UNIVERSITY OF CALIFORNIA
SANTA CRUZ

3 EXPERIMENTS IN 3 COMPLETELY DIFFERENT THINGS

A dissertation submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

ECONOMICS

by

Ciril Bosch-Rosa

June 2013

The Dissertation of Ciril Bosch-Rosa
is approved:

Professor Dan Friedman, Chair

Professor Carl Wlash

Professor Ryan Oprea

Tyrus Miller
Vice Provost and Dean of Graduate Studies

Contents

List of Tables

List of Figures

Abstract

3 experiments in 3 completely different topics

By Ciril Bosch-Rosa

Abstract 1: We present a three-player game in which a decision-maker, in the role of referee, accepts or rejects the offer made by a proposer to a passive receiver. If the offer is accepted, the split takes place as suggested, if rejected both proposer and passive receiver get $0. The payoff of the decision-maker, on the other hand, will be the treatment variable. Our results show a decision-maker that ignores his payoffs, but that is so concerned about equality among other players rejecting both selfish and generous offers. When we introduce a cost to rejecting proposals, we are able to show that inequality aversion is the only reason behind rejections.

Abstract 2: When should a necessary inconvenience be introduced gradually, and when should it be imposed all at once? The question is crucial to web content providers. In a setting where people eventually fully adapt to changes, the answer depends on the shape of the "survivor curve" S(x), which represents the fraction of a user population willing to tolerate inconveniences of size x.We report a laboratory experiment that estimates the shape of survivor curves in several different settings. Our key finding is that web content providers will generally find it profitable to introduce inconveniences gradually over time.

Abstract 3: There is consensus that the recent financial crisis revolved around a crash of the short-term credit market. Yet there is no agreement around the necessary policies to prevent another credit freeze. In this experiment we test the effects that contract length has on the market-wide supply of short-term credit. Our main result is that, while credit markets with shorter maturities are less prone to freezes, the optimal policy should be state-dependent, favoring long contracts when the economy is in good shape, and allowing for short-term contracts when the economy is in a recession. We also report runs on firms with strong fundamentals, and rich learning dynamics, with a text-book bubble and crash pattern in the short-term credit market.

Acknowledgments

Gracies a vosaltres y Gracias Dr. Arch. Cocho

I also want to thank my advisers Dan Friedman and Ryan Oprea, as well as Carl Walsh for guidance and invaluable advice. I also want to thank (in a random order) James Pettit, Johannes Müller-Trede, Thomas Campbell, Antoni Bosch-Domenech, Joaquim Silvestre, Pablo Branas, Teresa Garcia Munoz, Curtis Kephaert, Gabriela Rubio, and Alicia from the burrito truck near the department.

Chapter 1

Introduction

> "Essentially all models are wrong, but some are useful" Box and Draper (1987).

The deep implications from Box and Draper (1987) might elude the casual reader, but if one stops to think about it, its deep implications will strike: “Of course, all models are wrong...but we don’t really care!”. And, indeed that is the case; it doesn’t matter if we cannot solve a 3-body celestial problem, if we can land the Mars Rover, and we don’t really care if we cannot predict the exact path of every neutron in a nuclear reactor as long as our probabilistic model allows us to provide safe nuclear energy.

Any theory can be expanded by adding new free parameters, and this will give us better predictions. But is this our objective? Of course we prefer a more accurate prediction to a less precise one, but is the gain in precision worth the cost in complication? Was that new axiom worth it? The question should not be ”does adding this axiom make my model better?” but rather “was introducing this axiom worth the extra complication?” Was the marginal cost of smaller or equal than the marginal profit?

What Box and Draper are telling us is that the goal of science (and its models) is not to try and axiomize every single element of the physical medium, but rather to strike an equilibrium between complexity and generalizability[1]. Much like in a constrained utility maximization problem, the beauty of scientific modeling relies on balancing the extra precision added by a new axiom, against the extra complication of the model. As Karl Popper posed it: "Science may be described as the art of over-simplification, the art of discerning what we might with advantage omit." The question is not how inaccurate can a model be, yet still be useful.

This balancing of the pros and cons can only be measured through controlled lab experiments, where scientists, very much like cooks tasting their own dish, can asses whether a model needs some improving, or rather can do without some of the axioms.

And, while it is through *ceteris paribus* comparisons that experiments prove themselves most useful, one chapter of this dissertation will use lab techniques not to test a theory, but rather to estimate a parameter value. In this case lab techniques are necessary because only in a controlled environment can one acquire precise measures with no external interferences. While parameter measurement is certainly less ambitious than model refining, it is essential for any application of the theoretical models. For example, it is by as knowing that $9.80665 m/s^2$ is the standard acceleration for free fall (in the earth) that gravitational models can be used to put rockets in a sub-orbital flight, with a less precise measure it would be impossible to successfully use rockets.

In this doctoral thesis I will present three independent pieces of research that use experimental techniques applied to economic questions. In the first and third chapter I will test the predictions made by the standard inequality aversion theory used in the social preferences literature. In the second chapter (which is coauthored with Christina

[1] For example, Galileo's falling bodies postulates do not take into account the friction of the medium (air), even thought it was obvious for the author that it did have an effect on the acceleration of the bodies in motion

Aperjis, Bernardo Huberman, and Daniel Friedman) I will be estimating an essential parameter required by the theory that we develop in this same chapter. Finally in the third chapter I will look at the predictions made by a model of financial panics and will use the results to suggests potential policies to stabilize short-term credit markets.

The objective of this thesis is not only to be innovative in the methodology, but it also makes an effort to try and open new lines of research for experimental economics. While the first chapter is a classic topic in economics about which much has been written, it is in the second and third chapter where not only do I introduce completely new methodological instruments, but most importantly, I ask new questions, expanding the frontier of experimental research and showing that experimental economics has the flexibility to tackle all sorts of new economic problems. In all cases the models are incomplete, but some of them might come in handy.

Chapter 2

A Tale of Two Tails: rejection patterns of extreme offers in a three-player game

> "How selfish soever man may be supposed, there are evidently some principles in his nature, which interest him in the fortune of others, and render their happiness necessary to him, though he derives nothing from it except the pleasure of seeing it." The Theory of Moral Sentiments, Adam Smith (1759)

2.1 Introduction

The literature on other regarding preferences has come up with many reasons to explain the counter-intuitive rejections observed in ultimatum games[1], yet no experiment has ever tried to observe the preferences of a neutral third party with no stakes in the game to see what are the "primal" reasons behind these rejections. Using a novel three-player ultimatum game structure that separates the decision-maker's choices from his final payoffs, we are able to map the decision-maker's neutral preferences over the whole span of possible splits in the game. The results show a decision-maker that ignores his relative payoffs when making decisions, but that is very concerned about the equality of splits between the two other players. In fact, this concern is so deep that rejections are of both selfish and generous [2] offers. Finally, we report that while decision-makers show some worries over the selfish behavior of proposers, once we introduce a cost to rejecting offers, these worries disappear completely, with rejections being driven only by inequality concerns. The main contributions of this paper are, thus, the design of a game where the decision-maker has no strategic or monetary concerns, the mapping of the preferences of this "disinterested" decision-maker (which frequently rejects generous proposals), and finally showing that selfish intentions of proposers are hardly of second order of importance to decision-makers.

2.2 Literature Review

Three-player ultimatum games have been largely studied, being responsible for

[1] These reasons range from inequality aversion (Bolton and Ockenfelds (2000) or Fehr and Schmidt (1999)), to punishment of selfish intentions (Blount (1995)), or Rawlsian preferences (Charness and Rabin (2002) and Engelman and Strobel (2004)), and even to the need for signaling disconformity (Xiao and Houser (2005)).

[2] From now on we will consider any offer of more than$5 to be "generous".

some key insights in the much written about ultimatum game literature. In Knez and Camerer (1995), a proposer makes a simultaneous offer to two independent responders who can accept or reject proposals conditional on the offer made to the other receiver. The results show that receivers accept offers depending on their relative standing to the third participant, that is, responders are not willing to get offered less than their counterpart. In Güth and Van Damme (1998), a proposer splits the pie with a decision-maker and a passive "dummy" player who plays no role in the game; if the offer is accepted by the decision-maker, then the split goes as suggested, if rejected, then everyone receives zero. The result is that both proposer and responder end up ignoring the presence of the dummy player and split the pie between themselves. Finally, Kagel and Wolfe (2001) present us with a setup identical to Güth and Van Damme (1998) except that now, if the offer is rejected, the dummy player gets a consolation prize. As in Güth and Van Damme (1998), but against inequality aversion theories, the dummy seems to play no role for decision-makers even when he gets a high consolation prize. Another strand of literature that has a bearing in our experiment is third-party punishment. Fehr et al. (2005) report that third-party sanctions are not used to reduce inequality, but rather as retaliation for selfish actions. Yet, Leibbrandt and Lopez-Perez (2008) use a within-subject design to conclude that both second and third party punishments are driven by outcomes not intentions. Interestingly, and against Fehr and Fischbacher (2004), Leibbrandt and Lopez-Perez (2008) find that second-party punishment is not significantly higher than third-party punishment. More recently, Falk et al. (2008) have revisited the subject suggesting that while inequality has some effect on punishment, intentions are the main reason behind most punitive actions. Our conclusions are in stark contrast with this latter result. While we are not the first to report rejections of generous offers, we are the first to do so in a lab experiment (previous reports were field experiments with subjects from rural regions of Russia). Furthermore, these previous results had always been dismissed as an anomaly. For example, Bahry and Wilson (2005) report an "inverted-U" pattern [3] when comparing ultimatum game results across old Soviet Union regions, but they dismiss it as a result of Soviet education. The second paper in which the inverted-U pattern is mentioned is Güth et al. (2007), in which the authors gathered ultimatum game data through newspaper publications. In this case the results are not reported, but only informally mentioned due to the insignificant number of observations.

Finally, there has been some controversy about the validity of the strategy method, a technique which we use in our experiment. Brandts and Charness (2011) is a good survey on the subject and supports the use of the strategy method. In fact, if we had used a direct method instead of the strategy method, the inverted-U results might have been even more prominent as Brandts and Charness (2011) report that punishment rates are lower if the strategy method is used. Further, Brandts and Charness (2011) claim that "in no case do we find that a treatment effect found with the strategy method is not observed in the direct-response method". See also Brandts and Charness (2003) for more information on the matter.

2.3 Experimental Design

The experiment has two different game structures. The first one is a three-player ultimatum game (3UG) which is at the center of the paper, the second a two-player ultimatum game (2UG) which we use to show that decision-makers take seriously the possibility of generous offers being made.

The three-player ultimatum game (3UG) has a proposer (A) making an offer to a

[3] As we will see in our experiment, when plotting the ratio of accepted offers, if there are rejections of both generous and selfish offers, the graph looks like an inverted U.

dummy player (C) on how to split \$10. Meanwhile, the decision-makers (B), without knowing the actual proposal, fills a strategy profile (Figure 2.1) accepting or rejecting all potential offers from A to C. If the offer is accepted, then the split goes as suggested by A; if rejected, then both A and C get nothing for the round. B's payoffs are our treatment variables, which we divide into two groups of treatments. The first one is the "costless-rejection" group which has 3 treatments:

- Low (L): B gets paid \$3 for his decisions, whatever the outcome of the game
- Normal (N): B gets paid \$5 for his decisions, whatever the outcome of the game.
- High (H): B gets paid \$12 for his decisions, whatever the outcome of the game

The reason for having three treatments is to test if the decision-maker is payoff "neutral" or whether his relative payments affects his choices. For example; rejections of generous offers could be justified by a need to prevent anyone from getting a higher payoff than the receiver (i. e., rejecting because of disadvantageous inequality). If that were the case, then we should not observe rejections of generous offers in the H treatment (or at least significantly less than in the L treatment). On the other hand, if rejection patterns are not significantly different across treatments, then it means that decision-makers do not take into account their payoffs when making decision. This would imply that the results of the 3UG game are the decision-makers "pure" set of preferences over the splits of A and C; not only does the decision-maker not have any strategic concerns in the game, but he is not driven by his relative monetary standing, just his pure set of preferences.

Figure 2.1: Decision-Maker Screenshot

The second group will be what we call the “costly-rejection” treatments and its purpose is to test the robustness of our findings in the first family. The ”costly-rejection” group has two treatments:

- Low (L-1) : B gets paid $3 if A’s offer is accepted and $2 if rejected.
- High (H-1) : B gets paid $12 if A’s offer is accepted and $11 if rejected.

Figure 2.2 graphically lays out the general structure of the 3UG game for both families.

Figure 2.2: General Structure of the Game

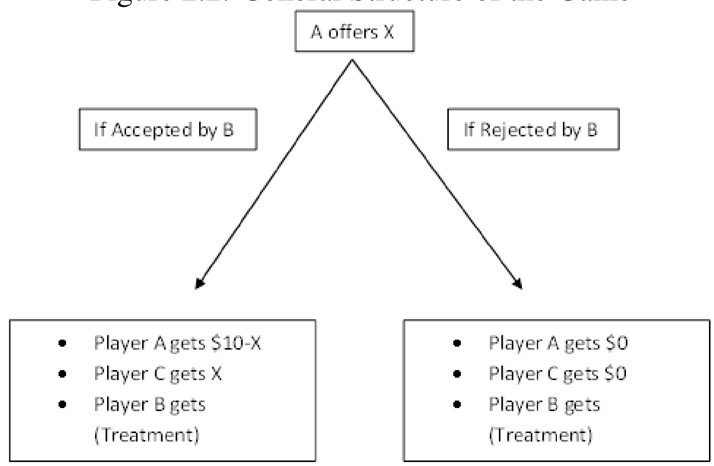

2.3.1 2UG

In the 2UG game, we keep the 3-player group design, but now A makes two independent offers on how to split $10; one offer to B, the other to C. As in the 3UG case, we will use the strategy method to elicit B and C's preferences over offers made to them. So, if B (C) rejects the offer that A made to him, then B (C) gets $0 for the round. If, instead B (C) accepts the offer, then the split goes as suggested by A. A's payoff is randomly chosen from one of the two different outcomes; if the selected game turns out to be a rejection, then A gets $0 for the round, if an acceptance, then A gets his part of the proposal. The purpose of the randomization of payoffs is to prevent portfolio effects, and to make payoffs fairer across subject types. The 2UG game was designed to verify if decision-makers took seriously the possibility of "hyper-generous" offers, and to validate our subject pool. The results of this game will be crucial to justify the credibility of some of our most interesting 3UG results.

2.4 Implementation

The experiment was run with a total of 237 undergraduates from both the Universitat Pompeu Fabra (UPF) in Barcelona, and the University of California Santa Cruz (UCSC), in Santa Cruz. Each session had three rounds and on average lasted 25 minutes. The mean earnings at UCSC were of $4.5 and at UPF of d’4.35 plus a show-up fee ($5 andd’3[4]) that was announced only at the end of the experiment[5].

Subjects were recruited through the ORSEE systems of each university, and were required not to have any previous experience in bargaining games. In total 15 sessions were run, UCSC sessions had 12 subjects[6] and UPF sessions 18 subjects[7].

As subjects arrived to the lab, they were seated randomly in front of a terminal and the initial instructions were read aloud. In these instructions we announced that:

1. Instructions for each round would be read immediately before each round started [8].

2. Each subject would be assigned a player type (A, B or C) which they would keep through the experiment.

3. Each round, subjects would be randomly assigned to a different group of three players (one of each type).

4. Only one of the rounds, randomly chosen by the computer, would be chosen for

[4] From now on, we will use the dollar sign to include both euros and dollars.

[5]While most subjects are aware of the rule of a “show-up fee” not announcing it until the end of the experiment adds pressure to the decision-makers would their decisions result in a rejection.

[6]Except 3 sessions that had 9 subjects.

[7]Except 2 sessions that had 12 subjects

[8]From experience, we prefer to read several times small amount of instructions rather than going over all instructions at the beginning of the session since subjects then get distracted. By breaking instructions into small concise parts we increase the likelihood that subjects are paying attention and, consequently, that they know what is expected of them in each round.

the final payoffs.

5. No feedback would be given until the end of the session[9], when they would be informed of the actions of subjects in their group for each round, as well as the round selected for the final payoffs.

Each session was composed of two 3UG rounds and one 2UG[10].

2.4.1 2UG

We summarize all of B subject's observations in Figure 2.3. In it we present the percentage of decision-makers accepting each potential offer from A to C (e.g. almost 60% of B subjects accept a hypothetical offer of $3 while only 30% accept one of 1). The acceptance results are slightly higher than those reported in the literature (Camerer and Thaler (1995)), but still within the range of what would be expected. The average offer was of $3.59, which is also what would be expected in an experiment like this. These results validate both our subject pool and the software interface, but most importantly, they show that decision-makers act consistently[11] when deciding about hyper-generous offers (i.e., subjects do not randomize or "experiment" within this range of offers). We take this as an indication that decision-makers take seriously the possibility of a generous offer.

[9]This was done to minimize learning effects and have results of a "one-shot game" in each round.

[10]Please check Appendix A for details on the ordering of treatments and number of observations in each session.

[11]Three subjects that rejected offers of $8 or more yet accepted all smaller offers. We believe that these subjects misunderstood the interface and were trying to reject offers smaller than $2.

Figure 2.3: Acceptances of 2UG

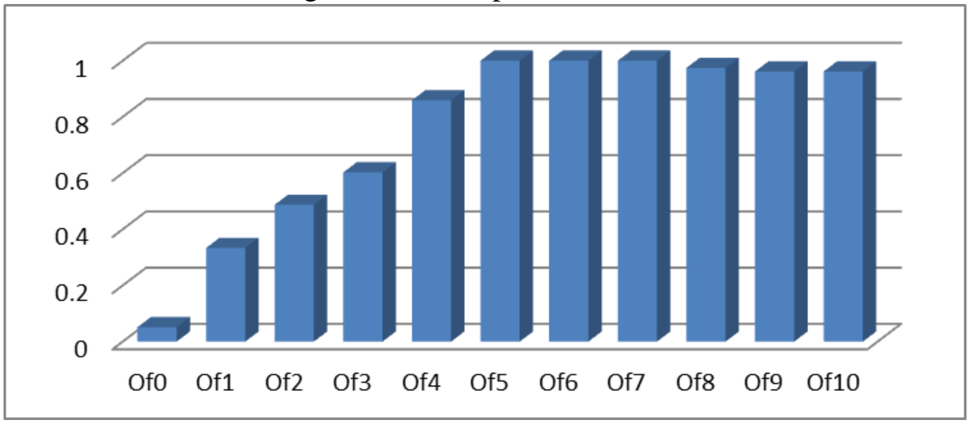

2.4.2 Free-Rejection Treatment

Figure 2.4 presents the results of the free-rejection group of treatments. In the upper-left corner we see the baseline treatment, N, and in a clockwise order the comparison between N and L, L and H, and finally between N and H in the lower-left quadrant. Two things stand out immediately from these graphs. First, decision-makers reject both generous and selfish offers. In fact, if an offer is generous, the more generous it is, the less likely that it will be accepted. It is this pattern of behavior that gives us the inverted-U shape that Bahry and Wilson (2005) first identified in their field experiments. The second striking feature is that all treatments seem to have identical effects, whether we pay $3 or $12 to the decision maker, the behavior is the same. In fact, if we run a two-sided Fisher test comparing the aggregated number of acceptances for each potential offer, we find no statistically significant differences across treatments (Table 2.1).

Table 2.1: Two-sided Fisher p-values comparing total acceptances across treatments.

P-values	$0	$1	$2	$3	$4	$5	$6	$7	$8	$9	$10
L=N	1	.775	.596	1	1	.141	.55	1	1	.81	
H=N	.355	.280	.202	.808	.604	.250	.759	.792	.226	.469	.636
L=H	.329	227	.089*	.789	.768	1	1	.768	.269	.412	.787

Further, if we use a Wilcoxon matched-pairs signed-rank test to compare the number of accepted offers in each treatments a subject has participated in, then we see that the number of accepted offers is not statistically different between the N and L treatments (p = 0.375) nor among the N and H (p = 0.161)[12].

Figure 2.4: Acceptances Free 3UG

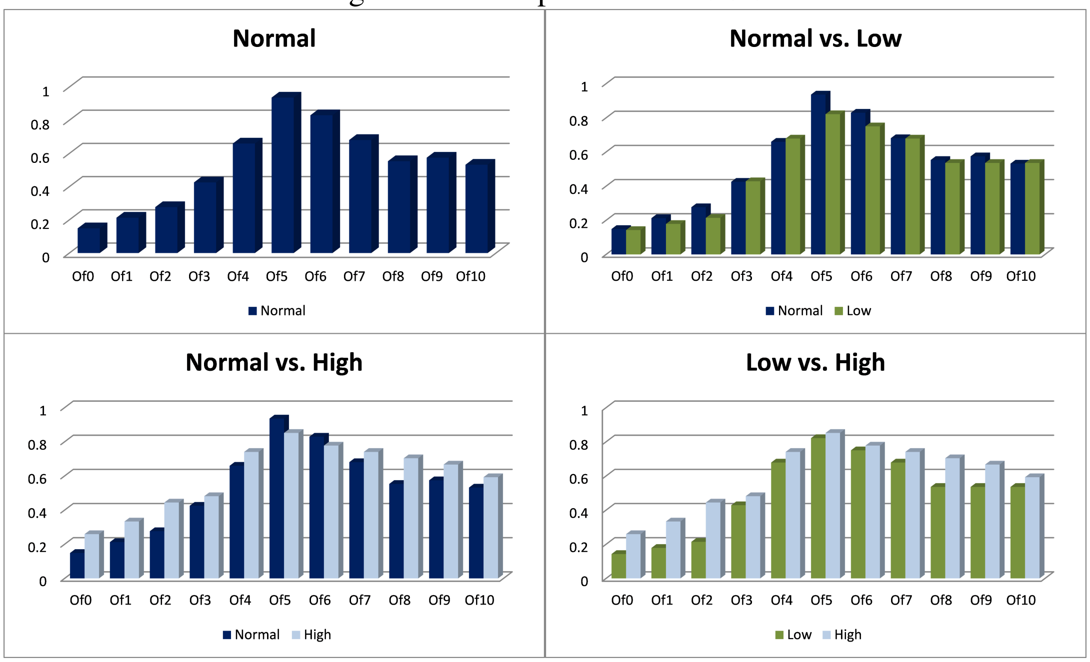

Finally, we run a regression of total accepted offers (Total) on dummies for location (Where), order (First), and treatment (High and Low). The results are shown in Table 2.2, in the first two columns we compare H directly to L, in the third and fourth we compare High and Low to the baseline N. The results show that payoffs and ordering[13] have no effect on the number of accepted offers, and neither does location

[12]On the other hand, the test becomes somewhat more significant when comparing L and H (p = 0.0825), probably because the number of subjects participating in both H and L is extremely low (n = 4). See appendix B for a lengthier discussion on this question.

[13] Column 2 shows some minimal order effects. We attribute these to the lack of first round H treatment observations. See Appendix B.

Table 2.2: Regression of total accepted offers by subject and treatment.

