciation had she not been induced to try it—indeed, she is likely missing out on other delicious vegetables due to mistaken beliefs and a lack of experience. In this study, we show that the need to learn preferences through experience can generally cause persistent welfare loss.

The idea of tastes that are not fully known to the decision-maker has received a small amount of attention in economics but much more in psychology, so our work is informed by past studies from both fields. Preference discovery has been little studied in either field because psychological models often do not feature stable underlying preferences (Ariely et al, 2003; Lichtenstein and Slovic, 2006), while models in economics typically implicitly assume stable preferences that are known to the decision-maker.[1]

[1]We distinguish between learning about objective circumstances and learning about

Preferences might not need to be discovered through experience if people can simply predict what they will like. As discussed by Kahneman et al (1997), Scitovsky (1976) argued that people are bad at predicting their utility from a prospective choice. Becker (1996) argued the opposite, and indeed, Kahneman and Snell (1990) note that, when experiences are familiar and immediate, people seem fairly good at predicting utility. Many results from psychology and economics support Scitovsky's claim, however. Loewenstein and Adler (1995) find people fail to predict changes in their own tastes, and Wilson and Gilbert (2005) review extensive evidence showing systematic errors in forecasting happiness.

A few papers have studied preference discovery from an economic perspective, but they all focus on the intensive margin (the updating process) while we focus on the extensive margin (what is and is not learned). Several theoretical studies explore the process by which people will sample consumption items if they must learn them from experience, including Easley and Kiefer (1988), Aghion et al (1991), Keller and Rady (1999), Piermont et al (2016), and Cooke (2017). However, these all focus on the experimentation and updating process, and each either includes assumptions that ensure full learning by making learning effectively costless, or does not focus on the completeness of learning.[2] Armantier et al (2016) use theory and a lab experiment to study preference discovery, focusing instead on the experimentation and updating process, testing different theories of learning, and do not consider the potential incompleteness of learning.

Plott (1996) noted that feedback should help the learning process, so we can find suggestive evidence about preference discovery in lab experiments demonstrating unstable choices that are ameliorated over time by feedback. For example, van de Kuilen and Wakker (2006) find that repeated trials without feedback do not reduce Allais paradox violations, but with feedback, the violations decrease. Weber (2003) finds that repeated plays of a strategic game exhibit more apparent learning when feedback is provided.

one's own tastes, which Braga and Starmer (2005) refer to as "institutional learning" and "value learning" respectively. Our focus is on value learning, so we assume the agent knows the objective features of all goods. Institutional learning is best separately modeled, e.g., in experimental consumption models (Kihlstrom et al, 1984) or the two-armed bandit problem (Rothschild, 1974).

[2]Brezzi and Lai (2000) show, in another theoretical study, that learning when facing a multiple-armed bandit will be incomplete, but it is for a different reason (discounting) than what we study.

Other economics experiments on repeated choice with feedback provide further suggestive evidence of preference discovery. Preference instability may be a marker of preference discovery, and there is a large literature debating the importance and interpretation of "preference reversals," e.g., Cox and Grether (1996). Noussair et al (2004) find that with repeated choice, people can converge to a true induced value. Similarly, errors and biases often decline with repeated choice, as observed in the gap between willingness-to-pay and willingness-to-accept, non-dominant bidding behavior, and strategic games (Coursey et al, 1987; Shogren et al, 1994, 2001; List, 2003).

By the same token, preference stability over longer periods is sometimes taken as evidence that discovery is not required, though we argue this is a mistaken inference. While studies like Eckel et al (2009) show that preferences are affected by outside conditions (mediated by psychological affect), other studies (including Andersen et al, 2008 and Dasgupta et al, 2017) look over longer time periods and find some evidence of stability and some evidence that preferences depend on conditions in predictable ways.[3] However, we argue that eventual stability in choices is expected even with preference discovery, and need not indicate fully discovered preferences: if you stop trying new things, you stop learning, and your choice behavior appears stable regardless of whether your preferences are known to you.

Our contribution is to explore a set of intuitive hypotheses about how preference discovery, as described in Plott (1996), might work.[4] We focus on the extensive margin, what items an agent will and will not learn her tastes for, because if items are never tried they can never be learned, thus yielding a large opportunity for welfare loss. We use an induced utility function with fictitious goods for which the subject receives a noisy "guess" about the true value of the good. We experimentally vary the subject's lifetime, their income, the amount of noise in the guesses, and the frequency with which a good appears.

We conduct the experiment through Amazon's Mechanical Turk. Our results support our intuition. First, we find that choice reversals (known in some of the literature as preference reversals) occur, but primarily in

[3]Chuang and Schechter (2015) find, in developing country contexts, very little stability in preferences within a person over years, except in survey measures of self-reported social preferences. However, their interpretation is that the experimental measures they study are not good measures of preferences in these contexts.

[4]In a related paper, Delaney et al (2019), we develop a formal theory exploring the extensive margin of preference discovery.

early rounds. Similarly, we find that subjects reduce their sampling of new goods after a period of time. Welfare loss from not-yet-learned goods declines across rounds but stabilizes at a non-zero level. Incomplete learning appears to persist as an equilibrium outcome.

We further show that this incomplete learning is more common for goods that are rare, have low prior value guesses, or appear in choice sets alongside other goods that seem more attractive. Failure to try a good is also more common for a person with a lower income or a shorter lifetime. More noise in the initial guesses of value cause high-value goods to be less likely to be tried, and low-value goods to be more likely to be tried, which accords with intuition. Over time, subjects develop a pessimistic bias in beliefs about goods' values, since optimistic errors are more likely to be corrected.

This paper proceeds as follows. First, we present intuitive hypotheses about how preference learning might proceed. We then describe the experiment design. Our experiment results follow, and we conclude.

2 Preference Discovery

Imagine that a person is born not knowing her own preferences, so that she has to discover them through consumption experience. Whatever the learning process, if she is willing to sample a good then, given enough time, she will eventually learn her preferences for it. Therefore, the decision to sample a good (or not) when given the opportunity is a central determinant of a person's ability to learn her preferences and thus maximize her utility. For this reason, we focus on the extensive margin of preference learning.

What will drive an agent's initial sampling decisions? If she is fully myopic, she will only sample items that appear immediately attractive. However, if she's forward-looking, she might experimentally consume (Kihlstrom et al, 1984): she might try something to see whether she likes it. Those models typically yield complete learning, but we suggest that in reality, the process of learning preferences has an opportunity cost: to try a new thing that might be good, we must forego something else, in many cases a known pleasure. If so, people should be willing to try some, but not all, of the possible new goods. This suggests some intuitive hypotheses about what this incomplete learning might look like.

First, goods that are rarely available are less likely to have been encountered at any point in time, and thus are less likely to be learned at any point in time. Second, goods that appear less attractive on first impression, or that are part of a choice set that contains other goods that are quite attractive, will be less likely to ever be tried. Third, since experimental consumption yields benefits through potential future consumption, an agent should be less likely to try a good if she expects herself to have less future consumption: if she has a shorter lifetime (or shorter remaining lifetime) or if she has a lower income. Other characteristics should matter as well, but we do not test them in our study: the agent's risk and time preferences and the price of the good. Finally, it might seem like difficulty in guessing how much one will like a good (noisiness of a guess) would only interfere with the likelihood of learning one's preferences, but this need not be true. If a good's true value is high, noise in the signal about its value can only make it less likely for the agent to try it, because the noise increases the chance that the good gives a negative first impression; however, if the good's true value is low, noise could make it appear attractive enough to try.

This pattern of trying items and learning preferences should yield some behavioral consequences. First, early in an agent's consumption life, she will appear to have unstable preferences, exhibiting choice reversals as she samples goods and learns their values. Second, this instability will decline over time and she should eventually have stable choice patterns. Such stability has been noted in studies like Andersen et al (2008) and Dasgupta et al (2017). However, even at that time, she may still have undiscovered preferences because of negative first impressions of goods that she would actually like if she tried them. This would cause persistent welfare loss, and it would mean her choices do not represent what is best for her. Third, she should also have a pessimistic bias in her beliefs about goods' values on average: she will have an accurate belief about items she's tried, but for items she has not tried, her perception is likely to be worse than the reality of the good because optimistically high expectations about goods are corrected as she tries goods, whereas pessimistically low expectations are not corrected.

3 Experiment Design

We present an experiment in which individuals face a decision environment based on our intuitive hypotheses above. The experiment tests the extensive margin of preference learning (what is and is not learned) using induced values for fictitious goods instead of actual consumption of goods. We do this to avoid satiation, to ensure subjects are at the same level of preference learning at the start of the experiment, and because homegrown preferences for actual goods will vary significantly across people, be complicated by preferences for moderation and the potential for variation in access to complementary and substitute goods, and be difficult to observe, thus limiting our ability to test the model's precise predictions.

In the experiment, the subject plays through a series of T rounds. In each round, she has a budget y to spend and is confronted with a basket of available goods, which are randomly chosen from the universe of N goods: each good i appears in each round with probability q_i. She has an induced utility function that is converted to dollars to determine her experiment earnings. The utility function has fixed parameters. The subject starts out not knowing these parameters but receives noisy guesses about them. Each guess is updated to the true value when she has had sufficient experience (which we call a nibble, or minimum meaningful consumption experience, m_i) with that good. This minimum meaningful consumption size ensures that there is a non-vanishing opportunity cost to learning a new good, which undergirds our intuition about why learning might not be complete.

Specifically, her utility is linear in the goods:

$u(x_1, x_2, ..., x_N) = zx_1 + \hat{\beta}_2 x_2 + \ldots + \hat{\beta}_N x_N.$

The values $\hat{\beta}_i$ for the goods are randomly chosen for each subject, and they remain fixed for that subject for all rounds. There is a numeraire good x_1 that is available in all rounds. It gives a known return z and costs 1 per unit. Half of the non-numeraire goods appear with low probability and the rest with high probability. The goods have fixed prices $p_i = 1$. A nibble (minimum meaningful consumption experience) is $m_i = 1$ for all goods.

While she makes her decision in each round t, she sees her true or guessed value β_i^t for each available good. When the experiment starts, these are the priors (guessed values) we assign to her, and as she learns values over the course of the experiment, priors are replaced with true values. We generate each prior by adding an independent random disturbance to the true value. For each subject and each good, the random disturbance is drawn from a discrete uniform distribution over $[-\sigma, +\sigma]$. We call these "starting guesses" and tell the subject that each starting guess value is her true value plus a positive or negative random number, so that it is related to, but generally not the same as, the true value.

In each round, from the set of available goods, the subject must choose a bundle that costs y or less. This decision is time-limited by our software: if

Table 1: Experiment Parameters

Variable	Description	Fixed or varied?
N	Number of goods in the universe	Fixed: 10
q_i	Probability good i appears in a round	Experimentally: 25% or 50%
p_i	Price of good i	Fixed: 1
y	Income	Experimentally: 3 or 6
T	Lifetime (number of rounds)	Experimentally: 10 or 20
z	Value of numeraire good	Fixed: 65
β_i	True value of good i	Random integer in [50, 80]
σ	Max disturbance in "starting guesses"	Experimentally: 25 or 49
m_i	Meaningful consumption experience	Fixed: 1
c	Conversion rate, points to dollars	Fixed: 1000

she does not choose an affordable bundle within a minute then she consumes zero of all goods, earning zero for the round. After the round, the software tells her what her total utility is in that round and reminds her what bundle she chose. For each good, it also tells her what its value or her guess of its value is, as appropriate. The software automatically updates with the correct value each good of which she consumed at least m_i in that round. Since we do not seek to study the subjects' ability to infer parameters of multivariate functions but, rather, whether and when different goods will be tried, our simple design reduces the "learning" problem to a "tasting" problem.

The subject's earnings in a round come from her utility in that round. After all rounds of the experiment are complete, the subject sees a summary of her earnings in each round and the sum of those rounds' earnings in points and in dollars. She then completes a short questionnaire about herself and about the experiment. Her total earnings for the experiment are the sum of her earnings in all rounds, converted to dollars with a conversion rate c, plus an additional \$0.50 for completing the questionnaire.

As shown in Table 1, we experimentally vary across subjects income y, lifetime T, and noisiness in priors σ, so that our experiment has eight cells. Across all cells, all subjects have the same likelihood of each good appearing (which is fixed for any given good but experimentally varied across goods within subject), number of goods, numeraire value, maximum disturbance size, conversion rate, and distribution from which values are drawn.

We gave each good the name of a fictional fruit and we called the nu-

Table 2: Number of Subjects in Each Treatment Cell

Lifetime $T = 10$	Income $y = 3$	Income $y = 6$
Noise $\sigma = 25$	76	91
Noise $\sigma = 49$	85	95
Lifetime $T = 20$	**Income $y = 3$**	**Income $y = 6$**
Noise $\sigma = 25$	74	71
Noise $\sigma = 49$	76	78

meraire good "bread" to make the experiment more engaging and game-like while still limiting their importation of beliefs and tastes from outside the experiment. See Appendix A for full instructions.

We programmed the experiment in oTree (Chen et al, 2016) and deployed it on Amazon's Mechanical Turk (mTurk). Subjects were screened on being US-based and having successfully completed a large number of past mTurk tasks (a 98% or greater success rate on at least 1,000 tasks).

4 Experiment Results

We ran the experiment in February 2018. In all, 1,252 potential subjects signed up to participate, of which 646 completed the experiment.[5] Table 2 shows the number of subjects in each treatment condition. Among these 646 subjects, subjects earned an average of 4,797 experimental points, or \$4.80 plus a \$0.50 participation payment. The first quartile of earnings was \$3.50 and the third quartile was \$7.34.[6]

Our analysis proceeds as follows. First, in Section 4.1, we validate the experiment by showing that subjects choose according to their beliefs often,

[5]Problems with the server caused fatal timeouts for some potential subjects. Of the 606 who did not complete the experiment, 547 (90.3%) had made no choices by the time they stopped. Most of these likely had server timeouts.

[6]The post-experiment questionnaire asked a comprehension question that posed a simplified version of the experiment's choice problem. 82.4% of subjects answered correctly. Including only those who answered correctly produces qualitatively identical results except that the Mann-Whitney test for the effect of noise on efficiency becomes insignificant and the effect of noise on efficiency becomes significant at the 10% level in the Tobit regression for $T = 10$. This paper reports results from the full sample of subjects.

but not always, with some deviations consistent with learning and others consistent with error. We next, in Section 4.2, show that choice reversals exist and decline with time. We then show in Section 4.3 that most, but not all, goods are tried. Then in Section 4.4 we show how this engenders different degrees of eventual preference learning. Next, we show that believed preferences become increasingly stable. We demonstrate the pessimistic bias predicted. Finally, in Section 4.5 we show that welfare loss occurs and declines over time, but, importantly, it does not decline to zero.