	(1) Total	(2) Total	(3) Total	(4) Total
Low	-1.093	-1.327	-0.330	0.165
	(0.848)	(0.817)	(0.666)	(0.796)
First		1.707*		1.008
		(0.947)		(0.717)
Where		-0.101		-0.263
		(1.281)		(0.979)
High			0.763	1.399
			(0.805)	(0.978)
cons	6.593***	6.318***	5.830***	5.114***
	(0.747)	(0.637)	(0.461)	(0.816)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(all of these results are later confirmed in Table 2.3). Therefore, in the free-rejection treatments, decision-makers appear not to take their own payoff as a reference when making decisions. It seems that in our three-player game structure, decision-makers assign to themselves the "dummy" position that Kagel and Wolfe (2001) and Güth and Van Damme (1998) had previously reported in three player ultimatum games. This is an important result, as it shows that the inverted-U pattern is truly a mapping of preferences from a decision-maker only concerned about the splits between A and C and not his own payoffs.

- **Result 1:** *Under the "costless-rejection" treatments, decision-makers do not take their own payoff as a reference when making decisions.*

To better analyze the results of the 3UG game we define "absolute inequality" as the absolute value of the difference between A and C's payoff, and label all offers to the left of \$5 as the Left-Hand-Tail (LHT), and all offers to the right of \$5 as the Right-Hand-Tail (RHT). We then run a Spearman rank correlation test (Appendix C) and show that as we move away from the even split (i.e. when absolute inequality increases), acceptance rates decrease in both directions (i.e. both in the RHT and the LHT). To have a more accurate idea of how absolute inequality affects the probability of rejection, we run a linear probability model[14] (Table 2.3). In it, the binary accept/reject outcome is the dependent variable and we have dummies for ordering (First), treatment (High, Low), location (Where), as well as dummies for distance that are coded with both distance to the even split and tail (Left or Right) they are located in. For example, dist3l is the dummy for the $2 offer (which is 3 dollars to the left of $5) and dist2r is the dummy for an offer of $7 (which is 2 dollars to the right of $5). Column 5 of Table 2.3 has the full specification of the regression and shows that all dummies for distance are not only negative, but also highly significant. Moreover, if we look at the coefficients for the distance dummies, the further away an offer is from $5 the lower is its probability of being accepted; the higher absolute inequality, the lower the probability of being accepted. This relationship is monotonic in both tails [15] ranging from an 8% lower probability of acceptance for an offer of $6 (dist1r) to a 33.3% lower probability of acceptance for an offer of $10 (dist5r)!

- **Result 2:** *In both tails, the greater the absolute inequality, the lower the probability of the proposal being accepted.*

On the other hand, from Figure 2.4 we can see that the inverted-U does not seem totally symmetric as for the same level of inequality LHT seem to be less likely of being

[14] With clustered errors at the individual level.

[15] Strictly monotonic in the LHT and weakly in the RHT.

Table 2.3: Linear Probability model of Accepted Offers.

	(1) Accept	(2) Accept	(3) Accept	(4) Accept	(5) Accept
Low	-0.0300	0.0150	0.0150	0.0150	0.0150
	(0.0547)	(0.0671)	(0.0673)	(0.0673)	(0.0674)
High	0.0693	0.127	0.127	0.127	0.127
	(0.0617)	(0.0803)	(0.0805)	(0.0805)	(0.0806)
First		0.917	0.917	0.917	0.917
		(0.0581)	(0.0582)	(0.0581)	(0.0584)
Where		-0.0239	-0.0239	-0.0239	-0.0239
		(0.0895)	(0.0897)	(0.0897)	(0.0899)
Dist1l			0.00327		-0.196
			(0.0501)		(0.0469)
Dist2l			-0.242***		-0.441***
			(0.0587)		(0.0625)
Dist3l			-0.379***		-0.578***
			(0.0558)		(0.0636)
Dist4l			-0.448***		-0.647***
			(0.0555)		(0.0603)
Dist5l			-0.507***		-0.706***
			(0.0578)		(0.0849)
Dist1r				0.340***	-0.088***
				(0.0447)	(0.0306)
Dist2r				0.242**	-0.186***
				(0.0493)	(0.0662)
Dist3r				0.134**	-0.294***
				(0.0530)	(0.0531)
Dist4r				0.134**	-0.294***
				(0.0536)	(0.0584)
Dist5r				0.0948*	-0.333***
				(0.0565)	(0.0625)
Cons	0.530***	0.465***	0.608***	0.379***	0.807***
	(0.0415)	(0.0739)	(0.0803)	(0.0708)	(0.0690)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

accepted. After all, in Table 2.3, the dummy for an offer of \$3 (absolute inequality = 4) has a coefficient of -0.44, while an offer of \$7 (with same absolute inequality) has a coefficient of -0.19. We attribute this to the concern for intentions of the decision-maker, as he seems more tolerant of inequality if this is the result of a generous offer. To check if intentions are significant, we run a linear probability model for each individual treatment, and compare the coefficients of those offers with same absolute inequality through a Wald Test (Table 2.4). The result shows that the tails are asymmetric, that is;

Table 2.4: P-values of Wald test for equality in within treatment regression coefficients.

Treatment	dist1l=dist1r	dist2l=dist2r	dist3l=dist3r	dist4l=dist4r	dist5l=dist5r
L	0.3357	0.0187***	0.0052***	0.0026***	0.0013***
H	0.5813	0.0066***	0.0186***	0.0016***	0.0016***
N	0.0107**	0.0021***	0.0022****	0.000***	0.0000***

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.5: Two-sided Fisher Test.

Treatment	$4=$6	$3=$7	$2=$8	$1=$9	$0=$10
L	0.768	0.106	0.026**	0.011**	0.004***
H	1	0.093*	0.098*	0.029*	0.027**
N	0.048**	0.011**	0.006****	0.001***	0.0000***

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

for same absolute inequality, the coefficients are significantly different.

For completeness, in Table 2.5 we show the results of a Two-sided Fisher test comparing the number of accepted offers for same absolute inequality proposals across treatments.

- **Result 3:** *In costless-rejection treatments, decision-makers are less willing to tolerate inequality when this is the result of a selfish offer.*

The three results presented above offer a picture of a decision-maker unconcerned about his own payoff, yet showing such inequality aversion that he is willing to leave both A and C with a $0 payoff rather than accepting an unequal offer even if it is a generous one! In fact, if an offer is generous (i.e. in the RHT), then, the more generous it is the less likely it is to be accepted. Corollary 3 on the other hand shows that inequality aversion is not the only reason behind the inverted-U pattern that we observe, as for the same level of absolute inequality those offers in the RHT are more likely to be accepted than those in the LHT. We interpret this as intentions playing a role, but behind stronger inequality concerns. Since the outcomes of our first batch of experiments were so striking, we decided to design a new group of treatments in which the decision-maker has to pay $1 if the game ends in a rejection. We label this the "costly-rejection" group, and it has two treatments, H-1 and L-1. In the first (second) one the decision-makers is paid $12 ($3) if the outcome is an acceptance and $11 ($2) if it's a rejection. The purpose of this new treatments is to test whether decision-makers are still willing to reject generous offers when a cost is introduced. In particular, we are interested in whether decision-makers continue to ignore their payoffs (and act similarly) even when the costs to rejecting are so different. Notice that in both treatments the cost of rejecting is the same ($1), but in the L-1 case this represents 1/3 of the payoff, while in H-1 it is only 1/12 of the payoff.

2.4.3 Costly Rejection Treatment

In Figure 2.5 we present the results of the costly-rejection treatments and compare them to their costless-rejection counterpart, and the N baseline. The most striking feature is that we still observe rejections on both tails[16] under costly treatments; the inverted-U pattern of acceptances is still there even under the "costly-rejection" treatment. As we can see in Table 2.6, a linear probability model with dummies for distance to the fair split continues to show a monotonic negative correlation between inequality and probability of acceptance in both tails. So, decision-makers continue to reject both selfish and hyper-fair offers even when these have a high cost (L-1).

[16] See Appendix B for Spearman Correlation results

Figure 2.5: Acceptance rates of L-1 and H-1 vs. Normal, Low, and High Treatments

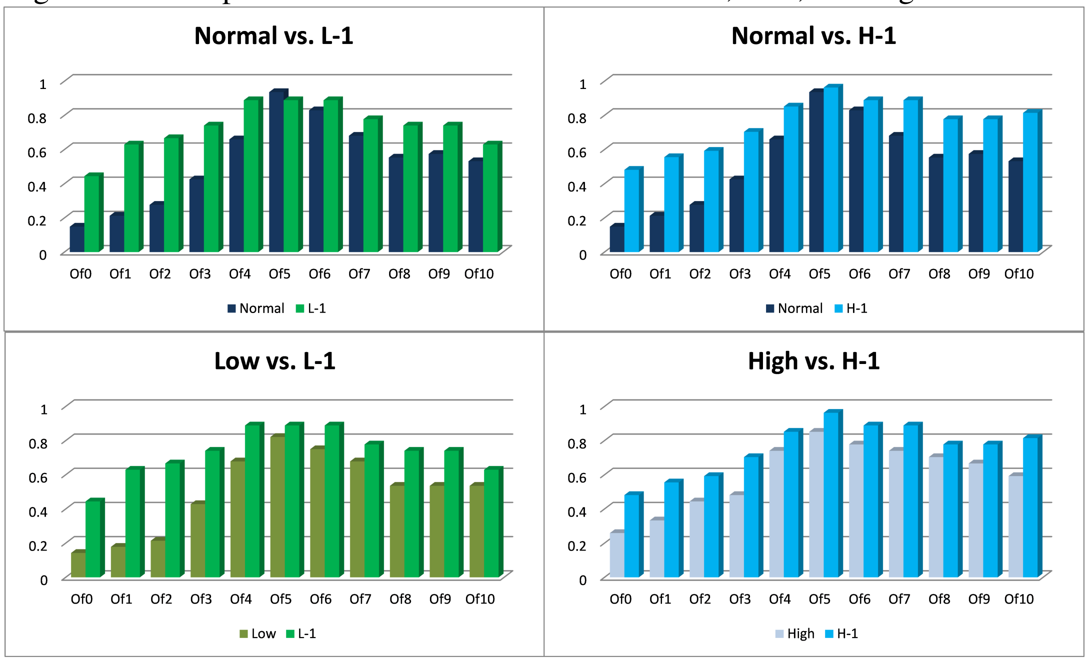

Another result that jumps out from Figure 5 is the similarity between H-1 and L-1 (Figure 2.6). It seems like even if the relative costs of rejecting are so wide apart, much like in the costless case, decision-makers behave in an identical manner across all costly treatments.

Figure 2.6: Acceptance Rate of L-1 and H-1

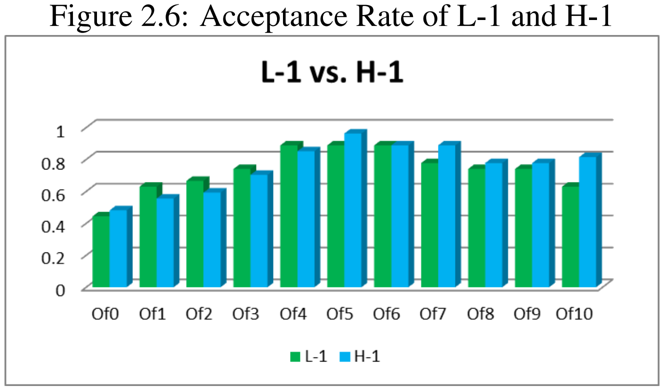

As in the costless treatment we run a Wilcoxon matched-pairs sign-rank test comparing the number of offers that each subject accepts in each treatment she participated

Table 2.6: One-sided Fisher P-values comparing total acceptances per treatment.

	\$0	\$1	\$2	\$3	\$4	\$5	\$6	\$7	\$8	\$9	\$10
L	0.01**	0.01**	0.01**	0.01**	0.05*	0.37	0.16	0.30	0.09*	0.09*	0.33
H	0.07*	0.08*	0.20	0.08*	0.25	0.17	0.23	0.14	0.37	0.27	0.06*

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

in, finding no statistically significant difference (p = 0.6172). Additionally, both the linear probability model of Table 2.6 and a Two-sided Fisher test, confirm that there exists no significant difference between treatments (Appendix D).

- **Result 4:** *In the costly-rejection treatments, decision-makers do not take their own payoffs as a reference when making decisions.*

Where we do see a difference is between both families; if we run a regression on total accepted offers comparing H to H-1 and L to L-1. The tests show a significant difference across treatment dummies (p= 0.002 and p = 0.000 respectively).

To be more precise about where the treatments differ we run a One-sided Fisher test and observe that the differences are mostly in the LHT (Table 2.7). As we can see from both Table 2.7 and Figure 2.5, it looks like once we introduce a cost to rejecting offers, concerns over the selfish proposals seem to fade away and only absolute inequality is driving rejection.

- **Resut 5:** *Costly-rejection treatments do not present the same pattern of acceptances than costless-rejection ones, as LHT acceptance rates are significantly higher in the former than in the latter.*

In order to prove that concerns for selfishness disappear once we introduce a \$1 penalty

Table 2.7: Linear Probability model of Accepted Offers.

	(1) Accept	(2) Accept	(3) Accept	(4) Accept	(5) Accept
High1	0.0236	0.0289	0.0289	0.0289	0.02889
	(0.0615)	(0.0671)	(0.0674)	(0.0674)	(0.0677)
First	0.0236	0.0289	0.0289	0.0289	0.02889
	(0.0615)	(0.0671)	(0.0674)	(0.0674)	(0.0677)
Dist1l			0.0556		-0.0556
			(0.0417)		(0.0412)
Dist2l			-0.0926***		-0.204***
			(0.0391)		(0.0675)
Dist3l			-0.185***		-0.296***
			(0.643)		(0.0820)
Dist4l			-0.222***		-0.333***
			(0.0595)		(0.0809)
Dist5l			-0.352***		-0.463***
			(0.0697)		(0.0849)
Dist1r				0.188***	-0.0370
				(0.0574)	(0.0461)
Dist2r				0.133**	-0.0926
				(0.0493)	(0.0662)
Dist3r				0.0585	-0.167**
				(0.0500)	(0.0763)
Dist4r				0.0585	-0.167**
				(0.0462)	(0.0660)
Dist5r				0.0216	-0.204***
				(0.0487)	(0.0675)
Cons	0.731***	0.713***	0.786***	0.672***	0.897***
	(0.0608)	(0.0840)	(0.0809)	(0.0893)	(0.0860)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.8: P-values of Wald test

Treatment	dist1l=dist1r	dist2l=dist2r	dist3l=dist3r	dist4l=dist4r	dist5l=dist5r
L-1	1.00	0.7536	0.5302	0.3466	0.1175
H-1	0.7410	0.0991^{*}	0.0991^{*}	0.0481^{**}	0.0032^{***}

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.9: Two-Sided Fisher P-values.

Treatment	$4=$6	$3=$7	$2=$8	$1=$9	$0=$10
L-1	1.000	1.000	0.766	0.559	0.275
H-1	1.000	0.175	0.241	0.148	0.021^{**}

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

to rejecting, we compare both tails of the acceptance distribution. Running a linear probability model for each treatment in the group we compare the coefficients for those dummies with the same level of absolute inequality through a Wald test. The results show a symmetric L-1, but a slightly unbalanced H-1 (Table 2.7). On the other hand, a Two-sided Fisher test shows both treatments as symmetric (Table 2.8). We consider these results to be a direct result of the weight that decision-makers put into the intentions of proposers; even the smallest cost (1/12 of the total payoff) has such an effect that the inverted-U pattern is almost perfectly symmetric as inequality aversion still persists as the main driver behind rejections.

- **Result 6:** *Under costly-rejection treatments, the intentions of the proposer play a minor role in the acceptance pattern of decision-makers, disappearing completely in the L-1 treatment.*

We introduced the costly-rejection treatments to see if decision-makers still would reject generous offers even if this entailed a cost. The results show that not only does the decision-maker continue to reject offers in this group of treatments, but that he behaves the same across treatment group even under very asymmetric payoffs. But the surprise comes from the tail where decision-makers decide to stop rejecting; it is not on the RHT, but rather in the LHT. It turns out that when a cost to reject is introduced in the 3UG game decision-makers stop putting an important weight to intentions and only reject offers based on absolute inequality. So, when we introduce a cost to rejecting offers not only is the inverted-U form still there, but now it is symmetric.

2.5 Conclusion

We design a three-player ultimatum game where a third party accepts or rejects all potential offers on how to split \$10 between two other subjects. If the offer is accepted, then the split goes as suggested; if rejected, then both proposer and receiver get \$0. The decision-maker payoff is our treatment variable. The first result of the experiment shows that in our experiment the decision-maker is completely detached form the game and does not care about his relative payoff when making decisions. This means that any result presented under our game structure is a mapping of the "primal" preferences of a decision-maker, the decisions of someone whose only concern is split between A and C. This is a new result in the social-preferences experiments, which we consider can be very useful not only in the social preferences literature, but also to other areas of research such as experiments involving arbitration or conflict resolution. Our second significant result is shows that decision-makers are so concerned with inequality that a significant number of generous offers are rejected. In fact, this phenomenon is so extreme that the results takes the form of an inverted-U pattern of acceptances; a shape that to our knowledge had never been observed in any lab experiment before. Furthermore, when we introduce a cost to rejecting offers, not only do we still observe rejections of generous offers, but we find that the inverted-U shape is now more symmetric. This symmetry comes from decision-maker rejecting generous and selfish offers in identical proportions. We use then use this symmetry to show that for the decision-maker concerns over the selfish intentions of the proposer are clearly of second order of importance when compared to concerns over absolute inequality. All in all, the results presented here seem to revive inequality aversion as the main behavioral driver around other regarding preferences. The fact that still now we were able to present surprising results on social preferences experiments means that experimental economics is still far from understanding how subjects behave in social environments. Clearly, (10 years later) more research funding is needed[17].

[17] See page 851 from Charness and Rabin (2002).

Chapter 3

Survivor curve shape and internet revenue: A laboratory experiment

3.1 Introduction

A goal of content providers is to turn attention to their websites into revenues that will at least offset their costs. Achieving this goal is not easy, even for providers with established audiences. Providers may charge subscription fees, present advertisements or some mix (Baye and Morgan (2000), Prasad et. al. (2003), Kumar and Sethi (2009)). But all revenue strategies take a toll — while some users see the nuisance as a fair exchange for the value obtained, other users see the nuisance as intolerable and leave the website, and some potential users are deterred from joining. The issue is especially acute with increasingly intrusive "rich media" advertising formats Godes et. al. (2009).

In this paper we do not investigate which revenue strategy is best, nor how to choose the optimal nuisance level in steady state; presumably the best choices are very situation specific. Instead we ask a simpler question: should a content provider introduce the necessary nuisance in gradual steps or all at once?

An incorrect answer can cause lasting damage. A case in point is Netflix's September 1, 2011 decision to raise fees abruptly on their basic service. Consumers left in droves, Netflix shares lost more than a third of their value in two weeks, and on September 18, CEO Reed Hastings apologized to everyone, saying "I messed up." The stock price has still not recovered. Another example is Digg's disastrous release on August 25, 2010 of its advertising-heavy v4. The issue is not confined to the internet, of course: witness swimmer's perennial debate of whether to jump into cold water or to wade in gradually.

After a brief review of other relevant literature, we begin in Section 3.3 by recalling the model of Aperjis and Huberman (2011, 2012). It establishes that, under a set of auxiliary assumptions, the answer to the question hinges on the shape of the survivor curve $S(x)$, the fraction of a human population willing to tolerate an inconvenience of magnitude x. If the logarithm of $S(x)$ is convex, then the content provider maximizes value by introducing the necessary nuisance all at once. If the logarithm of $S(x)$ is concave, then the nuisance is best introduced gradually according to a schedule that balances the number of long-term users against more rapid revenue acquisition.

Are survivor curves typically log-concave, log-convex, or neither? To the best of our knowledge, previous research provides no clear evidence. Behavior in natural settings is difficult to interpret because visitors leave for many reasons unrelated to the chosen inconvenience increment x, while new visitors arrive that may have different reactions to x and to the content. Moreover, when a change x is introduced, visitors may form beliefs about further inconveniences that may be introduced later, and such beliefs could vary widely across visitors. Competitors' adjustments in inconvenience might also have a major impact.

Laboratory experiments are especially helpful to answer the shape question, because one can control for all these confounding factors, and can systematically vary the nuisance size x. In section 3.4 we describe a recent experiment designed to discover the shape of the survivor function over a variety of domains. The experiment confronts 112 human subjects with six different tasks interrupted by nuisances of magnitude $x \in [x_{\min}, x_{\max}]$. It generates 636 binary observations of decisions whether to stay with an enjoyable activity or to leave after the nuisance has been imposed.

Section 3.5 collects the results. Summary statistics and preliminary analysis show that the chosen ranges $[x_{\min}, x_{\max}]$ are reasonably well calibrated, that order effects are unimportant, and that behavior is reasonably consistent across tasks. The main finding concerns the shape parameters in Weibull distributions estimated for data from each of the six tasks. Estimation requires extension of established techniques to deal (for the first time that we know of) with doubly censored data. Surprisingly (at least to some of the coauthors), the estimated shape parameters are all well inside the log-concave region.

The concluding discussion notes some caveats, suggests broader applications and implications for psychology and philosophy, and points to future research.

3.2 Related Literature

We know of no studies estimating the shape of survivor curves for scalable nuisances. Own-price demand elasticity is a distantly related topic. Within the vast literature on that subject, perhaps the most relevant article is Popescu and Wu (2007), which argues theoretically that firms with risk averse customers maximize profits by gradually increasing or gradually decreasing price. In an adaptation model, Fibisch et al. (2005) find that price elasticities increase over time, and that data suggest a faster adaptation for price decreases than for price increases.

Adaptation theory considers how users react over time to an introduced inconvenience. A number of papers consider adaptation in the context of repeat-purchase markets and characterize optimal dynamic pricing policies (Kopalle et al. (1996), Fibisch et al. (2003), Popescu and Wu (2007), Nasiry and Popescu (2010)) . In these papers, a firm (usually a monopolist) is facing consumers whose purchase decisions are influenced by past prices through reference price effects. The demand in a given period is assumed to be a function of the current price and the reference price (but does not depend on the number of people that purchased the product in the previous period). In a laboratory experiment Kahneman et al. (1993) suggest that duration plays a role in the recollection of aversive experiences, with reference points being formed at the peak and at the end of the negative experience [1].

There is an active theoretical literature on reference points (Kahneman and Tversky (1979), Frederick and Loewenstein (1999), Koszegi and Rabin (2006)) which has inspired many recent laboratory experiments, including Gneezy (2005) and Baucells et al. (2011). Abeler et al. (2011) find empirical evidence supporting Koszegi and Rabin (2006): payoff expectations seem to anchor reference points, as identified by subjects' effort choices. By contrast Heffetz and List (2011) find no support for the expectations reference point hypothesis. Closely related to this literature we find a number of experimental and empirical studies that focus on the formation of reference points (surveys are provided by Kalyanaram and Winer (1995), Mazumdar et al. (2005)). In these studies, the inconvenience is the price of a product, and thus the reference point is a reference price. Even though the role of historic prices in forming price expectations is supported in many of these studies, there has not been sufficient evidence for any specific model on how consumers update their reference prices.