4.1 Consistency with Believed Preferences

To show the extent to which subjects choose myopically according to their believed preferences, we construct a dummy variable for each subject for each round, and we give it a value of 1 if the subject chose the bundle that maximizes believed utility. The subject's believed utility is based on the parameters of goods they have learned up to that round, and the point estimates ("starting guesses") they have for goods they have not yet tried. For goods that are not yet learned, the assumption that a rational person will maximize current-period expected value based on these parameter point estimates is only strictly true for myopic risk neutral people, as risk averse or risk loving people will have a bias (against or in favor of, respectively) untried goods, and experimental consumption by definition will cause people to diverge as well. We use this as a starting point for our analysis, and discuss divergences from this simple myopic optimization below.

Pooled across all treatments and rounds, people maximize believed expected earnings in this way 61.0% of the time. In round 1, subjects choose in accordance with their believed preferences at a rate of 43.2% for $T = 10$ and 41.8% for $T = 20$. At the end of experimental lifetimes, that value is 65.4% in round 10 for $T = 10$ and 69.2% in round 20 for $T = 20$, a significant increase (within-subject signed-rank test: $p < 0.001$ in both cases, $n_{T=10} = 347$, $n_{T=20} = 299$).[7]

Recall that experimental consumption explains some choices that don't maximize believed preferences. Since experimental consumption has no further value in the final round, why would a subject make a non-myopic-

[7]Subjects made other non-maximizing choices as well. Of 103,950 good-round pairs, subjects chose a value between 0 and 1 (less than a meaningful consumption experience) 322 times (or 0.3% of the time), and a value less than 0 a total of 14 times (less than 0.1% of the time). 99.3% of the time, subjects chose an integer between 0 and 6.

maximizing choice in her last round? One potential reason is that subjects who are risk averse may choose a learned good with a lower parameter value over an as-yet-unlearned good with a higher "starting guess" to avoid uncertainty. Subjects who are risk loving may do the opposite, choosing an as-yet-unlearned good with a lower "starting guess" over a learned good with a higher parameter value.

Neither risk preferences nor experimental consumption can explain non-maximizing choices among goods that have already been learned. In 21% of all choices we observe, subjects choose a good they have tried before with a lower known value than another available good with a higher known value (in some cases also forgoing untried available goods with higher "starting guesses"). These choices likely indicate error. As noted above, these errors decline significantly over time. Further, the magnitude of most of these errors is small: of these choices, 63.7% choose a good that's only dominated by a small amount (between 1 and 9 units of absolute value). This means that in 92.3% of all choices, an error cannot be identified (although they may choose something with a lower prior) or we identify only a "small" error.

Further, subjects who make non-myopically-maximizing choices in the final period do not seem to suffer in our study: on average, those who behave inconsistently achieve 95.1% efficiency while those who behave perfectly consistently achieve 95.0% efficiency (not statistically different, $p = 0.908$).[8]

These results show that subjects are engaging in some optimizing choice as proposed in our model, but that they are quite a bit less sophisticated than they could be. To the extent to which error enters into our subjects' decision process, that introduces noise that makes it harder for us to detect the empirical results we find in the remainder of this paper.

Regarding the effects of treatment variables (T, y, and σ) on subjects' behavior, we conjectured that experimental consumption depends on current period sacrifice and discounted expected potential gain therefrom. Table 3 confirms these insights. In particular, for tests pooled across rounds, we see

[8] Choices inconsistent with believed preferences have little impact on our experiment's results. Excluding subjects who make these non-myopic-maximizing choices yields qualitatively similar results with only a few changes: In Table 3, the difference in efficiency across income levels becomes marginally significant ($p = 0.090$), while the difference in efficiency across noise levels ceases to be statistically significant ($p = 0.118$). In Table 4, the number of remaining rounds becomes significant at $p < 0.001$ in the second model, while in the third model, noise becomes marginally significant ($p = 0.064$) as does income ($p = 0.063$).

Table 3: Nonparametric Tests of Treatment Effects on Learning Outcomes

Lifetime T	Choices inconsistent with beliefs	Full discovery	Efficiency
10	0.444	0.159	0.846
20	0.358	0.502	0.896
p-value	0.000	0.000	0.000
Income y	**Choices inconsistent with beliefs**	**Full discovery**	**Efficiency**
3	0.375	0.270	0.878
6	0.432	0.361	0.861
p-value	0.001	0.013	0.355
Noise σ	**Choices inconsistent with beliefs**	**Full discovery**	**Efficiency**
25	0.405	0.340	0.876
49	0.404	0.296	0.863
p-value	0.941	0.237	0.035

All variables are aggregated to the subject level. N's can be inferred from Table 2. "Full discovery" captures whether a subject has tried every good by the end of the experiment. "Choices inconsistent with beliefs" is the proportion of rounds in which a subject's choices do not maximize expected utility given beliefs. "Efficiency" is the utility achieved as a proportion of the maximum achievable. p-values are from Mann-Whitney tests.

that subjects with longer lifetimes and lower incomes made fewer choices that are myopically inconsistent with their beliefs. If experimental consumption happens more in early than in later rounds, then a longer lifetime (as compared to a shorter lifetime) should give more rounds in which little experimenting is happening, thus explaining why longer lifetimes are associated with choices more consistent with beliefs. Higher income yielding more choices inconsistent with beliefs could happen because higher incomes should yield more experimentation. We return to the rest of the results in Table 3 later in this section.

We can seek evidence that some choices that diverge from maximizing

Table 4: Drivers of Utility Maximization Deviations and Choice Reversals

	Non-maximizing choice	Choice reversal (all rounds)	Choice reversal (rounds > 5)
Remaining rounds	0.017***	0.000	0.006***
	(0.001)	(0.001)	(0.002)
Lifetime $T = 20$	-0.170***	-0.037**	-0.081***
	(0.015)	(0.016)	(0.000)
Noise $\sigma = 49$	-0.003	0.009	0.027
	(0.014)	(0.014)	(0.017)
Income $y = 6$	0.053***	0.050***	0.050***
	(0.014)	(0.014)	(0.017)
Constant	0.340***	0.296***	0.295***
	(0.016)	(0.017)	(0.019)
R^2 (overall)	0.0399	0.0044	0.0087
n subjects	646	646	646
n subject-rounds	9,450	8,804	6,220

Robust standard errors in parentheses. *** $p < .01$, ** $p < .05$, * $p < 0.1$. Random effects OLS panel regressions at the subject-round level with errors clustered at the subject level. For treatment variables, we use dummies that are equal to 1 for the higher value.

believed utility are experimental consumption by regressing the dummy for deviation from believed preference maximization on factors that should affect the value of experimental consumption. Table 4 shows OLS results. (Logit and probit results are similar.) Belief-inconsistent choices increase with remaining lifetime and endowment and decrease with overall lifetime. These results are consistent with subjects making inconsistent choices early as they learn their preferences, and then increasing consistency as their understanding of their preferences improves. We return to the rest of Table 4 in the next subsection.

4.2 Choice Reversals

Now we turn to choice reversals. For each subject in each round, we infer whether her choice contradicted the ranking implied by a past choice, and we call such contradictions choice reversals. In other words, if goods A and

B were available in round 1 and the subject chose more of A than of B, but in round 2 when both were available she chose more of B than of A, that is a choice reversal.

We conjectured that choice reversals would occur and would decline over time. We test this conjecture in the latter two columns of Table 4, using OLS panel regressions at the subject-round level. Each experiment subject needs some time to build up a choice profile that can be contradicted. In the first round, it is impossible to observe a choice reversal because there is no past ordering to contradict. If the sets of goods available in rounds 1 and 2 are disjoint, then it is also impossible to witness a choice reversal in round 2. For this reason, the second column presents a regression model that includes all rounds except the first, and the third column includes only rounds 6 and up (when half or a quarter of the subjects' lifetimes have passed) to allow a choice profile to be established. This third column is our preferred specification.

While the specification in Table 4 that includes all rounds does not show an effect of remaining rounds on choice reversal rate, our preferred specification (which excludes the first five rounds) does. The latter indicates that choice reversals decline over time, as predicted, and the former indicates that this is confounded by the mechanical difficulty in observing reversals in early rounds.

In Table 4, we see that a longer lifetime reduces the rate of choice reversals, while a higher income increases the rate of choice reversals. Thus, the same experimental factors that drive inconsistent-with-belief choices also drive choice reversals, giving further evidence of experimentation.

4.3 Trying Goods

Next, we examine subjects' tendency to try goods. Most subjects try most, but not all, goods that they have the opportunity to try. Figure 1 plots over time the proportion of all goods of which at least a nibble (the minimum that must be tried to learn) has been consumed on average, as well as the proportion tried out of all goods that have appeared (and thus could be chosen). The raw proportion of total goods tried increases at a decreasing rate until it levels off at around 87% approaching Round 20. The proportion of possible goods tried shows a similar trend. By the end of 20 rounds, each subject had been presented with 98.5% of all possible goods on average, and 85.3% of subjects were presented with all 11 goods by their final round.

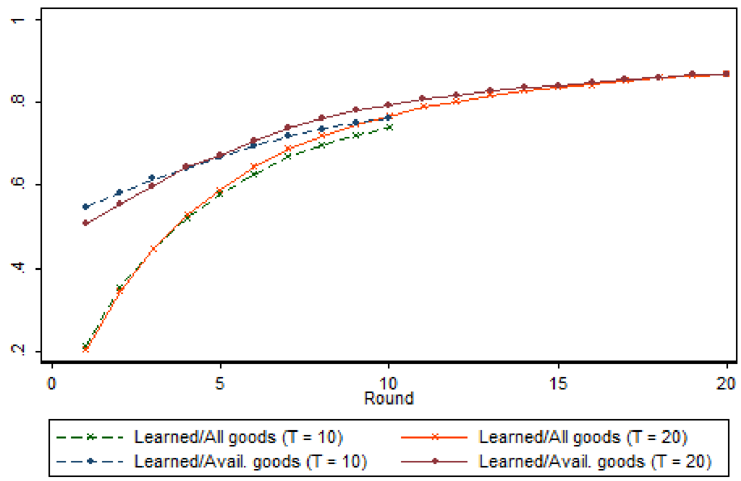

Figure 1: Percent of Goods Learned by Round

The vertical axis the proportion of all goods (or of goods that have appeared) that subjects have by the given round chosen at least $m_i = 1$ unit of, averaged across subjects.

This tendency to try more goods as time progresses (but at a decreasing rate) is not proof of our conjectures, as even random choice would yield this outcome, so we move on to more interesting results.

If subjects were myopic, i.e., if there was no experimental consumption, life length would not affect tendency to try goods, in which case the $T = 10$ and $T = 20$ lines in Figure 1 would coincide, but they do not. Those with shorter lifetimes try more of the available goods in the first round ($p < 0.001$, rank-sum test at subject level), but have tried a smaller proportion of the available goods in the tenth round, which is the last round for the subjects with shorter lifetimes but only half-way through for those with longer lifetimes (again $p < 0.001$, rank-sum test at subject level). This crossing of the lines is not consistent with perfectly forward-looking learning, which would imply that people with longer lifetimes get a higher benefit for trying new goods early in their lifetime. Specifically, while forward-looking learning comports with the result that the $T = 20$ line is steeper than the $T = 10$ line, it does not comport with the result that the $T = 10$ line starts at a higher intercept. We confirm this pattern using regression methods below.

For both lifetime lengths, both curves end far short of 100%. For subjects with a lifetime of 10 rounds, 15.9% try every good by the last round. This increases to 50.2% for subjects with a lifetime of 20 rounds. These values differ, $p < 0.001$, based on a rank-sum test at the subject level. While subjects in our experiment have finite experimental lifetimes, the fact that learning seems to flatten out before the end of life is supportive of our intuition that some goods will never be tried even in infinite time.

We suggested that agents with longer lifetimes and larger incomes would be more likely to learn their preferences, and the nonparametric tests in Table 3 confirm this intuition.

We also suggested characteristics that should predict whether a good is tried. In Table 5, we report a panel regression with one observation per subject per good per round, where the outcome variable is a dummy indicating whether this subject has learned her preference for this good as of this round, i.e., whether by this round she has tried it in at least the minimum size needed to learn. We show results from an OLS regression; results are similar for logit and probit. These regressions include the numeraire as a good.[9] Our preferred specification is III, which includes the numeraire and round-lifetime interactions (which we discuss below), but results are consistent across specifications.

We find again that subjects with longer lives and larger incomes try more goods. In Model II, we see that subjects have tried more goods as time passes, which appears to account for the effect of the longer lifetime; this is not one of our key results but is a reasonable sanity check. In Model III we examine time trends in the learning process by separately estimating the intercept and the time trend during the first ten rounds for the two lifetime treatments using additional dummy variables and their interactions with the round. The coefficient on the "Round x First 10 rounds" term indicates that those with longer lifetimes try goods more quickly than do those with shorter lifetimes ($p = 0.014$ on the slope interaction coefficient). They do not appear to try a larger number of goods at the outset ($p = 0.353$ on the the "First 10 rounds" dummy-intercept) despite, in principle, a higher payoff to learning. In other words, we have no explanation for why people with a longer lifetime don't try more goods in the first round as compared to people with a shorter lifetime; however, after that, people in the longer-lifetime treatment do have a faster rate of learning as we conjectured.

[9] Recall that the numeraire value is known with certainty and thus is always "learned." Results are similar excluding the numeraire. We also considered a version of specification III that used individual period dummies and found qualitatively similar results.

Table 5: Factors driving whether a good is learned

	I	II	III
Lifetime $T = 20$	0.145***	-0.014	
	(0.014)	(0.014)	
Income $y = 6$	0.065***	0.065***	0.065***
	(0.014)	(0.014)	(0.014)
Noise $\sigma = 49$	0.155***	0.155***	0.155***
	(0.045)	(0.045)	(0.045)
Prior	0.007***	0.007***	0.007***
	(0.001)	(0.001)	(0.001)
Prior x σ	-0.002***	-0.002***	-0.002***
	(0.001)	(0.001)	(0.001)
True value	-0.001***	-0.001***	-0.001***
	(0.001)	(0.001)	(0.001)
Average of	-0.006**	-0.006**	-0.006**
other values	(0.003)	(0.003)	(0.003)
Probability of	0.422***	0.423***	0.423***
appearance	(0.019)	(0.019)	(0.019)
Round		0.032***	0.055***
		(0.001)	(0.001)
L: $T = 20$			-0.013
First 10 rounds			(0.014)
L: $T = 20$			0.448***
Last 10 rounds			(0.019)
L: $T = 20$			0.004**
Round x First 10 rounds			(0.002)
L: $T = 20$			-0.046***
Round x Last 10 rounds			(0.001)
Constant	0.351*	0.176	0.0490
	(0.195)	(0.195)	(0.195)
R^2 (overall)	0.1250	0.2013	0.2190
n subjects	646	646	646
n subject-goods	7,106	7,106	7,106
n subject-good-rounds	103,950	103,950	103,950

Robust standard errors in parentheses. *** $p < .01$, ** $p < .05$, * $p < 0.1$
Random effects OLS panel regressions at the subject-good-round level with errors clustered at the subject level. For treatment variables, we use dummies that are equal to 1 for the higher value. Model III includes two dummies for the Lifetime $T = 20$ treatment group, one for the first ten periods and one for the last ten. It also includes interactions of these dummies with the round.