There is also a classic psychology literature on "just noticeable differences," which

[1] The empirical adaptation literature is also related to studies such as Ariely (1998) that examine how remembered pain relates to the time path of pain intensity. It may be worth pointing out that our own concerns are quite different: we shall examine empirically how stay/remain decisions (not recollections) depend on one-shot intensities (not time paths) of nuisances (not pain) in a variety of modalities.

is associated with failures in the transitivity of preferences as in the self-torturer example of Quinn (1990), or the Sorites paradox[2] Finally, there is field data suggesting that firms generally prefer subdividing price increases but not price decreases (Chen et al. (2008)).

3.3 Theory

We consider the setting of Aperjis and Huberman (2011). In discrete time $t = 1,2,3,\ldots$, each period the provider has the option to adjust the total inconvenience level (e.g., advertisement level, subscription cost) X_t. Let $x_t \equiv X_t - X_{t-1}$ denote the adjustment in inconvenience at time t.

Assume that in period t, users have a reference point r_t and use the website with probability $S(X_t - r_t)$, where $S : R \to [0,1]$. That is, we assume that this probability only depends on the difference between the current inconvenience and the reference point. We assume that S is a decreasing function: the larger the difference between total inconvenience and the reference point, the smaller the probability of using the website.

Aperjis and Huberman (2011) rely on adaptation theory to describe reference point dynamics. That theory says that as time goes on people tend to adapt and become less aware of past changes. In the present context, an increase in inconvenience by an amount x initially decreases a user's utility. However, as time goes by the user's reference point gradually adapts and, as a result, his experienced utility gradually increases if no additional inconvenience is experienced.

Here we focus on the special case of "complete" adaptation within a single period.

[2] In Greek, soros means heap. The paradox is attributed to Eubulides of Miletus, a disciple of the Megarian school of philosophy who presented the following paradox: "no one grain of wheat can be identified as making the difference between being a heap and not being a heap. Given then that one grain of wheat does not make a heap, it would seem to follow that two do not, thus three do not, and so on. In the end it would appear that no amount of wheat can make a heap." (Hyde 2011)

That is, we assume that $r_t = X_{t-1}$. In this case, the probability that a user continues using the website at time t is equal to $S(X_t - X_{t-1}) = S(x_t)$. Thus subsequent theoretical analysis assumes that the survivor curve S is the same in each period and depends only on the most recent change in inconvenience.

Other simplifying assumptions are straightforward. Once a user leaves, he never returns, so the fraction of users remaining on the site at time t is $\rho_t = \prod_{j=1}^{t} S(x_j)$. The provider wishes to maximize the present value of his profit stream, $\sum_{t=0}^{\infty} \delta^t \rho_t \pi(X_t)$, where δ is the provider's discount factor and the current per-unit profit level $\pi(X_t)$ is an increasing function of the current inconvenience level.

The main conclusion of Aperjis and Huberman (2011). is that, under current assumptions, the provider's optimal schedule of inconvenience changes $(x_1, x_2, \ldots)$ depends entirely on the shape of the survivor curve. There are two important cases.

Log-concave survivor curve. A function is log-concave if its logarithm is concave. All concave and linear functions are log-concave, but there also exist convex functions that are log-concave. Examples include $S(x) = e^{-x^k}$ with $k > 1$ and $S(x) = (1-x)^k \cdot 1_{\{x \in [0,1]\}}$ with $k > 1$, where $1_{\{\cdot\}}$ is the indicator function. An important property of a log-concave survivor curve is that

$$S(x+y)S(0) \leq S(x)S(y)$$

for any $x, y \geq 0$. Here $S(x)S(y)$ represents the probability that a current user will continue to be a user if inconvenience increased by x last period and then by y this period, while $S(x+y)S(0)$ represents the corresponding probability when the entire inconvenience change $x+y$ was introduced in the current period. Iterating the inequality, it is intuitively clear that when S is log-concave, more users will remain if an increase in inconvenience is introduced gradually than if it is introduced all at once.

Aperjis and Huberman (2011) confirms the intuition, and shows that in the log-concave case it will be optimal for the provider to increase inconvenience gradually in order to give people time to adapt to changes. That paper then derives a specific schedule of changes that optimizes the tradeoff between maximizing the number of users in the long term and achieving a higher revenue per user sooner.

Log-convex survivor curve. A function is log-convex if its logarithm is convex. For instance, this is the case if $S(x) = 1/(1+x)^k$ with $k > 0$ or $S(x) = e^{-x^k}$ with $k \in (0,1)$. If S is log-convex, then $S(x+y)S(0) \geq S(x)S(y)$ for any $x, y \geq 0$, and therefore a user is more likely to stay if an increase in inconvenience is introduced at once than if it introduced gradually. When the survivor curve is log-convex, it is optimal for the provider to increase inconvenience once; this is shown by Aperjis and Huberman (2011) in a more general setting than the one we consider here. Note that this is not a result of selection, because the function S is assumed to not change over time.

Figure 3.1: Comparison Log-Concave and Log-Convex

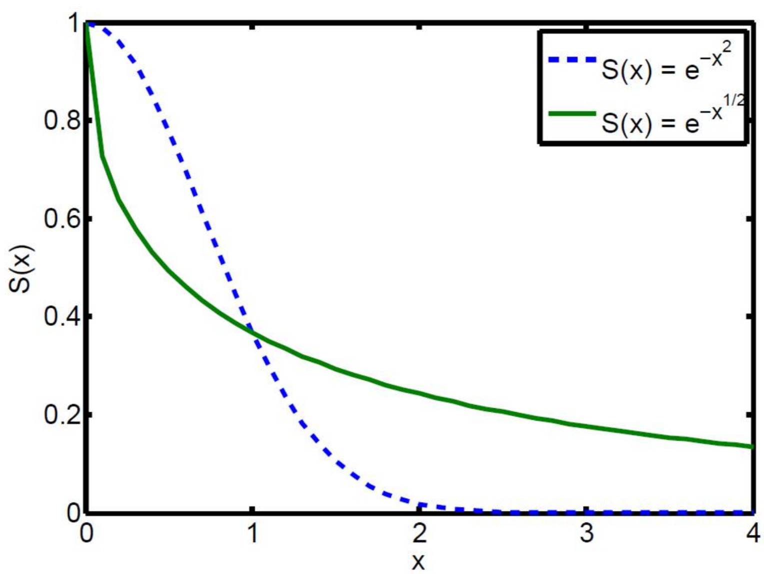

To get some intuition for the distinction between log-concave and log-convex survivor curves, consider Figure 3.1 which shows the log-concave function e^{-x^2} and the log-convex function $e^{-x^{1/2}}$. Note that for small deviations x from the reference point, the dashed line is above the solid line, indicating that a user is more likely to use the website when his behavior is described by the log-concave function. On the other hand, for large deviations ($x > 1$ in the Figure) the comparison is reversed, suggesting that if the survivor function is log-convex, it is better to make one large change.

Given that the optimal way to introduce inconvenience is so different for log-concave and log-convex survivor curves, it is important to understand whether one of the two shapes prevails. This motivates us to measure the survivor curves in the laboratory for a number of different activities and types of inconvenience.

3.4 Methods

The laboratory experiment presented subjects with tasks of the following sort. First, they engaged in a pleasurable activity, such as putting on earphones and watching an 8 minute video clip — their choice of an interview of John Stewart at The O'Reilly Factor, or a selection of the 10 most popular ads shown to viewers of the 2010 Super Bowl (Pilot experiments included a longer list of videos, but these two were the most popular). Then, after 100 seconds, an annoying computer-generated voice at $x \in [30, 80]$ decibels began reading the decimal expansion of $\pi = 3.14159....$ Subjects knew that the only way to escape the auditory nuisance was to click a button that immediately switched them to a bland activity, in this case watching a video of gentle waves breaking at La Jolla beach, for the remaining 6 minutes or so. Of course, a higher fraction of subjects switched when $x = 80$ decibels than when $x = 30$, and intermediate fractions switched at intermediate values of x.

Figure 3.2: Image of Game

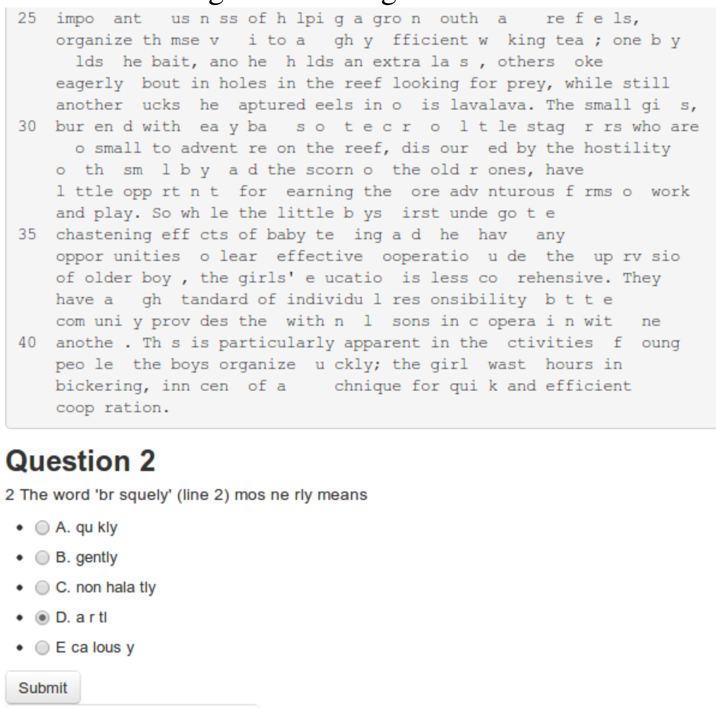

We also presented subjects with visual nuisances, like flashing pop-up ads that interrupted a video clip for 15 seconds every x seconds, with x ranging from $x = 5$ to $x = 30$. Figure 3.2 shows a text-based nuisance for the task of answering SAT questions, with a \$0.40-\$0.10 payment for each correct/incorrect answer. The nuisance is the random omission of each letter with probability $x \in [0.06, 0.21]$; in the Figure 3.2, $x = 0.15$. Subjects could escape the nuisance entirely by clicking a button, but then would be paid for the remainder of the 8 minute period at the much lower rate of $\$0.10/-\$0.02$.

We presented each subject with six distinct tasks that shared the common structure depicted in Figure 3.3. The subject starts with an engaging activity (A activity), which after a certain amount of seconds is interrupted by a scalable nuisance of size x that remains attached to the A activity thereafter. She can escape the nuisance at any time by clicking a button to switch to a "bland" activity (B activity) where she will remain for the rest of the 6-8 minute period. Her choice of whether or not to switch is a data

Table 3.1: Task specification.

Task	Activity A	Activity B	Inconvenience:	Range of x
Movie/Pi	Watch Movie	Watch Waves	π digits	[30, 80] decibels
Movie/Pop	Watch Movie	Watch Waves	15 sec Pop-up	every [5, 30] sec
Slug	Slug ($$$$)	Slug ($)	Jitter	[0.10, 0.25] rate
Read	Read Article	Count Bits	Drop Letters	[0.15, 0.30] rate
SAT	SAT ($$$$)	SAT ($)	Drop Letters	[0.06, 0.21] rate
Pay	Watch Movie	Watch Waves	Pay to Stay	[1, 23] cent fee

point that helps us estimate the shape of $S(x)$.

Figure 3.3: Image Game Structure

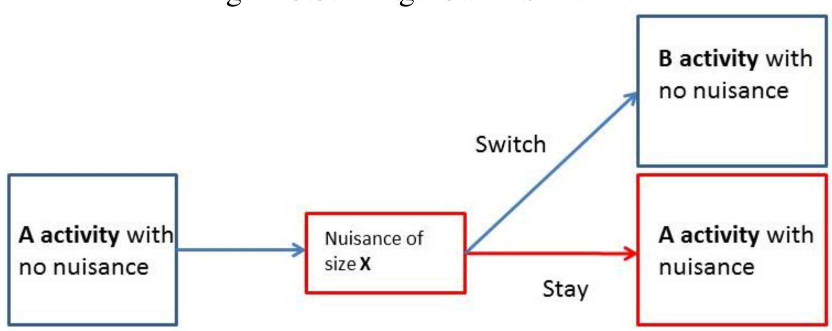

Table 3.1 summarizes the six combinations of A activity, scalable nuisance, and B activity presented to each subject. Of the entries not yet mentioned, Slug is a simple video game similar to Snake (see Appendix B for a detailed description of the activity), and the jitter nuisance involves a random turn (left or right) each pixel with probability $x \in [.10, .25]$. The Pay to Stay nuisance is a one time fee of x cents deducted from a 500 cent endowment, which can be avoided only by switching to the B activity. The B activity Count Bits is illustrated in Figure 3.4 below. Paid activities are indicated by ($$$$), and B activities paid at 1/4 the rate are indicated by ($).

The nuisance ranges $[x_{\min}, x_{\max}]$$ were chosen to avoid inefficient sampling where $S(x)$ is very close to 0 or 1. Based on a few pilot sessions, we aimed to have $S(x_{\min})$ in the vicinity of 0.8 and $S(x_{\max})$ in the vicinity of 0.20. Nuisance levels were chosen to span the range by six evenly spaced levels, as detailed in Appendix A.

Figure 3.4: Example of Bland Activity

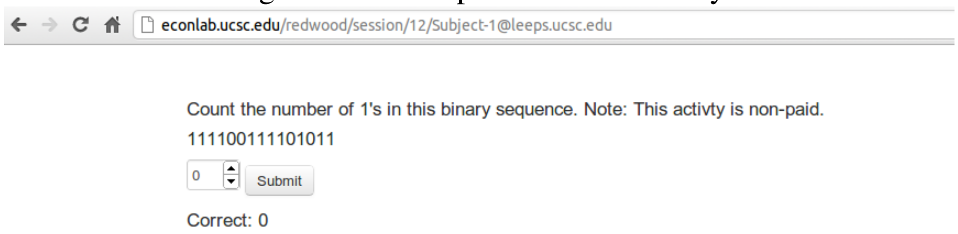

3.4.1 Procedure

We recruited 112 human subjects, from the LEEPS lab subject pool, most of whom are UCSC undergraduates, typically majoring in Economics, Biology or Engineering. Subjects would join one of the 16 sessions we ran, and none was allowed to participate twice. Sessions lasted 70 to 90 minutes including instruction reading and individual payments.

Upon arrival, each subject was assigned to an isolated computer terminal, and general instructions for the experiment were read (a copy is attached in Appendix C). Next, subjects practiced all B activities, in order to ensure that they knew exactly what they would do if they decided to switch to a bland activity. Subjects were then given specific instructions for the first of the six tasks, after completion they received instructions for the second round, completed it and were given instructions for the third task, etc. The order of the six tasks was varied in a balanced manner across sessions. In each session we randomly assigned each subject's nuisance level x, but limited the choice to one of the two nuisance bins that we created; either $x = 1,3,5$ or $x = 2,4,6$ in each session. These bins allowed us to have in each session a sizeable number of observations with the same treatment level in each activity.

Before each round it was announced whether A and B would be paid activities. If they were, then a detailed description of the payment system was given. If they weren't paid, then we emphasized it in the instructions. Subjects would know how much money they had made at the end of each paid round, and once the experiment was over, they were paid individually. Payoffs ranged from \$27 (some subjects proved very proficient at Slug) to \$12 (some were not that apt), including the \$5 show-up fee. On average subjects made around \$16.

3.5 Results

The experiment yielded 636 data points $(Y_{i,j})$, observations of whether or not subject i decided to switch after experiencing inconvenience level x in task j. Due to implementation glitches, we lost one Slug data point and the SAT data in two sessions (35 data points); hence the slight shortfall from the intended $6 \times 112 = 672$ observations. Figure 3.5 summarizes the data graphically.

Figure 3.5: Results

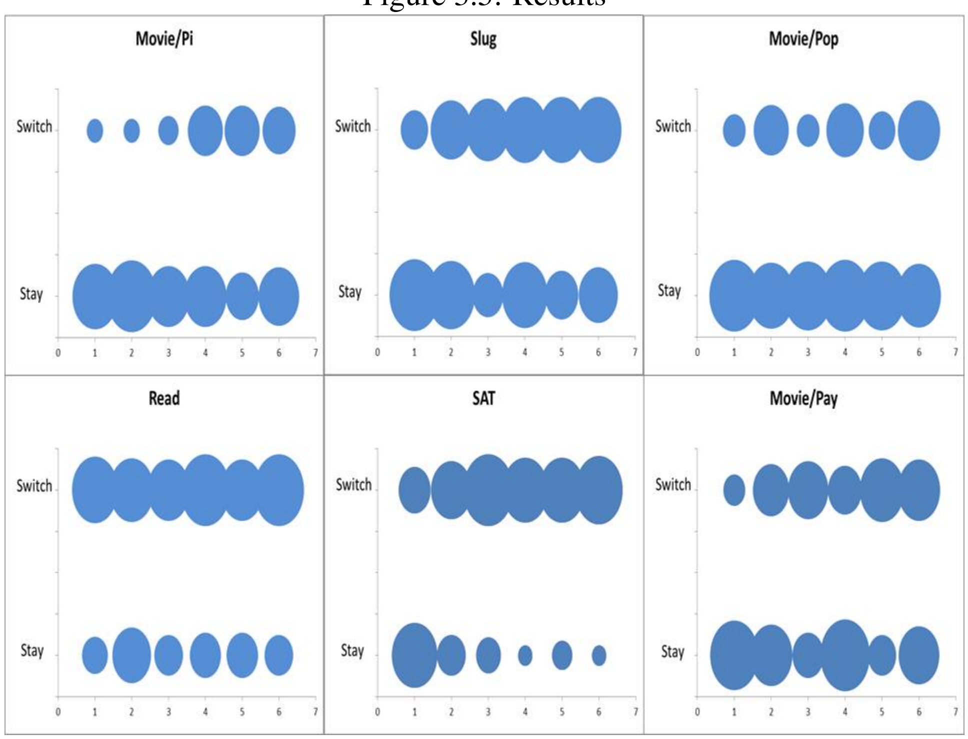

As a first step in the data analysis, we run a Probit regression of the binary outcome $(Y_{i,j})$ on dummies for inconvenience levels 2-6, task numbers 2-6 as in Table 3.2, and the task sequence or session.

As we can see in Table 3.2 all levels of inconvenience have a highly significant effect, as one would hope. So do most tasks, except Movie/Pop, which is not significantly different than the baseline task, Movie/Pi. Appendix A reports additional robustness checks, and confirms that there were no important session or sequence effects.

Finally, we use a Fisher Exact test to compare the proportion of subjects switching for each value of x across activities. The results show that for any value of x the difference in proportion is not statistically significant, pointing towards a similar underlying distribution of subject tolerance for nuisance levels across activities. This result will be ratified in our survivor curve estimates.

3.5.1 Estimation Strategy

The main objective of our experiment is to detect log-concavity or log-convexity of $S(x)$ separately in each of our six tasks. To do this we will consider each observation (switch or not) for each subject i as an independent observation for each separate curve j.

The main complication with our data comes from censoring. If $Y_{ij} = 1$, i.e., if subject j switches to the bland activity B when facing nuisance level x_i, then we infer that her switching threshold X is somewhere in the interval $(0, x_i)$, and thus observation is left censored (LC). Therefore the likelihood of the observation is given not by the density at x_i but rather by the cumulative distribution function F evaluated at that point: $F(x_i) \equiv P(X \leq x_i)$. On the other hand, if $Y_{ij} = 0$, i.e., if subject j stays in activity A, then we infer that his threshold is in the interval (x_i, ∞), and the observation is right censored (RC). The likelihood of such an observation is $1 - F(x_i) = P(X > x_i)$, where

Table 3.2: Switching Probit Model.

	(1) Switch	(2) Switch	(3) Switch	(4) Switch
2.treatment	0.299*		0.359*	0.735***
	(0.180)		(0.190)	(0.280)
3.treatment	0.578***		0.633***	0.649***
	(0.171)		(0.187)	(0.177)
4.treatment	0.524***		0.642***	0.968***
	(0.165)		(0.175)	(0.257)
5.treatment	0.746***		0.845***	0.913***
	(0.164)		(0.178)	(0.199)
6.treatment	0.737***		0.859***	1.265***
	(0.179)		(0.196)	(0.281)
2.activity		0.0511	0.0591	0.180
		(0.174)	(0.176)	(0.194)
3.activity		0.551***	0.583***	0.708***
		(0.174)	(0.176)	(0.185)
4.activity		1.029***	1.072***	1.322***
		(0.172)	(0.174)	(0.243)
5.activity		1.144***	1.238***	1.339***
		(0.199)	(0.191)	(0.238)
6.activity		0.428**	0.444**	0.554***
		(0.170)	(0.177)	(0.185)
Cons	-0.524***	-0.540***	-1.133***	-1.282***
	(0.125)	(0.125)	(0.186)	(0.221)
N	636	636	636	636
Order Dummies	No	No	No	Yes

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

$S(x_i) \equiv 1 - F(x_i)$ is the probability that the subject "survives" the introduction of the inconvenience.

This likelihood function applies to any parametric family of survivor curves. We use the standard two-parameter Weibull family. Recall that Weibull distribution has density $f(x;\gamma,\kappa) = \begin{cases} \frac{\kappa}{\gamma}(\frac{x}{\gamma})^{\kappa-1}e^{-(\frac{x}{\gamma})^{\kappa}} & if\, x \geq 0 \\ 0 & if\, x < 0, \end{cases}$ where $\kappa > 0$ is the shape parameter and $\gamma > 0$ is a scale parameter for the distribution. The corresponding cdf is $F\,(x;\kappa,\gamma) = 1 - e^{-\left(\frac{x}{\gamma}\right)^{\kappa}}$, and thus the survivor function is $S\,(x;\kappa,\gamma) = e^{-\left(\frac{x}{\gamma}\right)^{\kappa}}$.

Besides being standard, the Weibull family has the extremely convenient property that the shape parameter κ determines whether the survival function $S(x)$ is log-convex or log-concave (Bagnoli and Bergstrom (2005)):

- $S(x)$ is log-convex (and the hazard rate is strictly decreasing) if $0 < \kappa < 1$, and
- $S(x)$ is log-concave (and the hazard rate is increasing) if $\kappa \geq 1$.

Econometric packages usually include the Weibull distribution, and sometimes can deal with singly censored data, but we must build our own likelihood function to deal with doubly censored data. It follows from the preceding discussion that the likelihood function for data $Y = (Y_{ij})$ is:

$$\begin{aligned} L(\gamma,\kappa|Y) &= \prod_{Y_{ij}\varepsilon LC} P(X < x_i|\gamma,\kappa) \prod_{Y_{ij}\varepsilon RC} P(X > x_i|\gamma,\kappa) \\ &= \prod_{Y_{ij}\varepsilon LC} \left(1 - e^{-\left(\frac{x_i}{\gamma}\right)^{\kappa}}\right) \prod_{Y_{ij}\varepsilon RC} e^{-\left(\frac{x_i}{\gamma}\right)^{\kappa}}. \end{aligned} \tag{3.1}$$

We maximize function 3.1 over the parameter space using standard numerical techniques in Stata to obtain point estimates of the shape parameter κ. To obtain standard errors we use standard bootstrap procedures appropriate for finite samples. The results are reported in Table 3.3.

Table 3.3: Weibull estimation results.

Task	Shape Parameter
Movie/π	3.78±0.74
Movie/Pop	2.75±0.47
Slug	2.81±0.33
Read	2.81±0.22
SAT	2.40±0.20
Pay	2.94±0.37

All shape parameters are significantly greater than 1, so the estimated survivor curves for all six tasks are log-concave.

Finally, the hazard rate (in other contexts sometimes called the failure rate) h for switching is the density for switching at nuisance level x conditional on not switching at a lower level, i.e., $h(x) = f(x)/S(x)$. By choosing the Weibull family for our parametric distribution the only assumption we have made is that the hazard rate follows a monotonic path, which seems like a sensible assumption. But, what makes our results even more intuitive is that, in the Weibull family, if the shape parameter is greater than 1, not only is the shape of the survivor curve log-concave, but it also means that the hazard rate is increasing with the nuisance level. Symmetrically, had κ been smaller than one for some activity, not only would the corresponding survivor curve be log-convex, but it would also imply a hazard rate that is decreasing in the nuisance level (something that seems unlikely for any task in our experiment).

3.6 Discusion

The lab results are remarkably clear cut: all six tasks yield highly significant Weibull shape parameter estimates well inside the log-concave region. This is so even though the tasks and nuisances involve very different domains, including visual, auditory, text, fine motor and monetary.

The implication is clear within the theoretical framework of (Aperjis and Huberman (2011)): web content providers should introduce the necessary nuisances gradually to reach their target revenue. How gradually, of course, depends on the provider's discount rate and also depends on how rapidly users adapt to nuisances.

As with any empirical results, several caveats are in order. Our results were based on the decisions of more than 100 human subjects recruited from the LEEPS lab subject pool. a subject pool consisting mostly of undergraduate students in a US university. It is entirely possible that other populations would be more or less tolerant of nuisances than ours, and thus have survivor curves with different scale or location. However, it seems to us rather implausible that they would yield survivor curves with much different shape than ours, but of course that can only be confirmed through further research.

A second caveat is that we worked within the framework of a simple model, which neglected potentially important aspects of reality. For example, it ignored the arrival of new users. A slight extension of the model could easily incorporate them if their survivor curves resembled those of the original users. Although new users might differ from the originals in various ways, again there is no reason to suppose that their survivor curves have radically different shape.

Perhaps the more important caveat, and the most intriguing, is that the adaptation process may differ from that envisaged in the theoretical model. As noted in the literature survey, there is considerable recent empirical research on such matters, much of it inspired by Prospect Theory and in particular by Koszegi and Rabin (2006). So far the work seems inconclusive, but when a consensus emerges on reference point dynamics, it should be incorporated into a richer model of dynamic decision making.

Chapter 4

That's How We Roll: an experiment on rollover risk

4.1 Introduction

While the literature agrees on placing a run on short-term credit at the center of the recent financial crisis (e.g. Brunnermeier (2009), Krishnamurthy (2010)), there is much less consensus on how to prevent another panic. Brunnermeier *et al.* (2009) suggest that extending the maturity of short-term credits might help stabilize the market by making it less volatile. We use experimental tools to test the effects that this policy suggestion has on the market for short-term credit. To be more precise, we use experimental tools to investigate the effects of different maturity lengths on a market for Asset Backed Commercial Paper (ABCP), which is a specific type of short-term credit in which, if the issuing firm does not fulfill its promises, the holder of the ABCP can seize the posted collateral. Our results show that while on average markets with shorter credit maturities have a lower probability of freezing, a more detailed analysis of our data indicates that the optimal policy should be state-dependent, favoring long contracts when the economy is in good shape, and allowing for short ones during a recession. Our data also shows a significant number of firms with strong fundamentals being "locked out" of the credit market, as well as rich experimental dynamics with a consistent credit bubble and crash across sessions.

4.1.1 Why run an experiment on ABCP?

ABCP has been pointed as the necessary transmitter of the housing bubble into the financial system (Brunnermeier (2009)), so, while short-term credit is not a problem *per se*, ABCP played a central role in the financial meltdown and credit freeze of 2007. The argument is that ABCP (usually supported by structured subprime mortgages) took over the more "traditional" credit market in the years before the crisis, and by virtue of being cheap and unregulated, it exposed the market to a credit bubble, and to "excessive mismatch in asset-liability maturities". In Figure 4.1 (borrowed from Brunnermeier (2009)) we see how the market for ABCP almost doubles in size from 2005 to 2007 (the final years of the housing bubble), to crash and drop from an outstanding $1,200 billion to $750 billion in just six months in 2007. Yet, what is most interesting about this graph is not the increase in ABCP, but rather how unsecured instruments were only slightly affected at the peak of the 2007 crisis. This suggests that the problem revolved more around a change in the perceived value of collateral than around the use of short-term credit. In fact, even previous to this crash of the market, a whole literature had already developed around the idea of regulating collaterals (e.g., Geanakoplos (2009, 1996)).

Figure 4.1: Outstanding Asset backed and Unsecured Commercial Paper Comparison

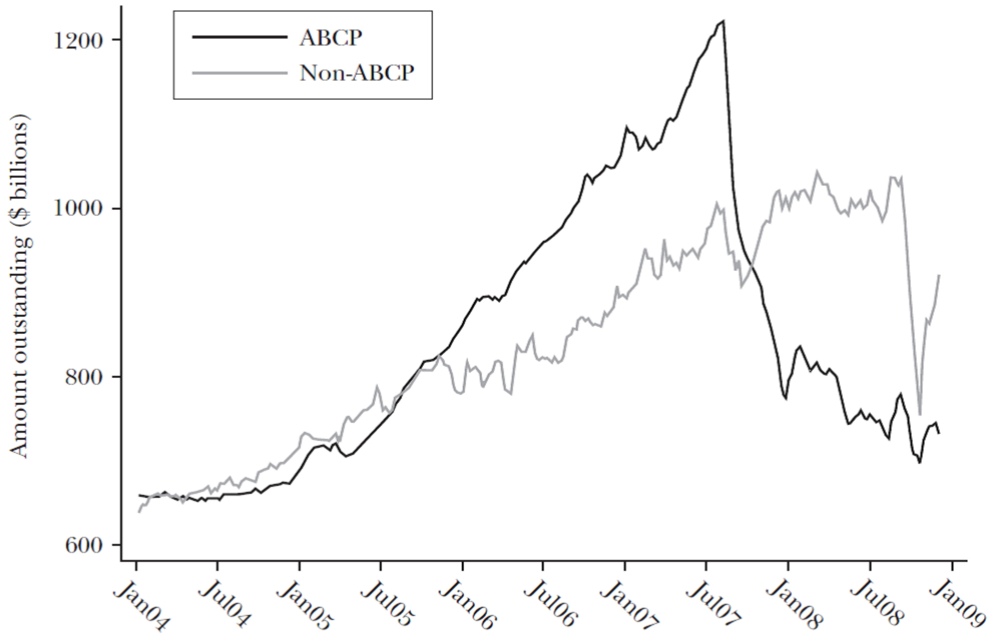

Source: Federal Reserve Board.

In a more specific analysis, Shin (2008) looks at the particular case of Northern Rock, and explains the shift in paradigm that modern financial markets, have brought to our understanding of "bank runs". As he puts it, while we all remember the lines forming at the doors of Northern Rock, the real storm had occurred weeks before, when non-depository creditors (mostly of ABCP) decided not to roll over their credits to the bank.[1] The important question, according to the author is thus "not so much why banks depositors are so prone to running, but instead why the plentiful short-term funding (...) suddenly dried up".

Finally, it looks like the Federal Reserve agrees with the diagnosis, and has pointed at credit markets as the main reason for the recent crisis (Bernanke (2009a), (2009b), (2008)), making it modify its policy to a "credit-easing strategy rather than a quantitative-

[1] As Shin (2008) puts it "The depositor run, although dramatic, was an event in the aftermath of the liquidity crisis at Northern Rock".

easing approach" (Bernanke (2009b)). In fact, in a 2008 speech[2] , Bernanke expressed this shift in the way the Federal Reserve would approach financial panics:

> "Bagehot defined a financial crisis largely in terms of a banking panic – that is, a situation in which depositors rapidly and simultaneously attempt to withdraw funds from their bank accounts. In the 19th century, such panics were a lethal threat for banks that were financing long-term loans with demand deposits that could be called at any time. In modern financial systems, the combination of effective banking supervision and deposit insurance has substantially reduced the threat of retail deposit runs. *Nonetheless, recent events demonstrate that liquidity risks are always present for institutions –banks and nonbanks alike–that finance illiquid assets with short-term liabilities*." (Emphasis added)

4.2 Our experiment in the context of the experimental literature

No experimental literature exists on the topic we are covering, so we use as references two strands of experimental research which are relatively close to our experimental design. The first one corresponds to continuous-time experiments, the second to "timing experiments", with a special emphasis on the experimental bank-runs literature. Continuous time experiments started years ago, with Friedman and Cheung (2009) and Morgan and Brunnermeier (2010) (whose working papers appeared around

[2]Given at the Federal Reserve Bank of Atlanta Financial Markets Conference on May 13 2008.

2003/04), but it has not been until recently that this experimental technique has taken off with Oprea et al. (2009) and Anderson et al. (2010) looking into strategic investment decisions, Oprea et al. (2011) studying the evolutionary equilibrium of the hawk and dove game, Friedman and Oprea (2012) experimenting with the effects of response delay in a repeated prisoners dilemma game, and Rabanal (2012) looking at mortgage default timing. While none of these papers directly address any of the questions of our paper, they are a good reference for the methodological design of our experiment.

The other relevant stand of literature is on experimental bank runs. To our knowledge, the first paper on this topic is Madies (2006), which is based on the theoretical model of Diamond and Dybvig (1983; DD). The results show that deposit insurance cannot avoid bank runs, and that the more experienced subjects are, the more often runs are observed. Garrat and Keister (2009), also test the DD setup but turn it into a repeated game by giving subjects the opportunity to exit several times per round. Schotter and Yorulmazer (2009) also adopt this technique. Both papers find that not only more experienced subjects are more prone to runs, but that the more opportunities to run within each round, the more likely runs are. Surprisingly, Garrat and Keister (2009) report having to exogenously force some subjects to exit, else no panics would occur. More recently Arifovic *et al.* (2011) look at how bank runs can be understood as a pure coordination problem. Finally, (Klos and Strater (2012)) approach bank runs from a Global Games perspective.

In summary, while there exists some experimental literature studying banking panics, most of it is based on models of "classic" bank runs, and none addresses the intricacies of modern financial markets, and the freezing of short-credit markets.

4.3 Theoretical Benchmark

Our experiment is inspired on the continuous time model by He and Xiong (2012;

HX). In it, a firm finances its long-term investment by issuing short-term credit to a continuum of creditors. Without loss of generality we will assume this credit to be of \$1. The value of the firm[3] follows a geometric Brownian motion and is perfectly observable by all agents. The Brownian motion can be written as:

$$\frac{dy}{y} = \mu dt + \sigma dZ \tag{4.1}$$

Where y_t is the value of the firm, μ is the drift, σ the volatility, and Z the standard Brownian motion.

Each creditor's debt matures with the arrival of an independent Poisson shock of intensity $\kappa > 0$, creating a uniform distribution of the maturities[4] , with all contracts having an expected duration of $1/\kappa$ at any point in time. This random maturities system is a simplifying assumption akin to Calvo pricing (Calvo (1983)), and avoids agents having to keep track of all other maturities when making the rollover decision, while still capturing all of the first order effects of other maturing contracts.

If within the time interval $[t, t+dt]$ enough creditors decide not to rollover their credit, then the firm draws from its cash reserves[5] (ϑ) and survives, on average, an extra $1/\vartheta\kappa$. Once the firm runs out of reserves it goes bankrupt and liquidates its assets at a discount value $\alpha < 1$, so the value of the asset is $\alpha F(y_t)$, where $F(y_t)$ is the present discounted value of the firm.

As payoffs, agents receive a stream of interests r until $\tau = \min(t_m, t_b, t_d)$ which is the earliest of three possible events. The first event (t_m) is the maturing of the long-term investment of the firm, in which case the agent gets back $\min(1, y_{t_m})$ and the firm ceases

[3]We will assume that the firm's only investment is on the long-term asset. Therefore, the value of long-term asset is the total value of the firm.

[4]Most firms spread out their maturities to avoid having large liquidity needs on any one specific date.

[5]He and Xiong (2012) describe ϑ as unreliable credit lines that the firm may tap, which is why the extra time is a function of the contract length. We believe that describing ϑ as cash reserves is more intuitive for our experimental purposes.

to exist. That is, the firm pays back the full principal of the credit if it can, or whatever it can pay back (but never more than the original \$1 credit). The second possibility (t_b) is a bankruptcy of the firm, in which case the creditor gets back $min[1, \alpha F(y_t)]$. Finally, the short-term credit can mature (t_d), at which point the creditor will decide to rollover his credit if the continuation value $V(y_{t_d}; y*)$ is higher than getting his credit back (\$1), where y_{t_d} is the value of the firm at the maturity point t_d, and $y*$ is the stopping threshold of other agents . The continuation value is thus written as:

$$\begin{aligned} V(y_t; y*) \quad = \quad & E_t \left\{ \int_t^{\tau} e^{(-\rho(s-t))} r ds + e^{(-\rho(\tau-t))} \left[\min(1, y_t) \mathbf{1}_{\tau=t_m} \right] \right\} + \\ & \min\left(1, \alpha F\left(y_t\right)\right) \mathbf{1}_{\tau=\mathrm{t_m}} + \max_{\mathrm{rollover\,or\,run}} \left\{0, 1 - \mathrm{V}\left(\mathrm{y_t}; \mathrm{y}*\right)\right\} \mathbf{1}_{\tau=\mathrm{t_d}} \end{aligned} \tag{4.2}$$

In equation (2) ρ is the discount value of the agent, and $\mathbf{1}_{\{.\}}$ is an indicator function which takes value 1 whenever the subscript is true, zero otherwise.

By evaluating the change in value of the continuation value (2) over a small time interval $[t, t+dt]$ the authors can write the Hamilton-Jacobi-Bellman:

$$\begin{aligned} \rho V(y_t; y*) \quad = \quad & \mu y_t V_y + \frac{\sigma^2}{2} y_t^2 V_{yy} + r + \phi \left[\min\left(1, y_t\right) - V\left(y_t; y*\right)\right] + \\ & \kappa \delta \mathbf{1}_{\{\mathbf{y_t} < \mathbf{y}*\}} + \delta \max_{\mathrm{rollover\,or\,run}} \left\{0, 1 - \mathrm{V}\left(\mathrm{y_t}; \mathrm{y}*\right)\right\} \end{aligned} \tag{4.3}$$

The left hand side represents the required return to the creditor, the first two terms in the right hand-side evaluate the fluctuation in the value of the firm. The equation also contains the continuation values of each of the three outcomes (long-term maturity, bankruptcy, short-term maturity) weighted by the probability of each one.

Finally, from equation (3) the authors show that agents will rollover the credit *if* *and only if* $V(y_t; y*) > 1$, that is, if the continuation value is greater than getting back the principal of the credit and "walking away". This result takes them to a unique *symmetric* equilibrium which is determined by the condition $V(y*; y*) = 1$, where no subject rolls over the short-term credit for a firm whose value is below $y*$, and always does so for values above $y*$.

What should be understood from this model is that unlike global games, subjects do not get a noisy signal, but a precise one. The strategic uncertainty comes from the asynchronous structure of the maturities, and the frequent change in value of the firm. It is precisely from these two key elements that agents can coordinate on a unique equilibrium, and this is why we can have results that would never happen in classic static models.

4.4 Experimental Implementation

4.4.1 Basic Design

Our experiment considers groups of 4 subjects where each member of the group provides a $1 short-term credit to a firm which has made a time varying long-term investment. Each group is composed of 4 subjects (which is a number close to the 5 individuals in the groups of Garratt and Keister [2009]), and the composition of the groups remains invariable during all of the 60 rounds that a session lasts. All subjects are informed of the size of their group and of its unchanged composition during the session.

The time unit of our experimentare ticks. Following Anderson et al. [2008], each tick is 1/5 of a second (i.e., 200 milliseconds). Each of the 60 rounds has a random end which is governed by a Poisson process, and has an expected length of 150 ticks (30 seconds), at which point the long-term investment matures and the firm ceases to exist. In HX, the value of the long-term investment (y_t) follows a geometric Brownian motion. Given the discrete nature of computer internal clocks, we will need to discretize this Brownian motion, and to do so we will use the procedure described in Anderson et al. [2008].

In each of the 60 rounds we will ask subjects to make one and only one decision, namely, whether or not to continue rolling over their credit to the firm. If at any time 2 subjects decide not to rollover their credit, then the firm will continue to run for a fixed θ of ticks before it goes bankrupt and has to liquidate its assets at a fire-sale value. This "extra time" θ is a linear function of the duration of short-term contracts and can be interpreted as the cash reserves of the firm[6] . The decision to choose 2 out of 4 subjects as the threshold for bankruptcy is again inspired by Garratt and Keister [2009], where the bank goes bankrupt if 3 out of 5 subjects decide to run.

The payoffs for each round will depend on the value of the firm at time $t, (y_t)$, and the decisions made by each subject in the group. To be precise, each round's individual payoffs will accrue from two different sources:

1. Flow payoff: For each tick that a subject keeps his investment in the firm she receives \$0.004 (i.e., \$0.6 for every 30 seconds invested).

2. End of round status: Depending on the decisions of the particular subject and the decisions of the other members of the group, the round could end in three different ways.

   (a) Exit: if a subject exits the project at time t_e, then she gets back her initial investment of \$1, independently of the value $y_{(t_e)}$ of the project at that point.

[6]This parameter comes directly from HX. In a future experiment we want to test the effects of changing θ.

(b) Bankruptcy: if at time t_b two subjects have stopped rolling over their credit, then the firm will run on its cash reserves for θ ticks, until finally going bankrupt at $t_{b+\theta}$, and being forced to sell its assets in the secondary market at a value$\alpha(F(y_{t_{(b+\theta)}}))$, where $F(.)$ is the present discounted value of the firm and$0 < \alpha < 1$. At this point the firm will pay all subjects still invested $Min[1, \alpha F(y_{t_{b+\theta}})]$ and then will cease to exist.

(c) Natural Ending: If the firm reaches its random "natural" ending $t_n$without going bankrupt, then all subjects still invested in the firm get $Min[1, y_{t_n}]$.

Subjects can keep track of both the firm's value (green jagged line in Figure 4.2), and of the fire-sale value (golden jagged line in Figure 4.2) in the graphical interface on their screen. Other useful information appearing on the screen are the values at which subjects in the group decided to stop rolling over their credit in the previous 15 rounds (upper right box in Figure 4.2), the \$1 threshold under which payoffs would be <\$1 (horizontal red line in Figure 4.2), and the moment they had exited, if they had decided to do so (vertical green line in Figure 4.2).

Figure 4.2: Screenshot

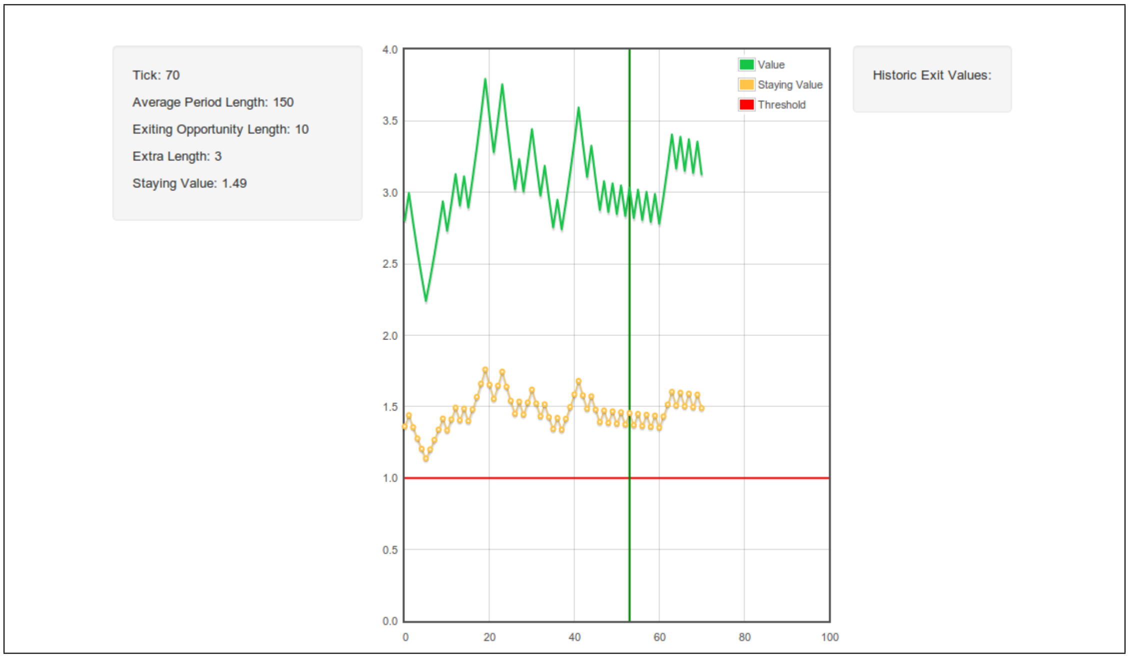

4.4.2 Credit rollover and credit maturities

The maturities system of this experiment is one of its unique aspects when compared to the experimental bank run literature, since in our game subjects decisions are not simultaneous. This creates several problems. The first is how to keep the game flowing (especially in the short maturity treatment) when subjects decide whether to rollover or not at every maturity. Our solution consists in having the credits rolled over by default, unless a subject decides otherwise.

To stop this automatic rollover, a subject will have to “connect” three numbered buttons on the screen by hovering over them in a precise sequence. We borrow this idea from Brunnermeier and Morgan [2010], where the purpose of the hovering mechanism is to avoid subjects making inferences from any clicks they can hear coming from other terminals. Our version is a slight variation over their mechanism: Not only do subjects have to hover, but they have to do it following a certain gradient, thus making it difficult to accidentally stop rolling the credit by inadvertently hovering over the designated area.

Subjects can decide to stop rolling over their credit at any time during the experiment. If they decide not to roll over, this decision will be implemented at the next maturity point. In fact, the second problem we confronted with our staggered maturity system was how to avoid turning the maturity points into focal points. The reason is that we are interested in observing at what value of the collateral agents decide to stop rolling over their credit, and we do not want this decision to be contaminated by the distance to the next maturity point. To avoid this confounding effect, we hide the maturity points from our subjects.