We also find evidence consistent with our intuitive hypotheses that goods are more likely to be learned if they have a higher prior belief or probability of appearance. A higher value of other goods decreases the likelihood that a given good has been tried.

Consider now the interaction between noise and priors. The prior can range from 1 to 129. Based on Specification III, the effect of noise ranges between $0.155 + (-0.00238) * 1 = 0.153$ for the lowest possible prior and $0.155 + (-0.00238) * 129 = -0.152$ for the highest possible prior. This confirms our conjecture that (for risk averse agents) goods with low priors would be more likely to be tried with more noise, and goods with high priors would be more likely to be tried with less noise.

4.4 Learning Benchmarks, Stability, and Pessimism

This tendency to try some but not all goods has predictable ramifications for the levels of discovery that our subjects achieve. We define *full discovery* as the state in which the agent has learned her preferences for all goods. We define *full relevant discovery* as the state of having learned all goods that are better than the numeraire; if one achieves this benchmark then one will not lose welfare because of misunderstood preferences. By the end of the experiment, only 25.9% achieve full relevant discovery for $T = 10$ while 57.9% achieve full relevant discovery for $T = 20$.

A weaker benchmark is *full voluntary discovery*, a state in which the subject has tried all things that are "attractive" in the sense that she might be willing to try them based on some form of utility maximization. This is an important learning benchmark because it captures a situation in which we would expect to observe choice stability; once the subject reaches full voluntary discovery, she should always choose a known good or the numeraire rather than any of the remaining unknown goods. From this state, the agent can still lose welfare from misunderstood preferences, but when she is in this state, her beliefs and optimization process provide no basis to try new goods and exit the state.

Defining this state requires some consideration as to how sophisticated the subject is; after all, results already presented show that while subjects do experimentally consume, they appear to be less than fully sophisticated. We track full voluntary discovery with two extreme cases. *Minimally experimental full voluntary discovery* represents full voluntary discovery for myopic agents. In this state, the subject has learned all goods with priors higher than the numeraire; in this case, since the numeraire is an always-present "outside option," the subject is stuck in this state if she is unmoved by the potential information that would be gained by trying an untried good. This informational value may be significant, so we also consider the very sophisticated *maximally experimental full voluntary discovery.* This is a state in which the subject has tried all goods above a much lower threshold: all goods with a prior value such that the expected cost of trying m_i units of good i is less than the expected benefit of future consumption of good i in the future given the possibility that experimentation would reveal it to be worth consuming. Since this threshold value depends on key factors about preferences and availability of other goods, we calculate it conservatively, i.e., under the assumptions that would maximize that benefit.[10] Only 59.4% achieve minimally, and only 33.4% achieve maximally, experimental full voluntary discovery for $T = 10$, while 89.3% achieve minimally, and 66.2% achieve maximally, experimental full voluntary discovery for $T = 20$. Figure 2 shows the proportion of subjects who reach full discovery, full relevant discovery, and these measures of full voluntary discovery over time. Since a longer lifetime is a more appropriate place to look for eventual learning outcomes, for visual simplicity we only show results for $T = 20$.

Recall that an agent who has achieved full voluntary discovery may stop trying goods she has not already learned. We declare a subject a candidate for persistent welfare losses if she has reached full voluntary discovery but not full relevant discovery. This is a relatively conservative definition, since given the flattening out of the learning curve, we infer that some subjects who have not achieved full voluntary discovery by our definition may be unwilling to sample new goods. This may be in part due to risk aversion. At the end of their experimental lives, 35.5% of subjects with $T = 10$ and 32.8% of subjects with $T = 20$ are candidates for persistent welfare loss using the minimally experimental definition of full voluntary discovery. These are not significantly different ($p = 0.476$ from a subject-level rank-sum test). Recall that the minimally experimental definition assumes agents are not forward-looking with respect to the information value of trying unpleasant-

[10]The calculation of this threshold is shown in an online appendix. A key assumption is that in all future periods, only the numeraire good and this good will be available. As we show in the online appendix, an agent has achieved maximally experimental full relevant discovery according to this definition if she has tried all goods i with priors at least as large as $X_i^t = 64.5 - \sigma + \sqrt{240.25 - 240 \cdot \frac{q_i \cdot y \cdot (T-t)}{1+q_i \cdot y \cdot (T-t)}}$. Note that this threshold moves higher as t progresses.

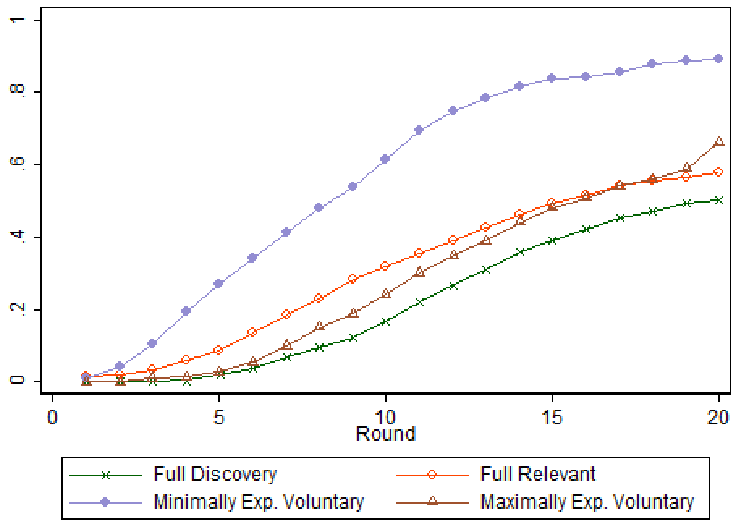

Figure 2: Achievement of Learning Benchmarks

seeming goods. Using the maximally experimental definition of full voluntary discovery, at the end of their experimental lives, 12.7% of subjects with $T = 10$ and 12.4% of subjects with $T = 20$ are candidates for persistent welfare loss. These are again not significantly different ($p = 0.907$ from a subject-level rank-sum test). A sizable proportion of subjects, regardless of their experimental lifespan, may have reached a point at which they are done experimenting in spite of incomplete learning. The extent of this problem depends on whether they are minimally or maximally experimental in their approach to new goods.

The leveling off of the learning curves in Figures 1 and 2 supports the idea that believed preferences eventually become stable even in our subjects' finite experimental lifetimes, but we can test that hypothesis explicitly. We construct a variable for each subject for each round (starting at round 2) that indicates how many parameters changed between this round and the preceding round. While the average number of changes in rounds 2-6 is 0.911 for $T = 10$ and 0.972 for $T = 20$, the average number of changes as subjects near the ends of their life spans is much smaller: in rounds 7-8 for $T = 10$ it is 0.389, and in rounds 15-18 for $T = 20$ it is 0.080. The difference is significant in both cases (sign-rank test at the subject level: $p < 0.001$ in both cases). Of the 299 subjects with $T = 20$, 256 (85.62%) chose no new goods in the final three rounds, and 224 (74.92%) chose no new goods in the final five rounds. Subjects appear to be closer to minimally rather than maximally experimental, which would imply that more are stuck in a state of persistent welfare loss.