Figure 4.3: Image of random maturity mechanism

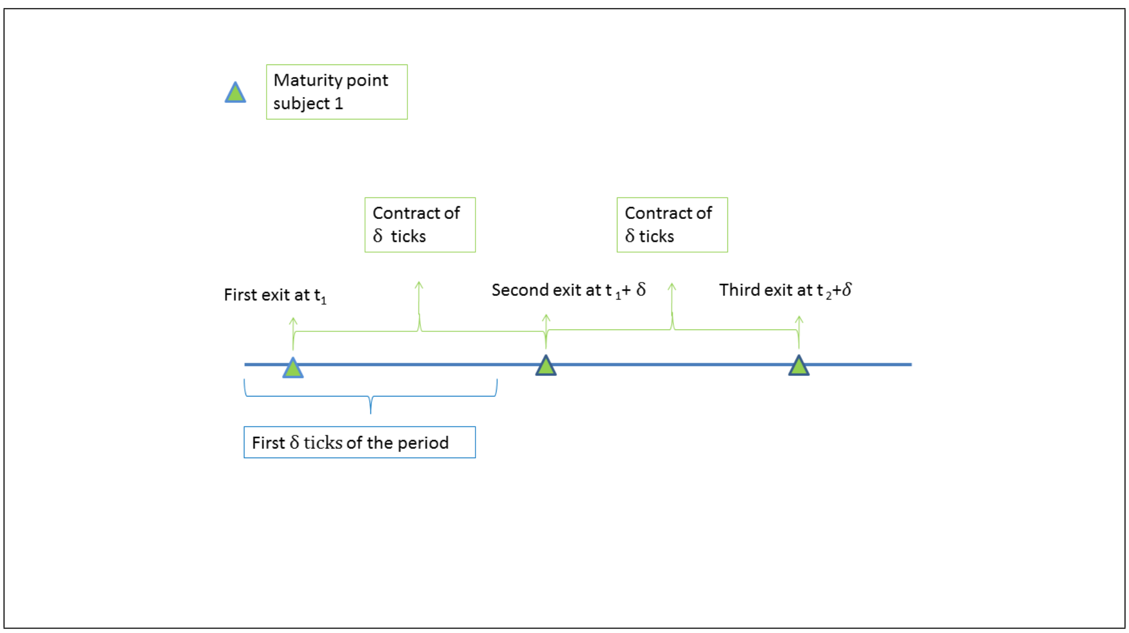

To hide the maturity points, we fix the length of credits to be δ ticks and have the computer randomly assign in each round j, and for each subject i, a random starting point t_{1ij} within the first δ ticks. From this initial (individual) point, maturities will happen every δ ticks. So, for example, for subject i in round j his first maturity point will be at $t_{1ij} \varepsilon [0, \delta]$, his second maturity point t_{2ij} at $t_{2ij} = t_{1ij} + \delta$, the third maturity t_{3ij} at $t_{3ij} = t_{2ij} + \delta$, etc. (see Figure 4.3). As a result, at every point in time the expected maturity of every subject is $\delta/2$ ticks away, avoiding a focal point problem.

Finally, we borrow the idea of the "pseudo-strategy method" from Anderson et al. [2008] and let all rounds play until their random ending, without providing any information to subjects of what other members of their group are doing. Once the round ends, all subjects get a screen shot describing all the events in the round, including other subjects (and own) requests to exit (green vertical lines in Figure 4.4), other subjects (and own) actual exit (orange vertical line in Figure 4.4), as well as round length and final payoffs, and a bankruptcy point (if there was one) shown as a red vertical line.

The implementation of this "pseudo-strategy" method will allow us to collect more information at every round, and avoid the problem of censored values that would occur if the rounds did not run until their random natural ending. For a review of the strategy method, see Brandts and Charness [2011b].

Figure 4.4: Pseudo-Strategy Method Screen

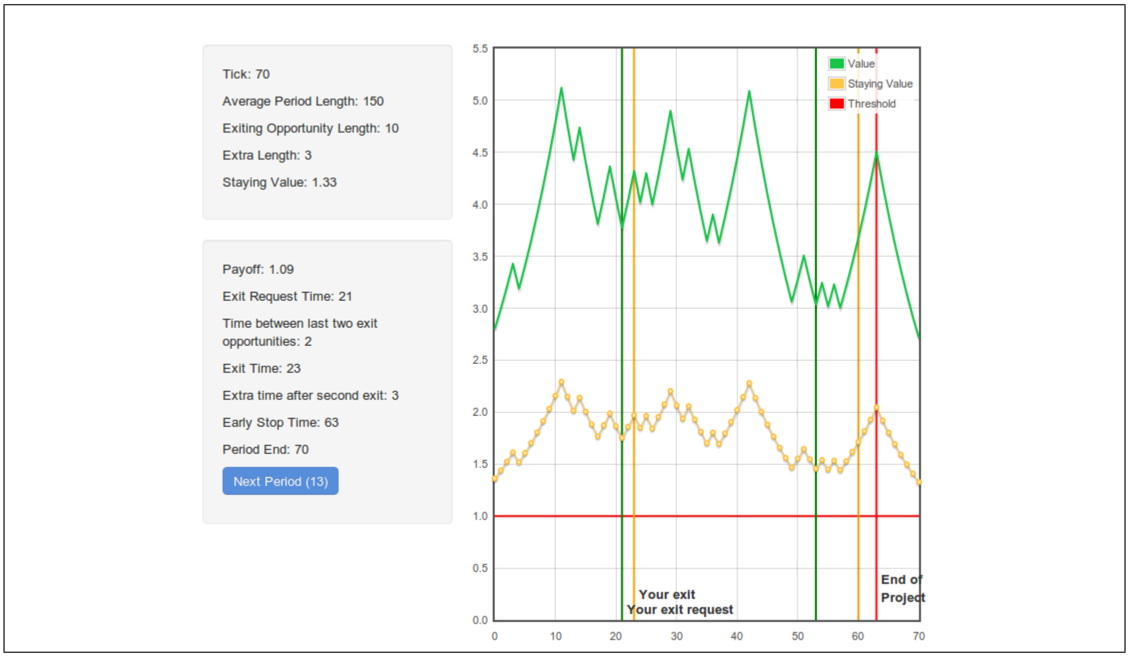

4.4.3 Parameters and Hypotheses

As mentioned above, the goal of this experiment is to test the effects that maturity lengths have on the market for short-term credit. We implement two treatments:

- Long treatment: Each contract is 8 seconds long (i.e., $\delta = 40$ ticks), and cash reserves last for 15 extra ticks after 2 subjects exit the market.
- Short treatment: Each contract is 2 seconds long (i.e., $\delta = 10$ ticks), and cash reserves last for 3 extra ticks after 2 subjects exit the market.

Table 4.1: Parameter Values

Parameter	Long Contract	Short Contract	Comment
δ	40 ticks	10 ticks	Contract Length
θ	15 ticks	3 ticks	Cash reserves
μ	0.0024	0.0024	Drift of the GBM
r	\$0.004 per tick	\$0.004 per tick	Per-tick flow payoff
σ^2	1.1	1.1	Volatility

Plugging these parameters into the He and Xiong [2012] model, it predicts that the optimal stopping threshold in the Short treatment will be higher than in the Long treatment. Therefore, we make the following prediction:

- Prediction 1: Subjects will stop rolling over their credit at higher values of the firm in the Short treatment than in the Long treatment.

Our second prediction is that we will see subjects stopping their rollover at values of the firm where, in case of a fire-sale, creditors would get back all of their investment (i.e.,$Min[1, \alpha F(y_{(t_{(b+\theta)})})] \geq 1$). This prediction also comes from He and Xiong (2012).

- Prediction 2: Credit freezes will happen even at values where a fire-sale would pay back the whole investment to all creditors.

4.5 Experimental Results

All session were run at the LEEPS lab of the University of California Santa Cruz, and all subjects were undergraduates from this institution. In total 72 subjects participated in the experiment, spread into 7 different sessions, and none played the game twice. In each session we had either 12 or 8 subjects for a total of 4,320 decisions (60 rounds × 72 subjects)[7]. From these observations we will ignore all stopping decisions for a value above $4, a value for which it is impossible to lose money in the Short treatment, and for which the probability of losing money in the Long treatment is <0.1%. In total we end up with 4115 observations. In Figure 4.5 we plot the CDF for both treatments along with the theoretical stopping threshold from He and Xiong (2012) (vertical dotted lines) where a horizontal orange line indicates the median value of the distribution. As we can see in Figure 4.5, the observed stopping values in the Long treatment are much lower than those in the Short treatment (Kolmogorov-Smirnov *p*-value = 0.000), which is in line with Prediction 1.

Figure 4.5: CDF plots for both treatments

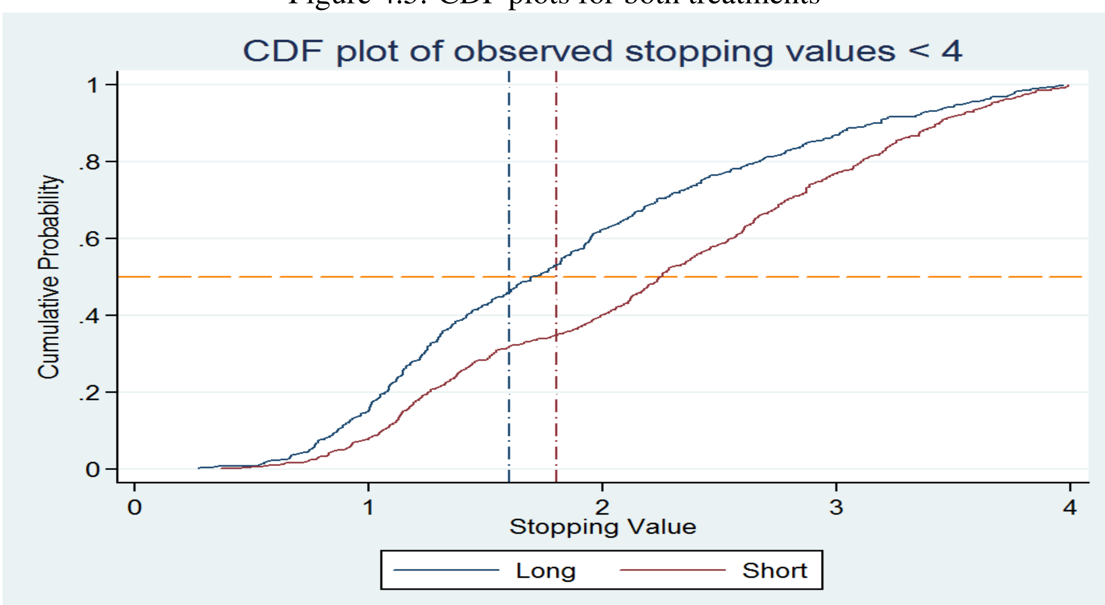

But Figure 4.5 is a CDF of the observed stopping values across all rounds in all sessions. To have a better description of the evolution of the stopping values during a session, we break the experiment into quarters (15 periods each quarter). As it is apparent from Figure 4.6 the stopping values across the two treatments continue to be

[7]Each session begins with some practice rounds whose results are not used in the analysis.

are widely separated in the fourth quarter (Kolmogorov-Smirnov corrected p-value = 0.008).

Figure 4.6: CDF Plots for Quarter 4

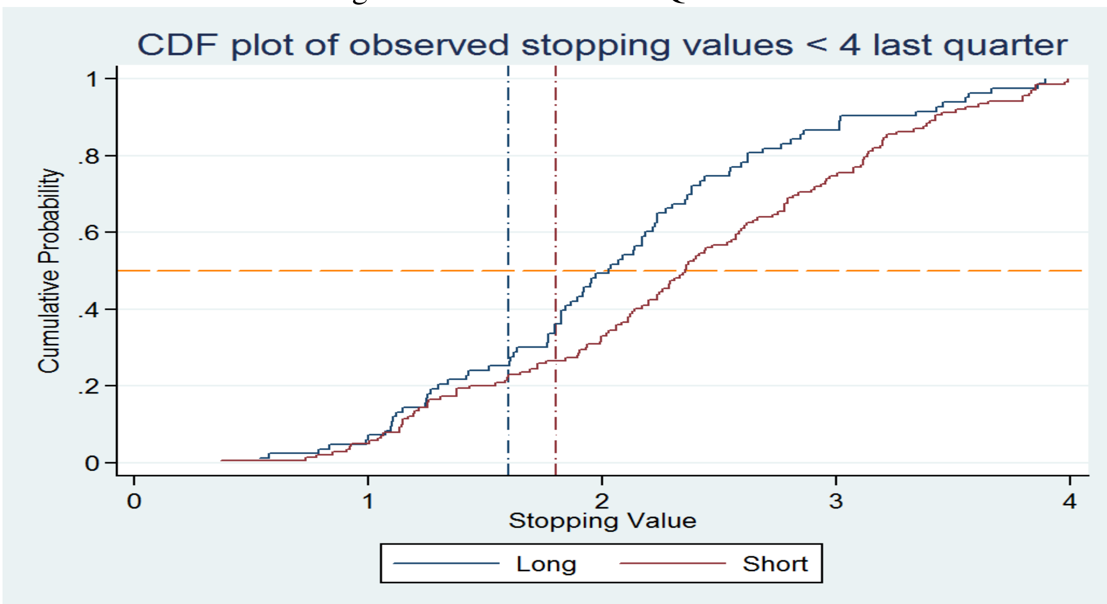

Also apparent from Figure 4.6 is that the CDF's of the stopping values for both treatments are significantly shifted to the right. If we compare the CDF's for the last quarter to the rest of the session, we get significant differences (Kolmogorov-Smirnov p-value = 0.029 in the Short treatment, and p-value = 0.001 in the Long treatment). It appears that as the sessions evolve, subjects try to undercut each other by stopping at higher values of the firm.

To check the effects of the quarters on the stopping values we run a simple linear regression, which confirms that contract length indeed has an effect on stopping values (Table 4.2)[8]. Our dependent variable is the stopping value (Stopvalue). "1.short" denotes a dummy variable for short maturity treatments and the other variables are quarter dummies.

[8]In Model 1, group dummies are not included in the regression, so clustering is at the quarter level. In Model 2, we include group dummies and cluster at the quarter-group interaction.

Table 4.2: Regression of stopping values and treatment

	(1)	(2)
	Stopvalue	Stopvalue
1.short	0.318**	0.932***
	(0.0756)	(0.109)
2.quarter	-0.253***	-0.205***
	(0.00388)	(0.0564)
3.quarter	-0.202***	-0.126**
	(0.00814)	(0.0587)
4.quarter	0.000974	-0.0457
	(0.00491)	(0.0633)
_cons	2.062***	1.766***
	(0.0527)	(0.0825)
N	1214	1214
adj. R^2	0.045	0.182
Group dummies	No	Yes

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

While the results of Table 4.2 confirm Prediction 1, namely, that Short contracts have significantly higher stopping values than Long ones, the surprising result is that the second and third quarters dummies are significant and *negative*. This means that for those quarters, the stopping values actually went down below the first quarter levels, while the dummy for the fourth quarter does not seem to be significantly different from the first quarter in the complete model (Model 2). A Kolmogorov-Smirnov comparing first and fourth quarter in the Short treatment confirms this p-value = 0.34, and p-value = 0.25 in the Long treatment.

When we plot the kernel density estimates for each quarter in both treatments (Figure 4.7) to have a visual idea of the distribution of the stopping values for both treatments, we see the surprising pattern of a text-book boom and crash of the credit market.

In both treatments (although clearer in the Long treatment) the first quarter is similar to a left skewed plateau, as stopping decisions are spread out across the whole strategy space with somewhat more incidence on the left tail. This suggests that there is a sort of "tâtonnement" learning process during which subjects presumably try to get better acquitted with both, the mechanics of the game and the strategy of other members of the group. It is only by the second quarter when subjects seem to realize that more money can be made from stopping at lower values, with the result that they start taking higher risks.

Figure 4.7: Kernel density by quarter for Long and Short treatments

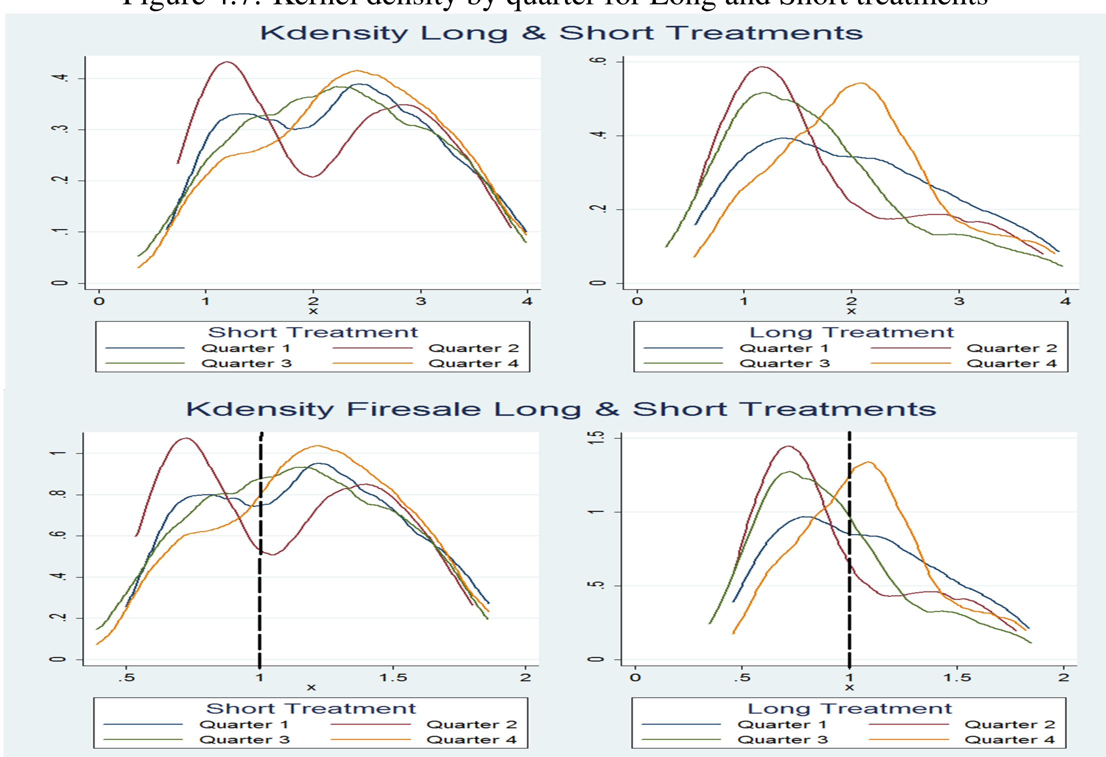

These risks are clearly seen in the lower half of Figure 4.7 where we plot the density estimates of the fire-sale values at the subjects stopping points. Looking at Quarter 2, the estimations show that a large number of subjects are stopping at values where the fire-sale value of the firm is well below the $1 break-even threshold. It is no surprise that this excessive risk-taking in the second quarter ends up in bankruptcies and, therefore, in losses for the subjects still invested, sparking, by the third quarter, a "full blown" panic that settles at stopping values above $2 in Quarter 4. (Appendix A has kernel densities broken into sixths of the session for a more precise description of the timing).

Table 4.3: Mean and SD of observed stopping values for each quarter by treatment

Quarter	Mean Long	Mean Short
1	1.92±0.898	2.26±0.901
2	1.35±0.887	2.11±0.924
3	1.48±0.839	2.23±0.893
4	2.02±0.774	2.35±0.851

In Table 4.3 we report the mean stopping values for each quarter and treatment. As we can see, this table confirms the bubble and crash story, and suggests why we didn't see a significant difference across the first and fourth quarter (Table 4.2): The rebound effect leaves the final stopping threshold very close to the starting average[9].

Given the parabolic evolution of the stopping values, we call these dynamics the "boomerang effect". Therefore, we can state:

- Result 7: *The dynamics of stopping values do not move in one direction, but rather have a "boomerang effect", resulting in a bubble followed by the crash of the credit market.*

While beyond of the scope of this paper, the analysis of these rich dynamics could be a starting point for future research in market dynamics, especially in the detection of bubbles and financial panics. Yet, what is clear from both the kernel densities and in

[9]A Mann-Whitney test confirms that there is no significant difference between the first and last quarter, yet second and fourth are significantly different in both treatments.

the empirical CDF's is that a significant number of subjects are stopping their credit at values that are above the fire-sale break-even point $\alpha F(y_t) = 1$. Therefore:

- Result 8: *Overall, 60% of the decisions to stop rolling over the credit are made even when firms have strong fundamentals, and can pay back the entire investment to all subjects even in the case of a fire-sale (i.e., Frantic runs occur).*

That frantic runs are common is particularly interesting when recalling that the unexpected freeze of securitized credit was one of the main reasons why the 2007 crisis was so damaging. Because frantic runs were so unexpected, firms had no contingency plan for a sudden freeze of the short-term credit market, making many firms unable to meet their liquidity needs.

The estimates above are useful as a first approximation to the qualitative results of our experiment on the short-term credit market. But, so far, we have not used all the data gathered in the experiment, since we have ignored the information contained in the decisions *not to stop rolling over the credit*. By not taking into account these non-stopping decisions our mean estimates are necessarily biased upwards.

We can overcome this bias by using the product limit estimator (Kaplan Meier (1958)), which corrects the bias in the estimations and, most importantly, allows the estimation of the survival curve for each treatment and lets us study the estimated stopping behavior of all the subjects across the whole range of firm values.

4.5.1 Hazard Rates and the Product Limit Estimator:

For any kind of policy analysis (and especially when dealing with financial panics) the state of the economy is a crucial factor. In our experiment we can take the value of collateral as a proxy for the state of the economy, so that the lower the value of the firm, the worse the economy is and vice-versa. To assess the behavior of our subjects for any state of the economy we will turn to estimating the hazard functions governing our experiment's dynamics, and to do this we will need to use the product limit estimator.

The product limit (PL) estimator is a non-parametric Maximum Likelihood Estimator of the distribution, which is adapted to dealing with censored data as in our case. As mentioned, the previous analysis ignored the decisions not to run. The PL estimator takes these non-stopping decisions as "censored" observations in the sense that we do not observe a subject stop rolling over his credit, because the value of the firm never got to his stopping threshold. For example, for a given threshold of subject i in round j, (t_{ij}), we only observe his stopping decision if the value of the firm for that round (y_j) reaches his specific threshold (i.e. $Min[y_j] \leq t_{ij}$). If this condition is never met, then we are left with a censored observation.

But while it is usually straightforward to use the PL estimator, in our dataset we have left-censored data, which is not standard. To deal with it, we need to "flip" the data and work with their mirror image. To flip the data, we have to find a constant, S, large enough such that $S > maxy_{ij}$ for all subjects i and rounds j. As explained with some detail in Appendix C, we decided to subtract all values from 4.

In Figure 4.8 we present the hazard and cumulative hazard estimates for the flipped data of each treatment. The hazard function can be understood as the "probability[10]" that a subject that has not stopped rolling her credit will do so within an infinitesimally small range of firm values. To be more precise, define the instantaneous hazard rate as a measure of the probability that a subject will decide to stop rolling over the credit

[10] It actually is the ratio between the probability density function of the event (running), and the survivor function.

within the (limiting) interval $\triangle y$, given that he has not yet stopped rolling over his credit:

$$h(y) = \lim_{\triangle y \to 0} \frac{e[y, y + \triangle y]/N(y)}{\triangle y} \tag{4.4}$$

Where $e[y, y + \triangle y]$ is the number of observed rollover stops in the interval $[y, y + \triangle y]$, and $N(y)$ is the number of subjects at risk for value y. It is clear that if we do not take into account the censored observations, then $N(y)$ will be higher, bringing down the real hazard rate for the value of the firm y, and consequently biasing our results.