We have shown, then, that subjects in our experiment are learning their preferences but are not learning them completely over the course of finite but long lifetimes. We hypothesized that this would lead to a pessimistic bias over time because positive misperceptions would be more likely to be corrected by experience. We test this by constructing a variable for each subject for each round that averages the subject's parameter belief errors, where each error is her current believed value minus her true value. In round 1, this error averages -0.156 across all subjects. This, as expected, is not significantly different from zero (t-test $p = 0.551$). At the end of subjects' experimental lifetimes, this value is significantly negative: -2.893 in round 10 for $T = 10$ and -2.038 in round 20 for $T = 20$. These values are significantly different from zero (t-test $p < 0.001$ in both cases).

Figure 3 shows how this pessimism evolves. The average error declines quickly as positive errors correct themselves faster than negative ones. The average error then levels off and starts to climb as subjects choose goods with small negative errors, correcting these errors. If we measure the bias among only undiscovered goods (thus eliminating the zero errors from the average), the mean error is -15.34 in round 20, 7.5 times the unconditional mean error at the same point. In other words, subjects' average beliefs about goods they have never tried steadily diverge from the true value, and display a persistent and sizable pessimistic bias.

4.5 Efficiency

Finally, we turn to the welfare implications of the learning process and its failures. We calculate an efficiency measure for each subject for each round as the utility achieved in that round divided by the maximum she could have achieved if she had chosen according to her true preferences. Averages of this measure by treatment pooled across rounds are shown in Table 3. Longer lifetimes and lower noise in priors both yield higher efficiency, which accords with results we have already shown about learning in those cases. Income does not affect efficiency (and we had no intuition that it would).

To look at how welfare evolves across rounds in the different treatments, we run a panel Tobit regression at the subject-round level, which we report

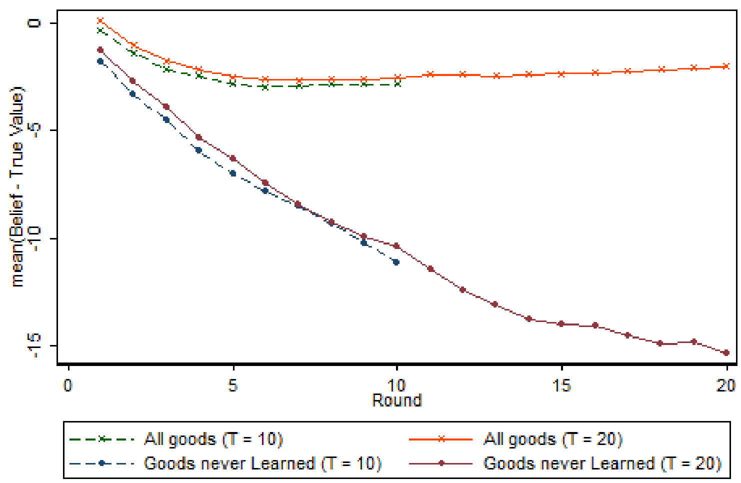

Figure 3: Average Error in Beliefs
The vertical axis is the mean error by round: the difference between subjects' prior value and true value, averaged across goods.

in Table 6. As time passes, efficiency loss declines. The effect of time is nonlinear, however: efficiency improves at a decreasing rate over time. Once we control for round number, life length ceases to have an effect, and in our regression, we see that the effect of income is only significant for $T = 20$.

A declining welfare loss is not particularly surprising and could have obtained as a result of other processes, as long as those processes involve optimization. The most important conjecture we made is that this welfare loss need not decline to zero even as time approaches infinity. While our experiment subjects are not infinitely-lived, as we show in Section 4.4, choices become quite stable by the end of our subjects' experimental lifetimes, particularly for those in the $T = 20$ treatment, and yet subjects do not discover their preferences for all goods that are better than the numeraire; therefore, if welfare is still being lost in the last period, this would confirm our most important conjecture. Indeed, we find that in period 10 of the $T = 10$ treatment, efficiency is 90.7% and in period 20 of the $T = 20$ treatment, efficiency is 95.4%. Thus, we have shown that our warning that as-yet-unlearned preferences could cause loss forever is borne out by our experiment.

Table 6: Determinants of Efficiency

	Pooled	$T = 10$	$T = 20$
Lifetime $T = 20$	0.003		
	(0.020)		
Income $y = 6$	-0.028	-0.011	-0.048*
	(0.019)	(0.027)	(0.027)
Noise $\sigma = 49$	-0.025	-0.034	-0.016
	(0.019)	(0.027)	(0.027)
Round	0.057***	0.072***	0.058***
	(0.002)	(0.008)	(0.003)
Round^2	-0.001***	-0.003***	-0.001***
	(0.0001)	(0.0007)	(0.0001)
Constant	0.698***	0.666***	0.702***
	(0.021)	(0.031)	(0.027)
Number censored at 0	554	242	312
Number censored at 1	3,979	1,063	2,916
n subjects	646	347	299
n subject-rounds	9,450	3,470	5,980
χ^2	1,607.14	464.17	1,093.86

Robust standard errors in parentheses. *** $p < .01$, ** $p < .05$, * $p < 0.1$. Tobit panel regressions at subject-round level with bootstrapped standard errors.

5 Conclusion

Most work in economics implicitly or explicitly assumes that people know what they like. We explore the idea that if self-knowledge is not endowed at birth but rather achieved through experience, as suggested by the discovered preference hypothesis of Plott (1996), then even the most rational and sophisticated people may fail to learn all of their preferences. At the heart of this failure is the fact that learning has an opportunity cost, and thus complete learning would be irrational. In this paper, we present the results of an induced-value experiment that supports this basic idea, and that explores the dimensions of that learning process, focusing on the extensive margin of learning: what goods are tried. We show that there are intuitively sensible patterns with regard to the types of goods, people, and information that yield better and worse learning outcomes.

Preference discovery processes can explain choice instabilities observed in observational and laboratory studies of behavior, especially in cases of items that are unlikely to have been "consumed" often by the agent. Moreover, stable choice behavior does not indicate that agents are choosing according to their true underlying preferences: they may simply have stopped experimenting. While goods in our study could be bought in continuous quantities, if choice items are discrete and have large consequences (like houses, jobs, or life partners), learning problems are likely to be worse; the analogy in our design is to goods that have a larger "nibble" (minimum consumption) size. Another element that would render learning particularly challenging is an agent's inability to directly assess a good's value even when she "consumes" it, as might be the case for credence goods, donations to charity, and environmental valuation. Indeed, the situations we suggest are most likely to give rise to learning failure correlate to the contexts that Thaler and Sunstein (2008) argue cause people to make bad decisions: cases where the agent is inexperienced and poorly informed, and where she will receive little feedback.

The preference discovery process must be studied in more detail and in more settings to understand how factors internal and external to the agent affect learning and thus welfare loss. It is possible that an agent's mental simulation of consumption can allow some learning without consumption, and if so, that would alleviate some of the issues we highlight. On the other hand, many other factors, like the tendency to forget once learned and any stochasticity in the experience of a good, would likely exacerbate learning problems.

In contexts in which learning one's preferences through direct experience is very difficult, our experiment results indicate that losses could persist; if the choices are important, like choices regarding a house or a job, the losses could be large, and, as Thaler and Sunstein (2008) note, policy-relevant. If agents are aware of the problems we identify, for important decisions, they may turn to other processes or criteria instead of discounted expected utility maximization based on beliefs. For example, people may reduce a complex housing decision to a simpler problem about their beliefs about the value of an asset appreciating over time. Future research could identify whether people do this and whether it seems to be welfare-enhancing, and could study whether specific nudges can help the preference learning process or can effectively replace it.

If we must learn through experience to know our own preferences, the implications are large. On the one hand, this concept can provide new insights on how to get people to try new things, whether in the case of a company marketing a product or a government or non-profit promulgating a green technology. On the other hand, it shows that cross-sectional choice data from any experimental or observational setting may be contaminated by unstable parameters. Worse, choices that appear stable and rational may not reflect what is actually best for the individual making the decision. A tenet undergirding most economics-based policy advice is that people know what's best for them; but if we have undiscovered preferences, that might not be true.

Acknowledgements

We are grateful for helpful comments from the editor and two anonymous referees. For advice early on, we thank Yongsheng Xu, Annemie Maertens, and participants at FUR 2012, SABE/IAREP/ICABEEP 2013, and seminars at Williams College and George Mason University, and we particularly thank CeMENT 2014 participants Brit Grosskopf, Muriel Niederle, J. Aislinn Bohren, Angela de Oliveira, Jessica Hoel, and Jian Li for detailed feedback. We gratefully acknowledge funding from the Williams College Hellman Fellows Grant.

6 References

References

Aghion P, Bolton P, Harris C, Jullien B (1991) Optimal learning by experimentation. The Review of Economic Studies 58(4):621–654, URL http://www.jstor.org/stable/2297825

Andersen S, Harrison GW, Lau MI, Rutstrom EE (2008) Lost in state space: Are preferences stable? International Economic Review 49(3):1091–1112

Ariely D, Loewenstein G, Prelec D (2003) "coherent arbitrariness": Stable demand curves without stable preferences. The Quarterly Journal of Economics 118(1):73–105

Armantier O, Lévy-Garboua L, Owen C, Placido L (2016) Discovering preferences: A theoretical framework and an experiment

Becker GS (1996) Accounting for tastes. Harvard University Press

Braga J, Starmer C (2005) Preference anomalies, preference elicitation and the discovered preference hypothesis. Environmental and Resource Economics 32(1):55–89

Brezzi M, Lai TL (2000) Incomplete learning from endogenous data in dynamic allocation. Econometrica 68(6):1511–1516, DOI 10.1111/1468-0262.00170, URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0262.00170

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

Chuang Y, Schechter L (2015) Stability of experimental and survey measures of risk, time, and social preferences: A review and some new results. Journal of Development Economics 117:151 – 170, DOI https://doi.org/10.1016/j.jdeveco.2015.07.008, URL http://www.sciencedirect.com/science/article/pii/S0304387815000875

Cooke K (2017) Preference discovery and experimentation. Theoretical Economics 12(3):1307–1348

Coursey DL, Hovis JL, Schulze WD (1987) The disparity between willingness to accept and willingness to pay measures of value. The Quarterly Journal of Economics 102(3):679–690

Cox JC, Grether DM (1996) The preference reversal phenomenon: Response mode, markets and incentives. Economic Theory 7(3):381–405

Dasgupta U, Gangadharan L, Maitra P, Mani S (2017) Searching for preference stability in a state dependent world. Journal of Economic Psychology 62(Supplement C):17 – 32, DOI https://doi.org/10.1016/j.joep.2017.05.001, URL http://www.sciencedirect.com/science/article/pii/S0167487016305840

Delaney JJ, Jacobson SA, Moenig TP (2019) A theory of preference discovery

Easley D, Kiefer NM (1988) Controlling a stochastic process with unknown parameters. Econometrica 56(5):1045–1064, URL http://www.jstor.org/stable/1911358

Eckel CC, El-Gamal MA, Wilson RK (2009) Risk loving after the storm: A bayesian-network study of hurricane katrina evacuees. Journal of Economic Behavior and Organization 69(2):110 – 124, DOI https://doi.org/10.1016/j.jebo.2007.08.012, URL http://www.sciencedirect.com/science/article/pii/S0167268108001741, individual Decision-Making, Bayesian Estimation and Market Design: A Festschrift in honor of David Grether

Kahneman D, Snell J (1990) Predicting utility. In: Hogarth RM (ed) Insights in decision making: A tribute to Hillel J. Einhorn, Chicago and London: University of Chicago Press, pp 295–310

Kahneman D, Wakker PP, Sarin R (1997) Back to Bentham? explorations of experienced utility. The Quarterly Journal of Economics 112(2):375–405

Keller G, Rady S (1999) Optimal experimentation in a changing environment. The Review of Economic Studies 66(3):475–507, URL http://www.jstor.org/stable/2567011

Kihlstrom RE, Mirman LJ, Postlewaite A (1984) Experimental Consumption and the 'Rothschild Effect.'., Studies in Bayesian Econometrics, vol. 5.

New York; Amsterdam and Oxford: North-Holland; distributed in U.S. and Canada by Elsevier Science, New York, pp 279 – 302

van de Kuilen G, Wakker PP (2006) Learning in the Allais paradox. Journal of Risk and Uncertainty 33(3):155–164

Lichtenstein S, Slovic P (2006) The construction of preference. Cambridge University Press

List JA (2003) Does market experience eliminate market anomalies? The Quarterly Journal of Economics 118(1):41

Loewenstein G, Adler D (1995) A bias in the prediction of tastes. The Economic Journal 105(431):pp. 929–937, URL http://www.jstor.org/stable/2235159

Noussair C, Robin S, Ruffieux B (2004) Revealing consumers' willingness-to-pay: A comparison of the BDM mechanism and the Vickrey auction. Journal of Economic Psychology 25(6):725–741

Piermont E, Takeoka N, Teper R (2016) Learning the krepsian state: Exploration through consumption. Games and Economic Behavior 100:69 – 94, DOI https://doi.org/10.1016/j.geb.2016.09.002, URL http://www.sciencedirect.com/science/article/pii/S0899825616300896

Plott CR (1996) Rational individual behaviour in markets and social choice processes: The discovered preference hypothesis. In: Arrow KJ, et al (eds) The rational foundations of economic behaviour: Proceedings of the IEA Conference held in Turin, Italy, IEA Conference Volume, no. 114. New York: St. Martin's Press; London: Macmillan Press in association with the International Economic Association, pp 225–250

Rothschild M (1974) A two-armed bandit theory of market pricing. Journal of Economic Theory 9(2):185–202

Scitovsky T (1976) The joyless economy: An inquiry into human satisfaction and consumer dissatisfaction. Oxford University Press

Shogren JF, Shin SY, Hayes DJ, Kliebenstein JB (1994) Resolving differences in willingness to pay and willingness to accept. American Economic Review 84(1):255–270

Shogren JF, Cho S, Koo C, List J, Park C, Polo P, Wilhelmi R (2001) Auction mechanisms and the measurement of WTP and WTA. Resource and Energy Economics 23(2):97–109

Thaler RH, Sunstein CR (2008) Nudge: Improving Decisions about Health, Wealth, and Happiness. New Haven and London:

Weber RA (2003) learning with no feedback in a competitive guessing game. Games and Economic Behavior 44(1):134 – 144, DOI http://dx.doi.org/10.1016/S0899-8256(03)00002-2, URL http://www.sciencedirect.com/science/article/pii/S0899825603000022

Wilson TD, Gilbert DT (2005) Affective forecasting: Knowing what to want. Current Directions in Psychological Science 14(3):131–134, DOI 10.1111/j.0963-7214.2005.00355.x, URL http://cdp.sagepub.com/content/14/3/131.abstract

A Appendix: Experiment Screens

Appendix B: Experiment Screen Shots

Welcome to the experiment! You will first read some instructions, then make a series of decisions, then answer a quick survey. This will take up to 10-20 minutes. Your earnings depend on your decisions and on chance, but will be between $0.50 and $10.