The Nelson-Aalen cumulative hazard rate is a non-parametric estimation of the cumulative hazard. Notice how the cumulative hazard can go above the value of 1, which means that for that value of the firm, if it occurred, a subject would have stopped rolling the credit more than once.

In Figure 4.8 we present both the estimated hazard function and the Nelson-Aalen estimates for the cumulative hazard of both treatments. What is striking in these graphs is how the hazard of the Long treatment starts below that of the Short treatment, but ends up way above it, after both hazard functions cross. This is extremely important as it shows that the hazards are not proportional, and thus there is an interaction between the treatment hazard ratios and the state of the economy. What we see is that when the economy is in good shape (the value of the firm if high), subjects are more willing to rollover the credit under long contracts. Yet, when the value of the collateral is close to "breaking the buck", then the hazard rate of the Long market explodes, overtaking that of the Short market. This suggests that the optimal approach to short-term credit regulation should be a *dynamic policy* that changes with the state of the economy[11].

[11] Of course this policy might be too expeinsive to implement. In this case it si apparent that overall a Short policy would be preferred to a Long policy.

Figure 4.8: Hazard estimates for both treatments

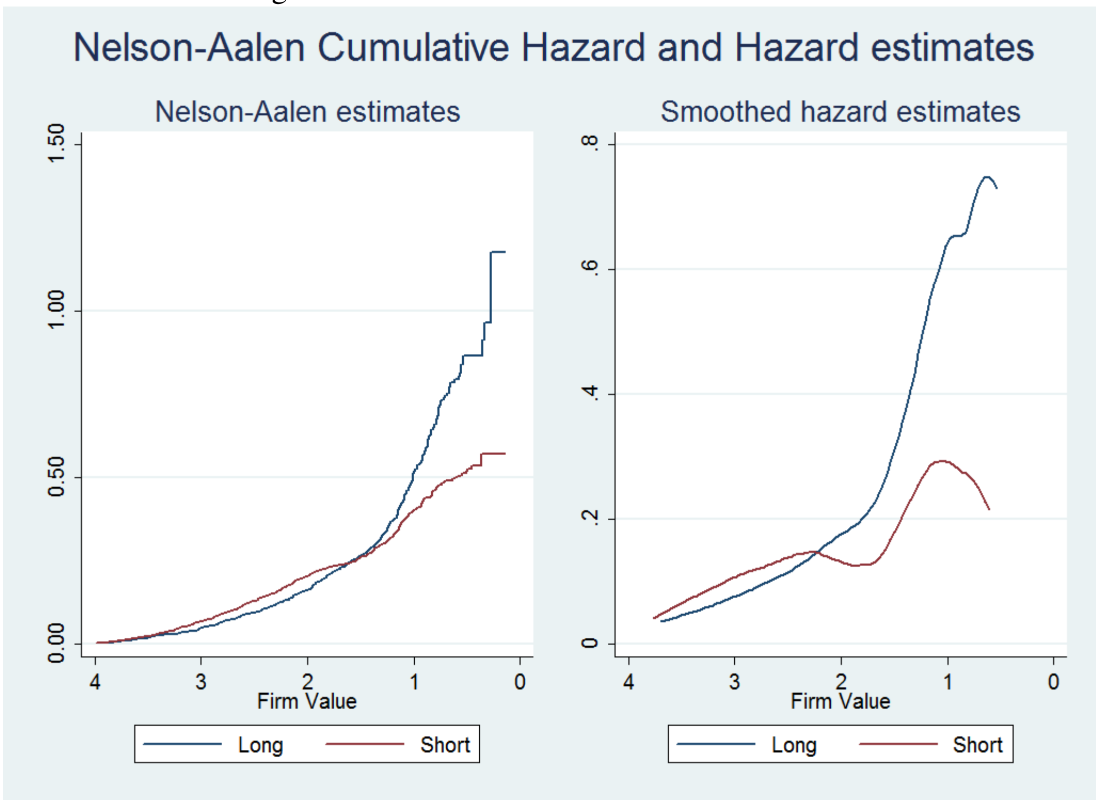

A Fleming-Harrington[12] test with all the weight put on the left tail (*p*-value =0.045) confirms that the differences across treatments are significant[13]. Therefore, even when the censored information is included in our analysis we see a highly significant difference between both treatments. Furthermore we see that the hazards are not proportional, but rather that their ratio is state-dependent. This fits with the bimodal shape of the densities that we saw in Figure 4.7, where the Short treatment had a peak of stopping decisions for high values of the firm.

A Tobit model estimation confirms the Fleming-Harrington tests by showing a highly significant treatment dummy (Short), and still showing the “boomerang effect”

[12] Precisely because the hazard estimates cross, our dataset violates the proportional hazards assumption, and the Fleming-Harrington test is the most appropriate.

[13] A Fleming-Harrington test putting equal weight in both tails, and another test with all weight on the right tail will also show highly significant differences (p-value = 0.000 in both cases).

Table 4.4: Tobit estimation of the stopping values

	(1)
	Stopvalue
Short	0.726***
	(0.0209)
2.Quarter	-0.197**
	(0.0832)
3.Quarter	-0.178*
	(0.105)
4.Quarter	-0.248*
	(0.134)
cons	0.0670
	(0.108)
N	3997
adj. R^2	
Group dummies	Yes

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

across quarters (2.Quarter, 3.Quarter, 4.Quarter), as we can see in Table 4.4.

Unfortunately, to have precise quantitative estimates of the PL estimator we need to work with a subset of our data. This is due to the heavy censoring which will bias the PL in the direction of the censored data points (Moeschberger and Klein (1985, 2003), Miller (1983)). Moreover, if there are censored points beyond the last precise observation (i.e., the last observed stopping value), then the bias will be even stronger (Moeschberger and Klein (1985)). With almost 70% of our data censored, and with some of these censored values to the right of the highest[14] observed stopping value, we will need to find a rule to reduce the censoring in our sample.

To avoid the bias described in Moeschberger and Klein (1985) we drop those ob-

[14]Remember that we have flipped the data.

servations that are censored beyond the latest observed stopping value (in total 38 observations). Then, following Anderson *et al.* (2009), we eliminate all observations in those sessions where the minimum value of the firm was not low enough. In our case this means all observations for rounds where the minimum value was (strictly) greater than \$0.9 (in total 3020 observations). Finally we also need to drop those subjects who run only 2 or fewer times in the whole session (155 observations). So in total we are left with a subsample of 902 observations, where 448 are stopping decisions and 454 are censored observations (almost a 50% ratio of censored data).

Running the PL estimate on this new subsample we still have a significant effect of the treatment variable (Fleming-Harrington with all the weight on the left tail has p-value = 0.011), and the hazard curves continue to cross in the same way as with the full sample. Therefore, while the bias produces quantitative differences across the full sample and the sub-sample, it does not have any effect on our qualitative results.

- Result 9: Both the subsample and the full sample present crossing hazard curves, with high hazard rates in the Short treatment for high values of the firm, but even higher hazard rates for the Long treatment when the value of the firm is near the break-even point.

Finally, we use the PL estimate to calculate the mean stopping value for each quarter. In both the full sample and the subsample the Short treatment has mean stopping values that are (in general) lower than in the Long treatment. Table 4.5 presents the results of the subsample estimates with their bootstrapped standard errors.

Table 4.5: Subsample mean and bootstrapped SE for each quarter

Quarter	Mean Long	Mean Short
1	0.84±0.04	0.74±0.06
2	0.92±0.04	0.90±0.05
3	1.01±0.07	1.05±0.08
4	1.15±0.09	0.90±0.12

Therefore, once censored observations are added, our analysis reveals two things: First, that on average the Short treatment has lower stopping values than the Long treatment. Second, that while the Short treatment might have lower stopping values, its hazard function is higher for high values of the collateral. That is, when the value of the firm is high, a significantly lower number of subjects decide to stop rolling over their credit in the Long treatment. But, once the value of the firm gets near the fire-sale break-even point, the hazard function of the Long treatment shoots upwards, making the market extremely prone to freezes.

- Result 10: The PL estimator shows that Short contracts have a lower mean stopping value than Long contracts.

4.6 Conclusion

In the *Handbook of Experimental Economics* (Kagel and Roth, eds., 1995), Al Roth explains in his Introduction why experiments are run. He mentions three reasons: To test a theory, to find new data that potentially complement the theory, and finally to offer policy advice ("whisper in the ear of princes"). This is exactly what we try to do in this paper.

To be specific, we try to compare the effects that maturity lengths have on the functioning of ABCP markets. Building on the theoretical model of He and Xiong [2012], we compare two markets where the only difference is the maturity of ABCP. Our main result is that while, on average, ABCP markets with short contracts are less prone to freezes (Result 4), the optimal policy should be state-dependent, favoring long contracts when the economy is in good shape, and allowing for short-term contracts when the economy is in a recession (Result 3). This latter result comes from estimating the survival curves of our markets, and observing the hazard functions of the Long and Short treatments cross at a value near the "break-even" point, with the Long market shooting up in risk when the value of collaterals is low.

Our second result (Result 2) is reporting, for the first time in a laboratory setting, "frantic runs", which are runs on firms that are able to pay all their debts even in a fire-sale (i.e., runs on firms with strong fundamentals). This is an important result as it cannot be observed in the canonical static models of financial panics. As Bernanke [2008] put it, the loss of access too secured borrowing was "surprising" and firms had no contingency plans prepared to face this situation. Being able to reproduce them in the lab is a first step to better understand how they work, and how they can be prevented.

Our last result is reporting rich experimental dynamics, with a consistent bubble-and-crash pattern across our sessions (Result 1). And, while analyzing the learning dynamics behind these results is beyond the scope of this paper, we do believe that the structure of our experiment can be an interesting starting point to analyze sophisticated learning in any kind of dynamic markets.

Finally, we hope this is the first of a series of papers on financial panic experiments, and that a buildup of similar papers might allow economists to whisper, with more confidence, in the ear of financial princes.

Appendix A

A Tale of Two Tails

A.1 Details on session structure

The treatment ordering for each session as well as the total number of subjects per session in Table A.1

Table A.1: Number of Subjects Participating in Each Type of Session

Treatment Order/Town	Barcelona	Santa Cruz
N2H	18	21
N2L	18	21
(H-1)2(L-1)	-	33
(L-1)2(H-1)	-	48
L2H	-	12
2NL	18	-
2NH	18	-
H2N	15	-
L2N	15	-

In Table A.2 we present the total number of actual decision-maker observations for each treatment:

Table A.2: Number of Decision-Maker (B) Observations

	Barcelona	Santa Cruz	Total
N	33	14	47
H	17	11	28
L	17	11	28
H-1	-	27	27
L-1	-	27	27

A.2 Ordering Effects

Due to a miscommunication between the Barcelona and Santa Cruz labs we have a very unbalanced number of sessions with H as the first round treatment (5) compared with third round H treatment (22). This unfortunately pollutes the ordering effects for the H treatments as a 2 tailed Fisher Test comparing first round treatments against other rounds in the experiment shows.

	$0	$1	$2	$3	$4	$5	$6	$7	$8	$9	$10
N	.75	.89	.34	.67	.17	1.0	.76	.49	.35	.92	.62
H-1	.70	1.0	1.0	.09*	.62	1.0	1.0	1.0	1.0	1.0	1.0
H	.09*	.03**	.01**	.16	.23	1.0	1.0	1.0	1.0	.13	.06*
L	.57	1.0	.35	.68	.40	1.0	1.0	1.0	.43	1.0	.43
L-1	1.0	.44	.69	1.0	.05*	.54	.54	1.0	.66	.66	.44

Table A.3: Two-Sided Fisher P-values Comparing First Round Treatments to all Other Treatments

While most treatments have no ordering effects, the LHT of the H treatment seems to be significantly affected by ordering. If we look at Graph A, we can see that while last round pattern of acceptances does look like those in the rest of treatments, first round H acceptances looks pretty random. As mentioned, we believe that this is due to the low number of observations of H in the first round, and that if we had more observations we would see no ordering effects.

Figure A.1: Acceptance Rates for H for First (n=5) and Third (n=22) Round

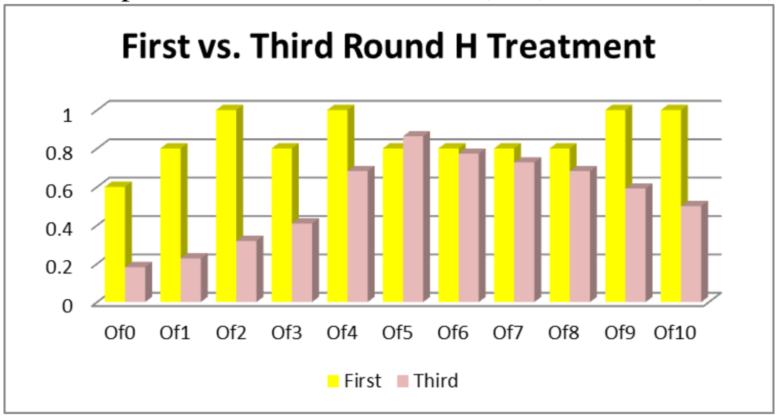

A.3 Spearman Rank Correlation

Table A.4: Spearman Rank Correlation Results for LHT and RHT of L-1 and N-1 treatments

	LHT (L-1)	LHT (H-1)	RHT (L-1)	RHT (H-1)
Spearman Rho	0.9856	1.000	-0.9710	-0.7495
Prob > |t|	0.0003	0.000	0.0012	0.059

A.4 Two-sided Fisher

	$0	$1	$2	$3	$4	$5	$6	$7	$8	$9	$10
L-1 vs. H-1	1.00	0.78	0.77	1.00	1.00	0.61	1.00	0.46	1.00	1.00	0.22

Table A.5: Two-Sided Fisher P-values Comparing First Round Treatments to all Other Treatments

A.5 Instructions

3UG:

Welcome! This is an economics experiment. You will be a player in many periods of an interactive decision-making game. If you pay close attention to these instructions, you can earn a significant sum of money. It will be paid to you in cash at the end of the last period. It is important that you remain silent and do not look at other people's work. If you have any questions, or need assistance of any kind, please raise your hand and we will come to you. If you talk, laugh, exclaim out loud, etc., you will be asked to leave and you will not be paid. We expect and appreciate your cooperation today. What's this all about: This experiment has three different rounds. Before each round the specific rules and how you will earn money will be explained to you. In each round there will always be three types of players: A, B and C. You will be assigned to a type in Round 1 and will remain this type across all three rounds. Only one of the three rounds will be used for the final payoffs. This round is chosen randomly by the computer. The outcomes of each round are not made public until the end of the session (i.e. after round 3). Each round the groups are scrambled so you will never make offers or decide for the same player in two different rounds.

Round 1:

The first thing that you will see on your screen is your player type.

You will then be assigned to a group consisting of three players: an A type, B type and C type.

Player A will be endowed with $10 which he will split with player C. In order to do so Player A will have to input the amount he is willing to offer Player C. Player A will only be able to make integer offers (full dollars), so A will not be able to break its offer into cents.

While player A is deciding how much to offer player C, player B will be filling out a binding "strategy profile". The strategy profile has an "accept or reject" button for each potential offer from A to C (from $0 to $10). Player B's binding decision to accept or reject A's offers to C will be done before he knows the actual offer made by A.

A's decision: How to split an endowment of $10 with Player C by making him an offer between $0 and $10. If the offer is of $X, A will be keeping for himself 10-X.

B's decision: Before knowing the offer from A to C, B will fill a binding "strategy profile" deciding whether he accepts or rejects every potential offer from A to C. This decision is made without knowing the offer from A to C.

Figure A.2: Diagram 3UG

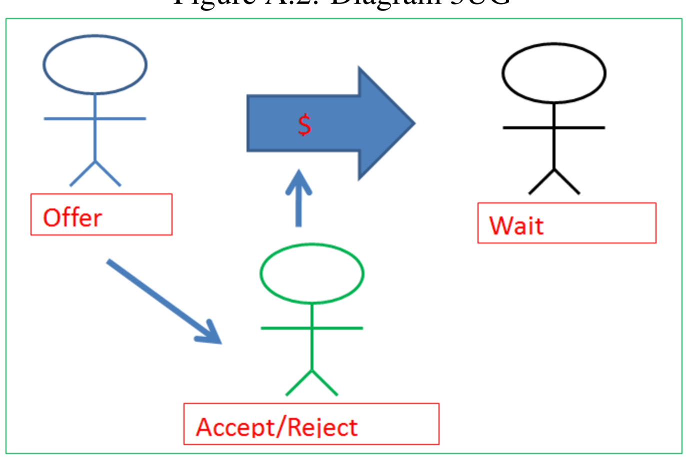

It is very important for A to realize that he is going to write the amount he wants to offer C and not how much he wants to keep.

Payoff for Round 1:

If B accepts the offer from A to C, then they split the $10 as suggested by A.

If B rejects the offer from A to C, then both (A and C) get $0.

B will get paid $3 no matter what is the outcome.

Timing and Payoffs:

1. B fills a strategy profile with all potential offers from A to C.

2. A decides how much to offer C (say X)

Figure A.3: Diagram of Payoffs

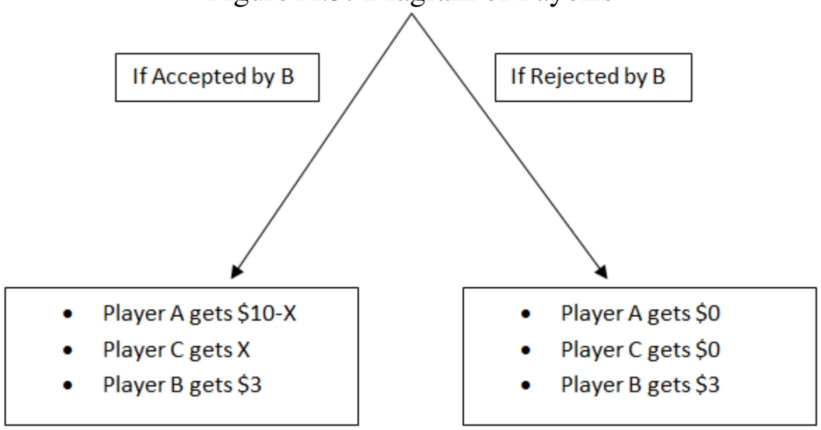

Round 2:

As mentioned at the beginning of the experiment you will keep your player type across the whole session. So A players are still A, B are B and C are C.

In this round type A players will be endowed with $20 and will have to make TWO offers:

1. How to split $10 with player B.

2. How to split $10 with player C.

As in Round 1 a binding "strategy profile" will be filled by B and C players before they know the offer made to *them*.

It is very important to notice that B and C players are making decisions concerning their own payoffs.

A's decision: How to split $10 with B and how to split $10 with A.

Each offer is independent. So the outcome of the offer to B has no effect on the outcome of the offer to C.

Payoffs for A will be as in Round 1 (if he offers X and the offer is accepted he gets $10-X, if the offer is rejected both him and the rejecting player get 0).

B and C players will get paid X or 0 depending if the accepted or rejected the offer made directly to them.

In order to make payoffs equitable for this round, A's payoff for this round will be chosen at random between one of the two outcomes (offer to B and offer to C). B and C's decision: Before knowing the offer made to them by Player A, B and C will fill a binding "strategy profile" deciding if they accept or reject *every potential offer made directly to them.*

If the offer from A is accepted, then the split is done as proposed by A. If the offer is rejected both the receiver and A get $0 as the outcome for this round.

Figure A.4: 2UG Diagram

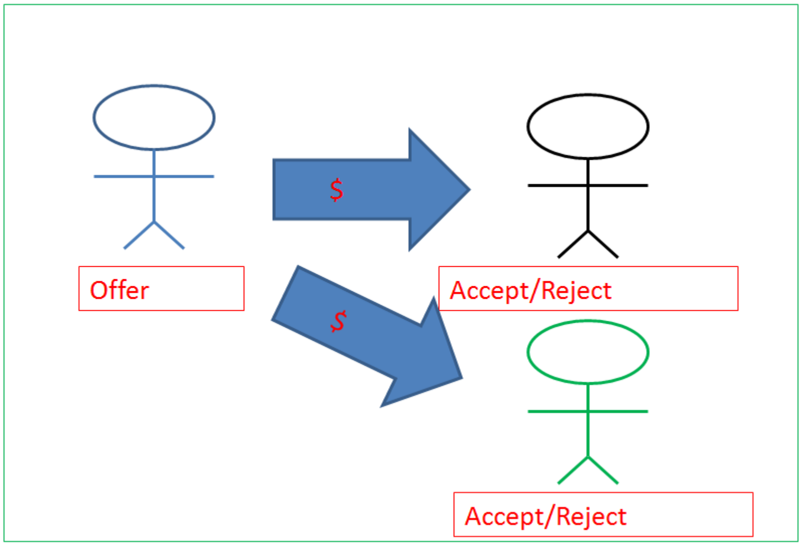

Timing and Payoff for Round 2:

1. Each receiver fills a strategy profile with all potential offers that A could make them.
2. A decides how much to offer C and B (say X)
3. Payoffs for B and C will be the outcome of their particular game with A.
4. To make outcomes equitable, the computer will choose randomly one of the two outcomes to be A's payoff for the round.

For each offer made from A to the other members of his group:

Figure A.5: 2UG Payoffs

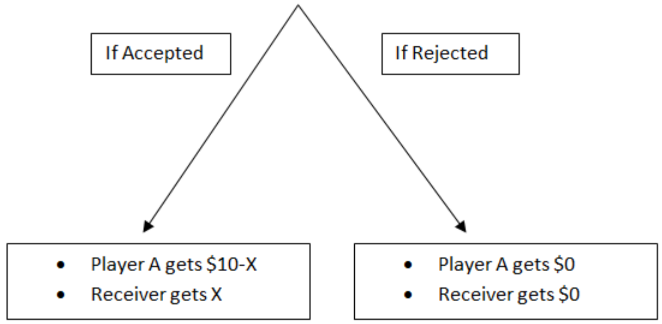

Round 3:

As mentioned at the beginning of the experiment you will remain your player type across the whole session.

This round is very similar to round 1. You will now be re-scrambled into groups of three subjects (one A, one B and one C subject).

A will be endowed with $10 and must decide how to split them with C.

B's role is exactly the same as that in round 1: Before knowing the offer from A to C, B will fill a "strategy profile" deciding whether he accepts or rejects *every potential offer from A to C*.

If the offer from A to C is accepted by B, then the split is done as proposed by A. If B rejects the offer, then both A and C receive $0 for this round.

B's payoff in this round is a flat $12 fee, whatever his decision and outcome of the round. So, the only change between Round 1 and Round 3 is that player B, is getting paid a different amount.

Figure A.6: 3UG (H) Diagram

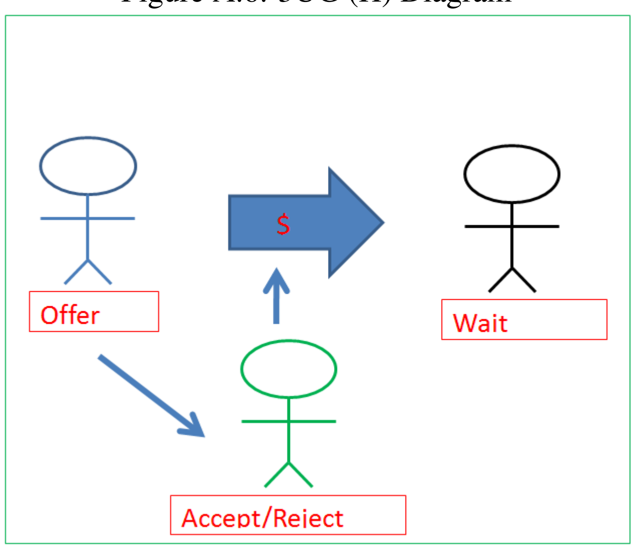

Timing and Payoffs:

1. B fills a strategy profile with all potential offers from A to B.

2. A decides how much to offer C (say X)

Figure A.7: Payment Diagram 3UG (H)

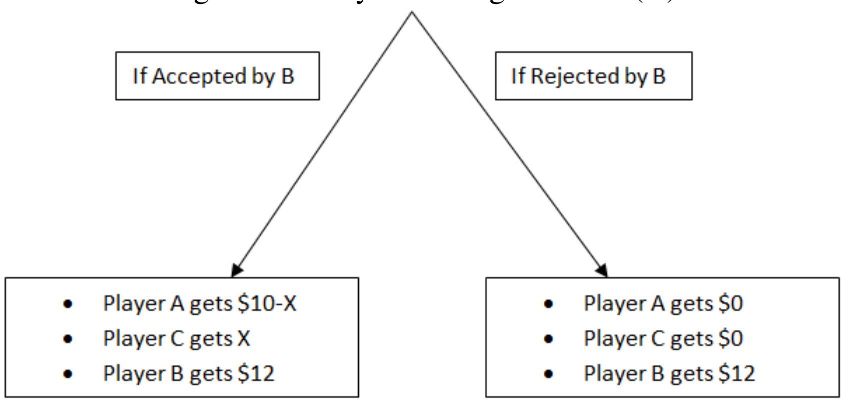

Appendix B

Survivor Curve and Internet Revenue

B.1 Details

In Table B.1 we present the different nuisance levels for each activity, and report the number of observations at each level.

Table B.1: Nuisance Values and [and Numbers of Observations].

Inconvenience	Level 1	Level 2	Level 3	Level 4	Level 5	Level 6
Pi volume	30[11]	40[20]	50[12]	60[22]	70[13]	80[20]
Pop-up	30[11]	25[19]	20[12]	15[22]	10[13]	5[21]
Jitter	.10[11]	.13[21]	.16[12]	.19[22]	.22[13]	.25[19]
Reading	.15[11]	.18[21]	.21[12]	.24[21]	.27[13]	.30[20]
SAT	.6[9]	.9[8]	.12[11]	.15[11]	.18[12]	.21[12]
Movie Pay	1[11]	5[19]	9[12]	13[22]	17[13]	23[21]

In Table B.2 we present the robustness checks for the probit model of Table 3.2 , and present the results for ordering ($i.order$), and a series of dummies $pibigpop_{i,j}$, $popbigpi_{i,j}$, $readbigsat_{i,j}$, $satbigread_{i,j}$ that test the effects of having similar activities with different levels of nuisance. For example, $pibigpop_{i,j}$ ($popbigpi_{i,j}$) is a dummy for the case when the nuisance for Movie/Pi (Movie/Pop) is bigger than that for Movie/Pop

(Movie/Pi); similarly $readbigsat_{i,j}$ ($satbigread_{i,j}$) is a dummy for the case where Reading (SAT) has a bigger nuisance level than SAT (Reading).

The results show that ordering has no statistical effect on the decisions of subjects, while different levels of inconvenience for similar activities seem to have an effect when Movie/Pi has a bigger nuisance than Movie/Pop (note that we only have 8 cases of this). Finally, for the 16 session dummies only one is significantly different (at the 5%) from our baseline.

We conclude thus that our results are robust, and even if we have a few dummies with significant effects, these are probably due to small sample bias.

Table B.2: Switching Probit Model, Continued.

	(4) Switch
2.order	-0.315
	(0.223)
3.order	-0.552
	(0.382)
4.order	-0.516
	(0.334)
5.order	-0.0737
	(0.287)
6.order	-0.248
	(0.356)
7.order	0.443
	(0.408)
1.readbigsat	-0.249
	(0.261)
1.satbigread	0.0678
	(0.515)
1.popbigpi	-0.302
	(0.408)
1.pibigpop	1.155**
	(0.536)
cons	-1.282***
	(0.221)
N	636
Session Dummies	Yes

B.2 Other activities

In this appendix we list all the activities that were not described in detail in the methods section.

Movie/Pop: Subjects were presented with a menu of two video clips (an interview of Zack Galifianakis by Letterman, and a clip on how to do crossover moves in basketball). After 100 seconds of visualization, a 15 second long pop-up would appear on the screen. This pop-up would partially cover the video clip, and have flashing colors; moreover, while the pop-up was on the screen, the movie clip would continue playing in the background but with no sound. The unit of the nuisance $x \in [5,30]$ is the number of seconds between consecutive pop-ups, e.g., if a subject was assigned a nuisance level of $x = 5$, then she would have a 15 second pop-up every 5 seconds. If the subject decided that the nuisance was too big, then she could switch to the bland activity which, as in all movie activities, was a video of gentle waves breaking at La Jolla beach. Once a subject switched to the bland activity she would remain there until the end of the round. Rounds lasted 8 minutes. Note on wave watching: The bland activity for all movie activities is "watching waves." We decided to use this video because as it has no plot, that is, its "replay value" is very high, allowing us to reuse it with almost no loss in its (relative) attractiveness.

Slug: Slug is a version of the classic video game Snake, which incorporates the university colors (the mascot of UCSC is a Banana Slug, and its color are blue and yellow). Snake was a popular arcade game in the 1970's but gained world-wide acceptance in 1998 as it became the standard pre-loaded game in Nokia phones. The game has been used as "Easter egg" by both Youtube and Gmail. In this game the objective is to get "food," which corresponds to colored pixels that appear at random points of the enclosed "playing space." Each time the player gets to food she earns points, but the slug increases in size, making it harder to maneuver. To get to the food subjects control the slug with the keyboard arrows. If the slug bumps into the walls of the enclosed playing space, or if it hits itself, the player loses. Losing has no cost in points, the subject just need to restart the game by pressing the refresh button (F5 on the keyboard), and the game starts over with the same amount of accumulated points. As mentioned, points are awarded by getting to food; 10 points for regular food and 40 points for bonus food. The difference between these two types of food is that bonus food only stays on screen during 10 seconds, while normal food is there until eaten. Food is color coded, with bonus food being yellow, and regular food blue. Each point was worth $0.01. The jitter nuisance would start 50 seconds into the round, and involves a random turn (left or right) each pixel with probability $x \in [.10, .25]$. The bland activity towards which subjects could switch was the same exact game without the jittering nuisance, but paying only one fourth of the amounts in the original activity (i.e., 10 points per bonus food, and 5 points for each piece of regular food "eaten"). Each round lasted 7 minutes.

Read: Subjects are given a menu with a series of articles from the New York Times (an article on the Proposition B for LA county, an article on veterans of the Iraq war coming back to the US, and an article on fee increase at the UC system). The nuisance $x \in [.15, .30]$ is the (independent) probability for each letter of being dropped. The first 15% of the text would be nuisance free. On the other hand, the text was presented broken into paragraphs. To ensure that subjects actually read the text, they could only move to the next paragraph by clicking a "next" button that would appear 10 seconds after the start of every new paragraph. The bland activity was counting bits, which presented subjects with a binary string of 15 digits, and asked them to count how many 1's were in the string. If the answer was correct, then a new string was generated. If the answer was wrong the subject would be given a new opportunity until he answered correctly. This would last until the end of the round, which was 6 minutes long.

SAT: Subjects could pick between two different texts taken from an SAT practice web-page. The text would be presented to subjects along with only one of the 8 multiple choice questions they had in this round. All answers were final, and once a choice was made the next question would appear, with no way of going back. This was a paid activity and each correct answer would pay \$0.40, while each incorrect answer would penalize \$0.10. The nuisance for this activity was letter dropping, and worked exactly as in the Read activity. In this case each letter was dropped with probability $x \in [0.06, 0.21]$. The bland activity was the same task with all the letters, but paying one fourth (i.e., \$0.10 for each correct answer and -\$0.02 per incorrect answer). If a subject decided to switch, she would not start over all the questions, but would start the bland activity at the same question where she switched to activity B.

B.3 Instructions

Upon entering the lab subjects were read an initial set of instructions that described the structure of the experiment but did not give any details on the activities or inconveniences they would encounter; subjects were told that detailed instructions would be given before each round.

We opted for this format of instructions because our experience tells us that subjects are prone to getting distracted with long instructions. We also believe that it is beneficial for them to have a little break between rounds to "start fresh." For the sake of conciseness we will change the format of instructions in this appendix, and instead of presenting instructions for each round in separate pages, we will list them one after the other.

B.3.1 General instructions

Welcome! This is an economics experiment. You will be a player in many periods of an interactive decision-making game. If you pay close attention to these instructions, you can earn a significant sum of money. It will be paid to you in cash at the end of the last period. It is important that you remain silent and do not look at other people's work. If you have any questions, or need assistance of any kind, please raise your hand and we will come to you. We expect and appreciate your cooperation today.

The Experiment: This experiment will have six different rounds. In each round you will begin with an enjoyable activity that we refer to as Activity A. At any time during the round you can switch to another activity, Activity B. The experimenter will announce the A and B activities for that round before it starts.

At the same time, the experimenter will also announce an "annoyance" that will accompany Activity A at some point during that round. If, after experiencing the annoyance, you think you would prefer Activity B, then simply click the button on your screen. It will immediately switch you to B, where you will remain for the rest of the round. You will never be interrupted by any annoyance in Activity B. Key points:

- You will start each round participating in an A activity.
- A activities will be interrupted by specific annoyances (announced before the round).
- At any point during the round you can switch from activity A to activity B (announced before the round)
- You can switch from A to B, but never from B to A.
- B activities do not have any interruptions.

Also note:

- Some rounds include a paid Activity and some do not.
- You automatically get to experience an A activity each round. To make sure that you are familiar with all with B activities, you will practice with all of them before the experiment starts.
- For some of the activities the audio output is needed. Please check if you have headphones attached to your computer. If you have your own, feel free to use them. You will be able to adjust the volume through the “speaker icon” on the upper right corner of your screen.
- Do not start Activity A until the experimenter announces that it is time to do so.

B.3.2 Specific activity instructions

Round Movie/Pi (8 minutes):

Activity A: Watching a video. You will choose it from a menu that will appear on screen.

Annoyance: While watching the video, at some point you will start to hear a computerized voice reading the first few thousand digits of the decimal expansion of $\pi = 3.14159\ldots$ This will continue at the same volume until the end of the round, or until you switch to activity B.

Activity B: Watching a video of waves breaking at La Jolla beach. This is not a paid round.

Round Movie/pop (8 minutes):

Activity A: Watching a video. You will choose it from a menu that will appear on screen.

Annoyance: While watching the video, at some point a pop-up will appear on your screen and mute the audio. These pop-ups are 15 second long, and will appear at regular intervals on your screen. The time remaining is shown on the pop-up.

Activity B: Watching a video of waves breaking at La Jolla beach. This is not a paid round.

Round Slug (7 minutes):

Activity A: Playing a game called “Slug”, very similar to the popular game “Snake.” Use your arrow keys to control a hungry slug. The slug gets longer as it eats food, and you earn points:

- Regular food (Blue Pixel): will stay on screen until you eat it, each piece that you eat which gives you 20 points.
- Bonus food (Yellow Pixel): gives you 40 points, will appear randomly and only lasts for 10 seconds on screen, if you don’t eat it during this time it disappears.

Your slug will “die” whenever it collides either with an edge of its rectangle or with its own body. But the points you earned are stored and accumulated, and you can begin again with a new slug. Just hit the refresh page key (F5) and the game will restart with a new short slug.

Annoyance: At some point the slug starts to “jitter.” That is, with some probability, it will change direction randomly each time it reaches a new pixel. The jitter rate (probability) will remain the same for Activity A the rest of the round.

Activity B: Playing the same game, “Slug,” but with two differences:

- The slug will not jitter
- You will earn points at 1/4 the previous rate: 5 points per blue pixel, 10 per yellow.

Round Read (6 minutes):

Activity A: Reading newspaper articles. You will choose one from a menu, and the text will appear on your screen. The text will be broken up into different pages. After

10 seconds "next page" button will appear. Just click the button to move to the next page. On the last page, please press the button to indicate when you are done reading the article.

Annoyance: In this activity the annoyance will be that some letters of the text will be missing. With a certain probability letters will be dropped from the article. This will apply to all the text, except the very beginning. As usual, press the button if you would rather go to the B activity than continue trying to read the article with missing letters.

Activity B: Counting the number of 1's in a string of 0's and 1's. If enter the correct number, then you will get 1 point and a new array of numbers will be randomly generated for you to count. If your answer is incorrect, then you will not get any points and will still have the same array of binary numbers for you to count. There is no limit to the number of attempts for each array. This is a activity — you get no money for the points!

Round SAT (8 minutes):

Activity A: Answering SAT questions. You will pick one of two sets of multiple choice questions. You will get paid 40 points per correct answer and will lose 10 points for incorrect answers. Your points are accumulated as you go and are shown on the screen. You will get to see 1 question at a time which you will be able to answer. Once you have answered a question you will NOT be able to change it, so you choice is always final.

Annoyance:: Except for the first question, some letters of the text will be missing. With a certain probability each letter will be dropped from each SAT question. As usual, you can press the button that takes you to activity B at any moment of the round.

Activity B: In this case the B activity will be the same SAT text, except it will have all the letters in the text, and it will pay you 10 points per correct answer and subtract 2 points if the answer is incorrect. If you switch to activity B you will start at the same point where you decided to change from A to B. So, for example, if you decided to switch at question 3, you will start activity B at question 3. Note that you can come out with negative earnings from this activity.

Movie/Pay (8 minutes):

Activity A: In this round you will be offered to pick from a series of clips to watch. On top of this you will be endowed with 500 points for you to keep.

Annoyance: Some seconds into the video you will be asked to pay a fee (in experimental points) if you want to continue watching the video.

Activity B: If you don't pay, the video will switch to waves breaking at La Jolla beach.

Appendix C

That's how we roll

C.1 Detailed Kernel Density

Figure C.1: Density estimates for both treatments

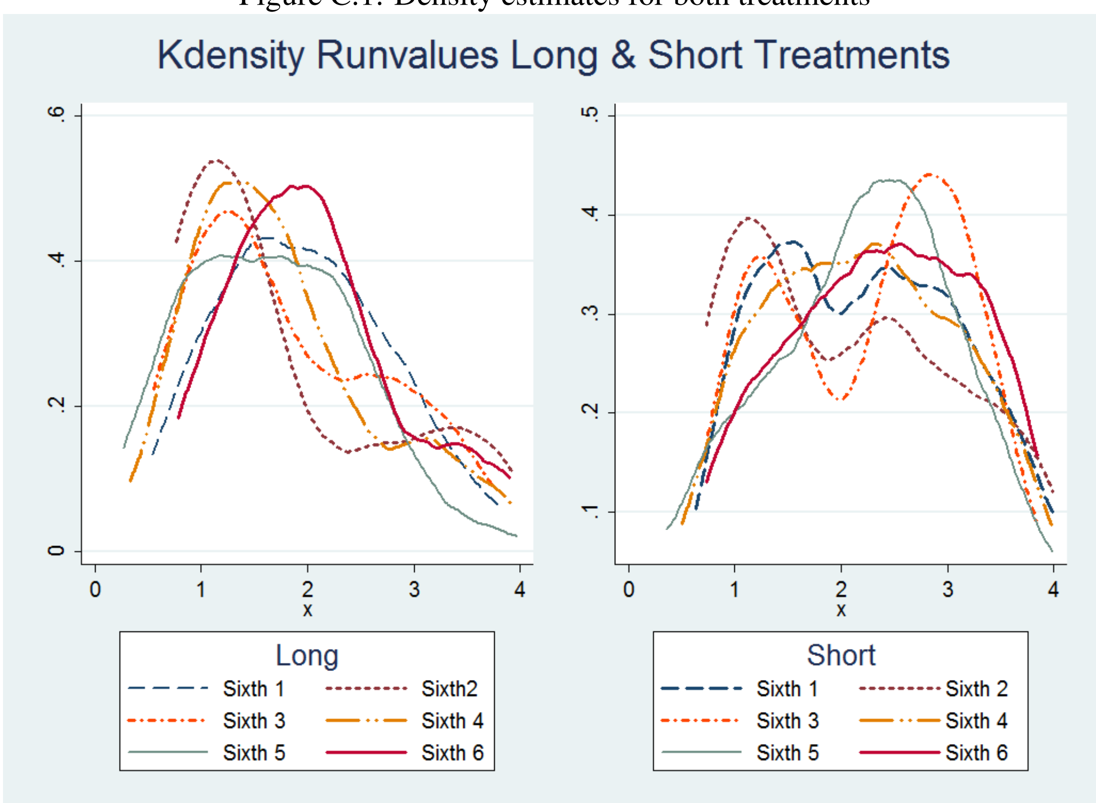

If we break up the sessions into sixth's (1/6) of a session, we have a more cluttered graph, but a more precise description of the dynamics of the experiment. As we can see, the underlying dynamics are the same, with the "boomerang effect" taking place. The difference now is that we can better see the timing of events. While the first sixth is again of a plateau-like shape, we see that risk taking actually starts by the second sixth (earlier than we expected) taking place by the second sixth. On the other hand, our first guess of a panic taking place by the end of the second quarter, and across the next periods was correct, as we can see from the fourth and fifth sixth of the data.

A very interesting result of this more detailed breakdown of estimated pdf's is that we can better appreciate how polarized are the stopping thresholds in the Short treatment. This is clear when observing that most sixths have a bimodal shape, especially in the "panic" (third) one, but also in the final ending distribution. While in the Long treatment this distribution is a sharp hill, in the Short treatment we see a more spread-out result.

C.2 Data Flipping

Flipping the data is a simple procedure where we just need to find a constant, S, large enough such that $S > maxy_{ij}$ for all subject i and round j. Since all data are in the interval $[0,4)$, $S = 4$ can be used to flip the data. Therefore, for every subject$i$ and round j, we can define z_{ij} such that $z_{ij} = 4 - y_{ij}$. A graphical explanation of the process is found in the Appendix B figure.

Figure C.2: Left to right censoring switch

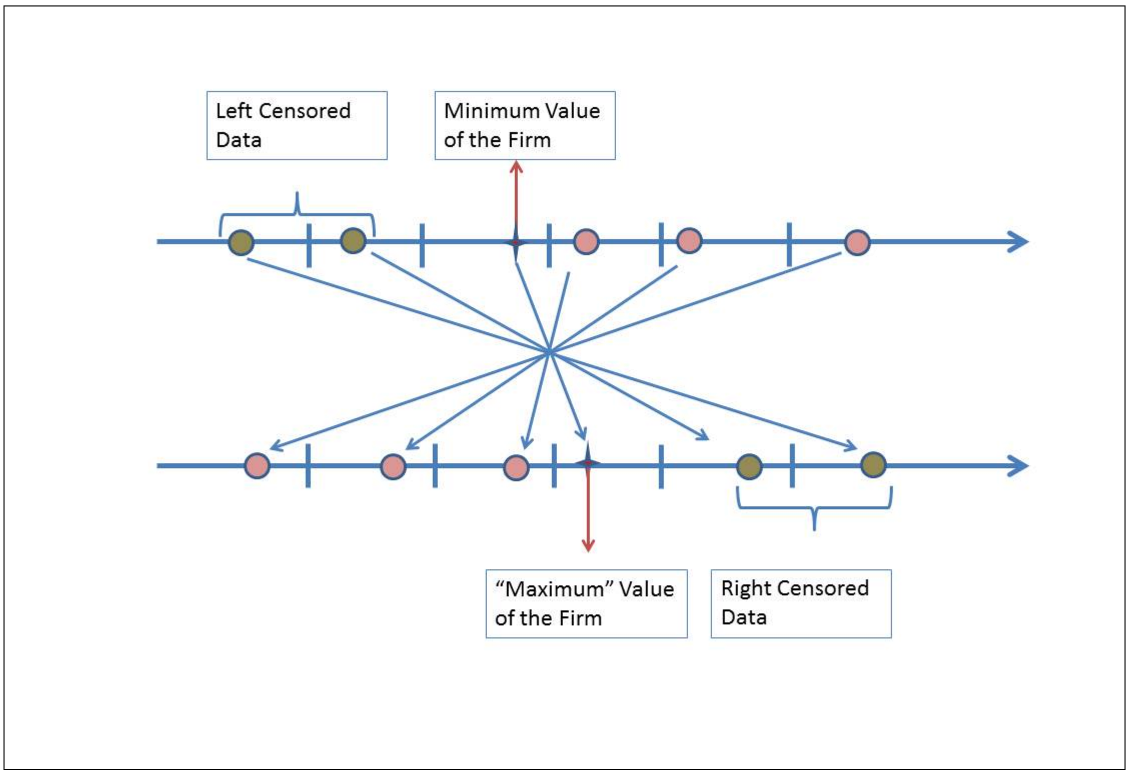

C.3 Full data set and Subset comparison

As the PL estimates graphed below show, our dropping of some data does not affect the underlying relationship of the hazard functions but, as predicted by the theory, it shifts the estimates in the opposite direction of the censoring

Figure C.3: KM Survival estimates for all data and subset

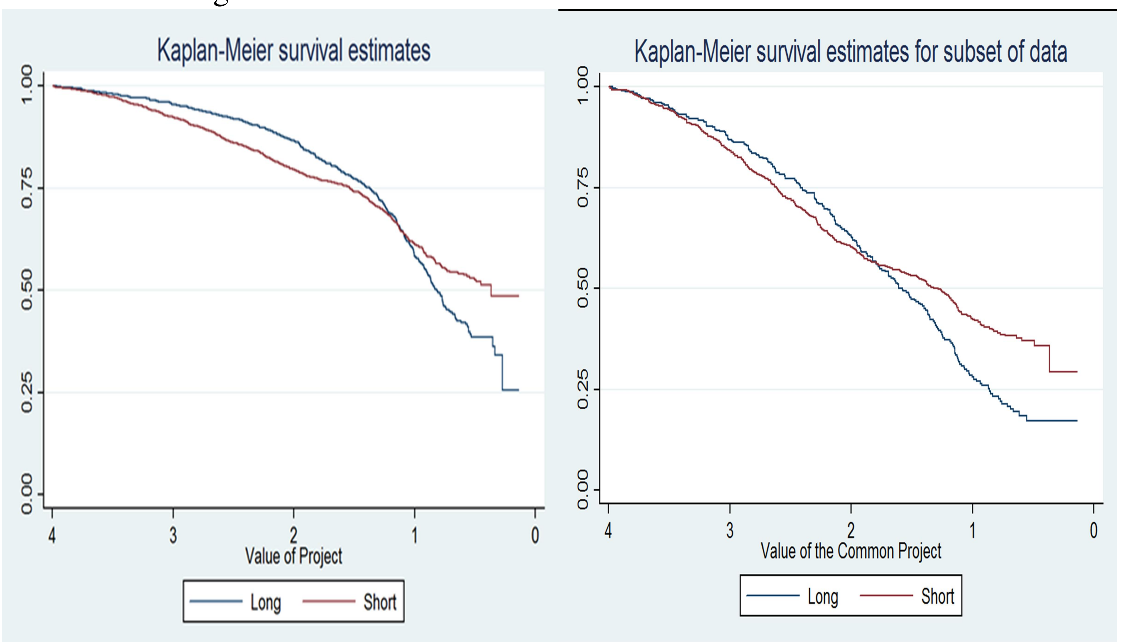

C.4 Instructions

Timing of the Experiment:

The session we will be running today has 60 rounds. At the beginning of the session you will be grouped with 3 other subjects with whom you will play all 60 rounds of the session.