Next

Instructions: Decisions

You will make decisions over a series of 10 rounds. In each round, you will start with some money: 3 francs. You will spend your money buying fruits, which each has a price of 1 franc, and/or bread, which also costs 1 franc. You will have up to a minute to make your decision in each round! If you do not submit a decision in a round, you will buy no goods and thus earn nothing for that round.

You will get value (in points) from the foods you buy. Bread always earns you 50 points points per franc you spend on it. There are several fruits; these are not normal fruits you see every day, but fruits with names we made up. Each fruit gives you a particular value, and each fruit's value stays the same for the whole experiment. Some fruits will appear more often than others, but the chance that a given fruit will appear stays the same across all rounds. Not all the fruits will be available in all rounds.

Your food earnings in a round is the sum of the values you get from all the fruits and/or bread you buy. Your total earnings in the experiment is the sum of your food earnings in each round plus $0.50 for filling out a short survey at the end of the experiment.

For example, imagine you have 4 francs. Imagine that an apple gives you 200 points of value, and an orange gives you 100 points of value. Bread, as stated above, gives you 50 points points of value. If you buy one apple, two oranges, and one bread, how much value do you earn in points?

Next

Instructions: Value

In the example above, we told you what your values were for each fruit. But in the experiment, you will not be told those values.

At the beginning of the experiment, you will be given a "starting guess" for the value for each fruit. That will be related to the fruit's real value for you: it will be the actual value plus or minus some random number.

At the end of each round, you will learn how much value you got from bread and from each of the fruits you chose. You will only learn your value from any fruit you chose at least 1 unit of. Your guesses for each fruit will be updated with these values.

In future rounds, to help you make your decisions, you will see all of the values for goods whose values you have learned, and you will see your starting guess for the goods whose values you haven't learned yet.

Next

Instructions: Summary and Earnings

In summary, in each of 10 rounds, you will choose how to spend your 3 francs buying fruits and/or bread. Bread is always available, but whether each fruit appears depends on chance. You will earn money based on the values you get for each fruit and/or bread you buy. You will start out not knowing for sure the values you get for each fruit, but you will be given starting guesses that are related to your actual values (they are the true values plus or minus a random number). Bread always gives you 50 points per franc.

After you choose how much you want to spend on each fruit and/or on bread, you will learn the value you got from any fruit you bought at least 1 unit of, and that will update your guesses with these actual values.

In each round, you will have one minute to make your food choice, and 30 seconds to review the information on values, so make sure you're paying attention! If you do not choose some fruit and/or bread by the time a round ends, you will get none of the fruit and no bread and thus earn no value that round.

Your earnings in points for each round is the value you get from your fruit and/or bread plus the bonuses you earn. Your earnings for your decisions are calculated as: the sum of your earnings in each round times the conversion factor of 0.001 dollars per point. After your decisions, you will answer a survey that will take a few minutes, and you'll receive another $0.50 for your completion of the survey.

For example, if you earned 4,000 points across all of the rounds, that would give you 4,000*0.001 = $4 for your decisions, plus $0.50 for the survey, for a total payment of $4.50. (Your payment will be rounded to the nearest cent if necessary.)

Decision

Time left to complete this page: **0:44**

This is round 1. You will play this game for 10 rounds in total.

Instructions reminder: spend all your money buying fruits and/or bread for 1 franc each. Bread always gives you 50 points points per franc; you start out with guesses about how many points per franc you get for each fruit. Your starting guesses before you try the good are your actual values from the fruits plus or minus random numbers. At the end of the round you'll learn your earnings from each fruit you buy at least 1 of. Your values for those fruits will be updated for you to see in future rounds. Your payment for this experiment depends on the values you earn in each round!

You have 3 francs to spend.

Choose how many of each of the foods you would like to buy:

Food name	Value or guess	Guess?	How much would you like to buy?
Merooki	48 points	guess	Merooki: 0.0
Bread	50 points		Bread: 0.0

Next

Decision Round Report

Time left to complete this page: **0:16**

This was round 1.

You earned 0 points.

Here are your updated values for all of the fruits. If the word "guess" appears, the value is your starting guess. If it does not appear, this is a value you've learned in this or a past round.

Food name	Guess?	Value
Frutana	guess	60 points
Jojofruit	guess	90 points
Banello	guess	69 points
Niblunda	guess	80 points

Decision Round Report

Time left to complete this page: **0:17**

This was round 2. You bought:

3.0 Niblunda

You earned 165 points.

Here are your updated values for all of the fruits. If the word "guess" appears, the value is your starting guess. If it does not appear, this is a value you've learned in this or a past round.

Food name	Guess?	Value
Frutana	guess	60 points
Jojofruit	guess	90 points
Banello	guess	69 points

This was round 8. You bought:

3.0 Banello

You earned 225 points.

Here are your updated values for all of the fruits. If the word "guess" appears, the value is your starting guess. If it does not appear, this is a value you've learned in this or a past round.

Food name	Guess?	Value
Frutana	guess	60 points
Jojofruit		69 points
Banello		75 points
Niblunda		55 points
Danutia	guess	40 points
Yegrevy	guess	45 points
Merooki		67 points
Oggerydot	guess	65 points
Zellitan	guess	44 points
Valavoo	guess	54 points
Bread		50 points

Next

End of Decisions Report:

Here are your earnings from the decision rounds:

Round	Food Earnings
1	0 points
2	165 points
3	201 points
4	234 points
5	225 points
6	156 points
7	207 points
8	225 points
9	234 points
10	201 points

Your total food earnings are the sum of your earnings in those rounds, which is 1848 points. Since you earn 0.001 dollars per point, this means your food earnings are worth \$1.85. Your total earnings are that plus \$0.50 for filling out the survey that you are about to start, or a total of \$2.35.

Click Next to start the survey!

Next

Questionnaire page 1

Please answer all of the questions below. Your answers will not affect your payment but will help us understand our results.

1. What do you think the experimenters will learn from this experiment?

2. Imagine you have 100 francs. If apples and bread each cost 1 franc, you know that apples earn you 5 points per franc and bread earns you 4 points per franc, how much would you buy of each if you want to earn as many points as possible?

Apples:

Bread:

3. Imagine we are throwing a five-sided die (with sides numbered 1, 2, 3, 4, and 5) 50 times. On average, out of these 50 throws how many times would this five-sided die show an odd number (1, 3 or 5)?

4. Out of 1,000 people in a small town 500 are members of a choir. Out of these 500 members in the choir 100 are men. Out of the 500 inhabitants that are not in the choir 300 are men. What is the probability that a randomly drawn man is a member of the choir?

   (Please indicate the probability in percent):

5. Imagine we are throwing a loaded die (6 sides). The probability that the die shows a 6 is twice as high as the probability of each

6. In a forest 20% of mushrooms are red, 50% brown and 30% white. A red mushroom is poisonous with a probability of 20%. A mushroom that is not red is poisonous with probability of 5%. What is the probability that a poisonous mushroom in the forest is red?

   (Please indicate the probability in percent):

Next

Questionnaire page 2

Please answer all of the questions below. Your answers will not affect your payment but will help us understand our results.

7. What is your gender?

   --------- ▾

8. What country were you born in?

   --------- ▾

9. What is your age?

Next