The time units of the round are "ticks" (1/5 of a second). Each round has a probability of 1/150 per tick of maturing; this means that on average each round will last 30 seconds.

The Common Project:

In each round, everyone in your group will start by investing 1 florin (lab currency) into a common project. Every tick the value of the common project will change. To be precise:

- The value of the common project will go up with probability: 0.5001, and down with probability: 0.4999.
- The change in value (whether up or down) will always be 7% of the current value of the investment.

You will be able to track the value of the firm on your screen:

[Image on projector]

Your Decision:

In each round you will make only **ONE** decision:

- To stay in the common project
- To exit the common project

How to exit a project: To exit the common project you will need to slide (not click) your mouse over the counter at the bottom of your screen and connect the numbers 3, 2, and 1

[Image on projector]

Once you have done so a green line will appear on your screen. This green line marks your "**exit request**" and you will exit at the next "**exit gate**" after your exit request.

Exit gates are individual (so no two players share the same exit gate), and happen every 8 seconds. To be more precise:

- In each round, every member of a group is assigned a first "exit gate" within the first 8 seconds
- After that, his next exit opportunities will happen every 8 seconds.
- Example: imagine your first exit opportunity is in second 2 of the round, then your next exit opportunity will be in second 10, then 18, then 26 etc.

[Image on projector]

To stay in the project you do not have to do anything.

Overview:

1. In this experiment you are grouped with three other subjects across 60 rounds.
2. In each round you all start with an investment of 1 florin in the common project
3. Each round you are asked to make one decision: whether or not to stay invested in the common project

4. To exit you need to swipe your mouse over the 3,2,1 countdown area.

5. This swiping will record an exit request and you will exit at your next exit gate

6. To stay you do not need to do anything

Payoffs:

Your payoff in each round will come from two different sources:

- Constant Return

- Original investment return

How much you make from each income source will depend on your decision to stay or to exit, and on the staying or leaving decisions of the other investors in your group.

Constant Return: For every "tick" that you keep your investment in the project, you will get a constant return. This constant return is of 0.004 florins per tick. This means that if you keep your investment for 30 seconds you will get 0.6 florins from the constant return (so a 60% return for every 30 seconds).

Original investment of 1 florin: Of the original investment of 1 florin that you made at the beginning of the round you can get back either the original florin you invested, or a part of the florin you invested, but never more.

This payoff will depend on:

1. Your decision to stay or to exit

2. The decisions of the other investors in your group

3. When and how the round ends.

The round can end in three different ways:

1. You Exit the project: if, at some point, you decide to exit the project, and are able to do so, you will get your 1 florin back independently of the value of the common project. On the other hand, you will stop getting paid the constant return per tick for the rest of the round.

2. Premature end of the project: if 2 investors in your group exit the project, then the project will continue running for 2 extra seconds before it "ends early" and pays all of the remaining investors a "staying value". How much the staying value pays back will depend on where the jagged yellow line is at the moment of the premature ending: a) If the jagged yellow line is above 1, then you will be paid 1 florin. b) If the jagged yellow line is below 1, then you will be paid the value of the line at that point.

3. Maturation of the project: as mentioned, the common project has a probability of 1/150 per tick of maturing. If the common project matures before an early stop happens, then all investors will be paid depending on the value of the common project (green jagged line): a) If the jagged green line is 1 or greater than 1, then all players that are still invested get their 1 florin back b) If the value of the common project is below 1, then all players that are still invested will get back the value of the common project at that point.

You can track both the value of the project and the premature ending value of the project on your screen.

[Image on projector]

<u>**Overview of the payoffs**</u>:

1. Your payoffs come from two different sources: a. Constant payoff b. Individual end of the round

2. The constant payoff gives you 0.004 florins per tick as long as you are invested and the round has not finished (there has not been a premature ending or a maturation of the project)

3. Individual end of projects has 3 different ways of taking place: a.

- You withdraw your investment and get back your entire 1 florin independent of the value of the common project

- The project has a premature ending, in which case those investors that are still in the project get back 1 florin if the yellow jagged line is above 1, or the value of the jagged line if it is below the value of 1

- The project matures, at which point all those still invested get back 1 florin if the green line was above 1, the value of the green line if it was below 1

Important things to notice:

All rounds will continue ticking until the project's maturation, so even if there are premature ending, you will not be given this information until the end of the round. You will also not be told when other investors are leaving the common project nor will you be told where your exit gates are. The information that you will see while the round is ticking will be:

- Value of the project

- Staying value

- Past exit requests by all investors in your group (upper right corner of screen)

[Image on projector]

Once the project has matured, then a screen will appear showing the whole unraveling of the round which includes:

- The exit requests made by all players (green lines)
- The actual exits at each individual exit gate (yellow lines)
- You will also be informed about your exit request and your allowed exit tick.
- Finally, if there was an premature ending it will be shown as a red line.

[Image on projector]

In summary:

Your goal each round is to decide whether you leave or not the project balancing the advantages and disadvantages of staying invested, the probability of a natural end and the behavior of other investors.

But not all rounds are paid. Not all rounds will count for your final payoffs. Although you will see how much you made at the end of each round, 10 of the 60 rounds will count towards your final payoffs. These 10 rounds are randomly chosen by the computer.

Practice: Before the session properly begins, we will have 6 practice rounds so that you get used to the mechanics of the session, so you should practice exiting. These rounds will be shorter than the rounds during the experiment.

While the instructions are somewhat long and complex, it is very important that you understand how the game works. You don't need to really understand all of the probabilities and numbers that we give you, as you can learn from experience, but you should make sure that you understand the mechanics of the game.

FAQ:

1) Is there a pattern in the change of value of the common project? No, we really tried to make it random. No matter what is the history of values that the common project took the probabilities of going up or down on value are always the same.

2) If values over the threshold of 1 always pay me back 1 Florin, why do you show them to me? We show you these values because we think you might be interested in knowing how far away you are from the 1 florin threshold.

Please feel free to ask as many questions as necessary to make sure that you have a full understanding of the instructions. To ask a question, just raise your hand to call my attention.

Bibliography

London Times readership drops 90% and the New York Times could be next. *Investor-Place*, July 22, 2010.

Johannes Abeler, Armin Falk, Lorenz Goette, and David Huffman. Reference points and effort provision. *The American Economic Review*, 101(2):470–492, 2011.

Chris Anderson. *Free: The Future of a Radical Price*. Hyperion, 2009.

Steven Anderson, Daniel Friedman, and Ryan Oprea. Preemption games: Theory and experiment. *Available at SSRN 1282245*, 2008.

Christina Aperjis and Bernardo A. Huberman. Adaptation and the provider's dilemma. *Available at SSRN 1672820*, 2011.

Christina Aperjis and Bernardo A. Huberman. Introducing revenue-generating features. *Patent. US 8271312*, 2012.

Dan Ariely. Combining experiences over time: The effects of duration, intensity changes, and on-line measurements on retrospective pain evaluations. *Journal of Behavioral Decision Making*, 11:19–45, 1998.

Dan Ariely. *Predictably Irrational: The Hidden Forces That Shape Our Decisions*. HarperCollins Publishers, 2008.

Mark Bagnoli and Ted Bergstrom. Log-concave probability and its applications. *Economic Theory*, 26(2):445–469, 2005.

Donna L. Bahry and Rick K. Wilson. Confusion or fairness in the field? rejections in the ultimatum game under the strategy method. *Journal of Economic Behavior & Organization*, 60(1):37–54, May 2006. URL http://ideas.repec.org/a/eee/jeborg/v60y2006i1p37-54.html.

Manel Baucells, Martin Weber, and Frank Welfens. Reference-point formation and updating. *Management Science*, 57(3):506–519, 2011.

M. R. Baye and J. Morgan. A simple model of advertising and subscription fees. *Economics Letters*, 69:345–351, 2000.

Ben Bernanke. Financial regulation and supervision after the crisis: the role of the federal reserve. In *speech at the Federal Reserve Bank of Boston 54th Economic Conference, October*, volume 23, 2009a.

Ben S Bernanke. Financial regulation and financial stability. In *Speech given at the Federal Deposit Insurance Corporations Forum on Mortgage Lending for Low and Moderate Income Households, Arlington, Virginia, July*, volume 8, 2008.

Ben S Bernanke. Reflections on a year of crisis. In *speech at the Federal Reserve Bank of Kansas CityâĂŹs Annual Economic Symposium, Jackson Hole, Wyoming*, 2009b.

Antoni Bosch-Domènech and Joaquim Silvestre. Measuring risk aversion with lists: A new bias. *Theory and Decision*, pages 1–32, 2012.

Samuel Bowles and Herbert Gintis. Behavioural science: homo reciprocans. *Nature*, 415(6868):125–128, 2002.

Stewart Brand. *The Media Lab: Inventing the Future at M.I.T.* Penguin, 1988.

Jordi Brandts and Gary Charness. Hot vs. cold: Sequential responses and preference stability in experimental games. *Experimental Economics*, 2(3):227–238, 2000.

Jordi Brandts and Gary Charness. The strategy versus the direct-response method: a first survey of experimental comparisons. *Experimental Economics*, 14(3):375–398, September 2011a. URL http://ideas.repec.org/a/kap/expeco/v14y2011i3p375-398.html.

Jordi Brandts and Gary Charness. The strategy versus the direct-response method: a first survey of experimental comparisons. *Experimental Economics*, 14(3):375–398, 2011b.

Markus K Brunnermeier. Deciphering the liquidity and credit crunch 2007-08. Technical report, National Bureau of Economic Research, 2008.

Markus K Brunnermeier and John Morgan. Clock games: Theory and experiments. *Games and Economic Behavior*, 68(2):532–550, 2010.

Markus Konrad Brunnermeier, Charles Albert Eric Goodhart, Avinash Persaud, Andrew Crockett, and Hyun Shin. The fundamental principles of financial regulation. 2009.

Antonio Cabrales, Rosemarie Nagel, and Roc Armenter. Equilibrium selection through incomplete information in coordination games: an experimental study. *Experimental Economics*, 10(3):221–234, 2007.

Guillermo A Calvo. Staggered prices in a utility-maximizing framework. *Journal of monetary Economics*, 12(3):383–398, 1983.

A. Colin Cameron and Pravin K. Trivedi. *Microeconometrics: Methods and Applications*. Cambridge University Press, 2005.

First Round Capital. The penny gap. March 10, 2007.

Hans Carlsson and Eric Van Damme. Global games and equilibrium selection. *Econometrica: Journal of the Econometric Society*, pages 989–1018, 1993.

Gary Charness and Matthew Rabin. Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, 117(3):817–869, 2002.

H.A. Chen, D. Levy, S. Ray, and M. Bergen. Asymmetric price adjustment in the small. *Journal of Monetary Economics*, 55(4):728–737, 2008.

Haipeng (Allan) Chen and Akshay R. Rao. Close encounters of two kinds: False alarms and dashed hopes. *Marketing Science*, 21(2):178–196, 2002.

Yin-Wong Cheung and Daniel Friedman. Speculative attacks: A laboratory study in continuous time. *Journal of International Money and Finance*, 28(6):1064–1082, 2009.

M. Cleves, W. Gould, R. Gutierrez, and Y. Marchenko. *An introduction to survival analysis using Stata*. Stata Press, 2008.

David Cooper and John H Kagel. Other regarding preferences: a selective survey of experimental results. *Handbook of Experimental Economics*, 2, 2009.

Christopher T Dawes, James H Fowler, Tim Johnson, Richard McElreath, and Oleg Smirnov. Egalitarian motives in humans. *Nature*, 446(7137):794–796, 2007.

R. Dewan, M. Freimer, and J. Zhang. Managing web sites for profitability: Balancing content and advertising. *Hawaii International Conference on System Sciences*, 7, 2002.

Douglas W Diamond and Philip H Dybvig. Bank runs, deposit insurance, and liquidity. *The journal of political economy*, pages 401–419, 1983.

Dirk Engelmann and Martin Strobel. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *The American Economic Review*, 94(4): 857–869, 2004.

Dirk Engelmann and Martin Strobel. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Reply. *The American economic review*, 96(5):1918–1923, 2006.

Armin Falk and Urs Fischbacher. A theory of reciprocity. *Games and Economic Behavior*, 54(2):293–315, 2006.

Armin Falk, Ernst Fehr, and Urs Fischbacher. Driving forces behind informal sanctions. *Econometrica*, 73(6):2017–2030, 2005.

Armin Falk, Ernst Fehr, and Urs Fischbacher. Testing theories of fairnessâĂŤintentions matter. *Games and Economic Behavior*, 62(1):287–303, 2008.

Josef Falkinger. Attention economies. *Journal of Economic Theory*, 133:266–294, 2007.

Ernst Fehr and Urs Fischbacher. The nature of human altruism. *Nature*, 425(6960): 785–791, 2003.

Ernst Fehr and Simon Gächter. Altruistic punishment in humans. *Nature*, 415(6868): 137–140, 2002.

Ernst Fehr and Bettina Rockenbach. Detrimental effects of sanctions on human altruism. *Nature*, 422(6928):137–140, 2003.

Ernst Fehr and Klaus M Schmidt. A theory of fairness, competition, and cooperation. *The quarterly journal of economics*, 114(3):817–868, 1999.

Gadi Fibich, Arieh Gavious, and Oded Lowengart. Explicit solutions of optimization models and differential games with nonsmooth (asymmetric) reference-price effects. *Operations Research*, 51(5):721–734, 2003.

Gadi Fibich, Arieh Gavious, and Oded Lowengart. The dynamics of price elasticity of demand in the presence of reference price effects. *Journal of the Academy of Marketing Science*, 33(1):66–78, 2005.

David M Frankel, Stephen Morris, and Ady Pauzner. Equilibrium selection in global games with strategic complementarities. *Journal of Economic Theory*, 108(1):1–44, 2003.

Shane Frederick and George Loewenstein. Hedonic adaptation. In Daniel Kahneman and Ed Diener, editors, *Well-being: The foundations of hedonic psychology*, pages 302–329. 1999.

Bruno S. Frey and Alois Stutzer. What can economists learn from happiness research? *Journal of Economic Literature*, 40(2):402–435, 2002.

Daniel Friedman and Ryan Oprea. A continuous dilemma. *The American Economic Review*, 102(1):337–363, 2012.

Rod Garratt and Todd Keister. Bank runs as coordination failures: An experimental study. *Journal of Economic Behavior & Organization*, 71(2):300–317, 2009.

John Geanakoplos. The leverage cycle. *NBER Macroeconomics Annual*, 24(1):1–66, 2009.

John GeanakoplosW. Promises promises. 1996.

Uri Gneezy. Deception: The role of consequences. *The American Economic Review*, 95(1):384–394, 2005.

David Godes, Elie Ofek, and Miklos Sarvary. Content vs. advertising: The impact of competition on media firm strategy. *Marketing Science*, 28(1):20–35, 2009.

Avi Goldfarb and Catherine E. Tucker. Privacy regulation and online advertising. 57 (1):57–ï£¡71, 2011.

Itay Goldstein and Ady Pauzner. Demand–deposit contracts and the probability of bank runs. *the Journal of Finance*, 60(3):1293–1327, 2005.

W. Gould, J. Pitblado, and W. Sribney. *Maximum likelihood estimation with Stata*. Stata Press, 2005.

Werner Güth, Carsten Schmidt, and Matthias Sutter. Bargaining outside the lab–a newspaper experiment of a three-person ultimatum game*. *The Economic Journal*, 117 (518):449–469, 2007.

Zhiguo He and Wei Xiong. Dynamic debt runs. *Review of Financial Studies*, 25(6): 1799–1843, 2012.

Ori Heffetz and John A. List. Is the endowment effect a reference effect? *NBER Working Paper 16715*, 2011.

Frank Heinemann, Rosemarie Nagel, and Peter Ockenfels. The theory of global games on test: experimental analysis of coordination games with public and private information. *Econometrica*, 72(5):1583–1599, 2004.

Charles A Holt and Susan K Laury. Risk aversion and incentive effects. *The American Economic Review*, 92(5):1644–1655, 2002.

Bernardo A. Huberman and Fang Wu. The economics of attention: maximizing user value in information-rich environments. In *ADKDD '07: Proceedings of the 1st international workshop on Data mining and audience intelligence for advertising*, pages 16–20, New York, NY, USA, 2007a. ISBN 978-1-59593-833-6. doi: http://doi.acm.org/10.1145/1348599.1348602.

Bernardo A. Huberman and Fang Wu. Comparative advantage and efficient advertising in the attention economy. *Proceedings of the 22nd European Conference on Operational Research (Euro XXII)*, 2007b.

Dominic Hyde. Sorites paradox. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Winter 2011 edition, 2011.

John H Kagel and Katherine Willey Wolfe. Tests of fairness models based on equity considerations in a three-person ultimatum game. *Experimental Economics*, 4(3): 203–219, 2001.

D. Kahneman, B.L. Fredrickson, C.A. Schreiber, and D.A. Redelmeier. When more pain is preferred to less: Adding a better end. *Psychological Science*, 4(6):401–405, 1993.

Daniel Kahneman and Richard Thaler. Economic analysis and the psychology of utility:

Applications to compensation policy. *The American Economic Review*, 81(2):341–346, May 1991.

Daniel Kahneman and Richard H. Thaler. Anomalies: Utility maximization and experienced utility. *Journal of Economic Perspectives*, 20(1):221–234, 2006.

Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.

Gurumurthy Kalyanaram and Russell S. Winer. Empirical generalizations from reference price research. *Marketing Science*, 14(3):161–169, 1995.

Botond Kőszegi and Matthew Rabin. A model of reference-dependent preferences. 121 (4):1133–1165, 2006.

John P Klein and Melvin L Moeschberger. *Survival analysis: techniques for censored and truncated data*. Springer, 2003.

Alexander Klos, Norbert Sträter, and Finance Center Münster. How strongly do players react to changes in payo parameters in an experimental bank run game? 2012.

Praveen K. Kopalle, Ambar G. Rao, and Joao L. Assuncao. Asymmetric reference price effects and dynamic pricing policies. *Marketing Science*, 15(1):60–85, 1996.

Lakshman Krishnamurthi, Tridib Mazumdar, and S. P. Raj. Asymmetric response to price in consumer brand choice and purchase quantity decisions. *Journal of Consumer Research*, 19:387–400, 1992.

Arvind Krishnamurthy. How debt markets have malfunctioned in the crisis. *The Journal of Economic Perspectives*, 24(1):3–28, 2010.

Subodha Kumar and Suresh P. Sethi. Dynamic pricing and advertising for web content providers. *European Journal of Operational Research*, 197:924–944, 2009.

Jaron Lanier. Pay me for my content. *The New York Times*, November 20, 2007.

Andreas Leibbrandt and Raúl López-Pérez. The envious punisher. Technical report, University of Zurich working paper, 2008.

Philippe Madies. An experimental exploration of self-fulfilling banking panics: Their occurrence, persistence, and prevention*. *The Journal of Business*, 79(4):1831–1866, 2006.

Tridib Mazumdar, S.P. Raj, and Indrajit Sinha. Reference price research: Review and propositions. *Journal of Marketing*, 69:84–102, 2005.

Rupert G Miller Jr. What price kaplan-meier? *Biometrics*, pages 1077–1081, 1983.

ML Moeschberger and John P Klein. A comparison of several methods of estimating the survival function when there is extreme right censoring. *Biometrics*, pages 253–259, 1985.

Stephen Morris and Hyun Shin. Global games: theory and applications. 2001.

Javad Nasiry and Ioana Popescu. Dynamic pricing with loss averse consumers and peak-end anchoring. *INSEAD Working Paper No. 2010/102/DS*, 2010.

Maurice Obstfeld. Models of currency crises with self-fulfilling features. *European economic review*, 40(3):1037–1047, 1996.

Axel Ockenfels and Gary E. Bolton. Erc: A theory of equity, reciprocity, and competition. *American Economic Review*, 90(1):166–193, March 2000. URL `http://ideas.repec.org/a/aea/aecrev/v90y2000i1p166-193.html`.

Ryan Oprea, Daniel Friedman, and Steven T Anderson. Learning to wait: A laboratory investigation. *The Review of Economic Studies*, 76(3):1103–1124, 2009.

Ryan Oprea, Keith Henwood, and Daniel Friedman. Separating the hawks from the doves: Evidence from continuous time laboratory games. *Journal of Economic Theory*, 146(6):2206–2225, 2011.

Ioana Popescu and Yaozhong Wu. Dynamic pricing strategies with reference effects. *Operations Research*, 55(3):413–429, 2007.

Ashutosh Prasad, Vijay Mahajan, and Bart Bronnenberg. Advertising versus pay-per-view in electronic media. *Intern. J. of Research in Marketing*, 20:13–30, 2003.

W.S. Quinn. The puzzle of the self-torturer. *Philosophical studies*, 59(1):79–90, 1990.

Luis Rayo and Gary S. Becker. Evolutionary efficiency and happiness. *Journal of Political Economy*, 115(2), 2007.

Arthur J. Robson. The biological basis of economic behavior. *Journal of Economic Literature*, 39, 2001.

Bettina Rockenbach and Manfred Milinski. The efficient interaction of indirect reciprocity and costly punishment. *Nature*, 444(7120):718–723, 2006.

Kristina Shampanier, Nina Mazar, and Dan Ariely. Zero as a special price: The true value of free products. *Marketing Science*, 26(6):742–757, 2007a.

Kristina Shampanier, Nina Mazar, and Dan Ariely. Zero as a special price: The true value of free products. *Marketing Science*, 26(6):742–757, 2007b.

Hyun Song Shin. Reflections on northern rock: the bank run that heralded the global financial crisis. *The Journal of Economic Perspectives*, 23(1):101–120, 2009.

Herbert A. Simon. *Designing Organizations for an Information-Rich World*, pages 37–72. 1971.

Amos Tversky. Intransitivity of preferences. *Psychological Review*, 76(1):31–48, 1969.

Erte Xiao and Daniel Houser. Emotion expression in human punishment behavior. *Proceedings of the National Academy of Sciences of the United States of America*, 102(20):7398–7401, 2005.