UNIVERSITY OF CALIFORNIA,
IRVINE

Essays on Auctions, Conflict, and Social Networks

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Economics

by

Michael Caldara

Dissertation Committee:
Professor Michael McBride, Chair
Professor Igor Kopylov
Professor David Porter
Professor Stergios Skaperdas

2014

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I thank Jean-Paul Carvalho, Igor Kopylov, Michael McBride, David Porter and Stergios Skaperdas for guidance, helpful comments and support. I thank Deans Barbara Dosher and Bill Maurer for their continued support of the Experimental Social Science Laboratory, my home while at UCI. I have benefited greatly from the support of several administrators, particularly John Sommerhauser, Pat Frazier, and Nancy Ford. I acknowledge and thank the Center for Economics and Public Policy, Department of Economics, Economic Science Institute, Experimental Social Science Laboratory, Institute for Humane Studies and Institute for Mathematical Behavioral Sciences for grants and fellowships. I also acknowledge and thank the Air Force Office of Scientific Research (Award No. FA9550-10-1-0569) and the Army Research Office (Award No. W911NF-11-1-0332) for financial support. Lastly, I thank my wife, Grace Caldara, for her unwavering patience and support throughout the Ph.D. program.

CURRICULUM VITAE

Michael Caldara

EDUCATION

Doctor of Philosophy in Economics **2014**
University of California, Irvine *Irvine, CA*

Master of Arts in Economics **2012**
University of California, Irvine *Irvine, CA*

Bachelor of Science in Management **2009**
Boston College *Chestnut Hill, MA*

RESEARCH EXPERIENCE

Research Associate **2013–present**
Economic Science Institute *Orange, CA*

Graduate Student Researcher **2011–2013**
University of California, Irvine *Irvine, CA*

TEACHING EXPERIENCE

Instructor **2014**
Chapman University *Orange, CA*

Instructor **2011**
University of California, Irvine *Irvine, CA*

Teaching Assistant **2010–2011**
University of California, Irvine *Irvine, CA*

REFEREED JOURNAL PUBLICATIONS

"The Efficacy of Tables versus Graphs in Disrupting Dark Networks: An Experimental Study," *Social Networks*, 35(3), July 2013, with Michael McBride.

ABSTRACT OF THE DISSERTATION

Essays on Auctions, Conflict, and Social Networks

By

Michael Caldara

Doctor of Philosophy in Economics

University of California, Irvine, 2014

Professor Michael McBride, Chair

The broad objective of this dissertation is to understand human behavior in large group interactions. More narrowly, I apply game theory and experimental methods to three distinct settings: pay-to-bid auctions, strategic network formation, and stateless societies. In each case, I study the relationship between individual behavior and group outcomes. These seemingly unrelated settings do share several elements in common: (1) groups are relatively large, (2) groups benefit from cooperation, and (3) individual incentives are misaligned, potentially leading to adversarial interactions (i.e., conflict). Additionally, each of the three studies, detailed in Chapters 2, 3, and 4 respectively, makes a stand-alone contribution to understanding its topic with implications for future research. Thus, the contribution of this dissertation is twofold: the contribution of each individual Chapter to its subject and the collective contribution of all Chapters to understanding large group interactions.

Chapter 1

Introduction

Chapter 1 provides an overview of the dissertation. In Chapter 2 of this dissertation, I experimentally study bidding behavior in pay-to-bid auctions, also known as "penny auctions." I find that auction revenues above the value of the prize can largely be explained by inexperienced participants and that the earnings of individual participants are highly correlated with strategic sophistication. In Chapter 3 of this dissertation (joint with Michael McBride), we experimentally study the dynamics of strategic network formation with limited observation. We find that the networks formed under limited observation are highly efficient and less prone to positional jockeying among actors when compared to networks formed under full observation. In Chapter 4 of this dissertation, I develop a model of the origin of the state. Starting from an anarchic environment, where property is insecure and must be defended, I show how evolutionary forces would favor the selection of a violent elite. When the distribution of violence capacities is relatively asymmetric or when the gains from cooperation are low, this elite must be bought off to ensure cooperation. Lastly, Chapter 5 concludes.

1.1 Bidding Behavior in Pay-to-Bid Auctions

In Chapter 2, I experimentally study bidding behavior in pay-to-bid auctions, also known as "penny auctions." The auction mechanism has become increasingly popular among online retailers in recent years and typically generates revenues well in excess of the value of the prize, contrary to theoretical predictions. Several plausible explanations have been proposed, but there is a dearth of empirical evidence to substantiate these explanations. To help distinguish between these explanations, I conduct a series of controlled laboratory experiments. I find that auction revenues above the value of the prize can largely be explained by inexperienced participants and that the earnings of individual participants are highly correlated with strategic sophistication. Persistent money losers learn to quit, suggesting that pay-to-bid auctions can only generate revenues above the value of the prize so long as they can find new, inexperienced bidders.

1.2 Network Formation with Limited Observation

In Chapter 3 (joint with Michael McBride), we experimentally study the dynamics of strategic network formation with limited observation. Limited observation is a real feature of many network formation settings, and yet much of the academic work on network formation assumes that the individual actors have full information about the network structure. As a result, the impact of limited observation on connection decisions and on the network structure is not well understood. We conduct the first ever experimental study of network formation with limited observation. We find that the networks formed under limited observation are highly efficient and less prone to positional jockeying among actors when compared to networks formed under full observation. The networks do exhibit some inefficient overconnection, but redundant connections are generally outside of the observational range.

1.3 The Origin of the State as a Stationary Bandit

In Chapter 4, I develop a model of the origin of the state inspired by the seminal theory of Olson (1993, 2000). Olson suggests that the state originated from the incentives of a *stationary bandit*, who resorts to moderate tax-theft in the short-run to preserve long-run production incentives. Olson's framework is particularly well-suited to study predatory states, a commonality throughout recorded history and still prevalent in some parts of the world today. Starting from an anarchic environment, where property is insecure and must be defended, I show how evolutionary forces would favor the selection of a violent elite. The impact of this dynamic depends greatly on the initial conditions. When the distribution of violence capacities is relatively asymmetric or when the gains from cooperation are low, the elite must be bought off to ensure cooperation and the cycle of predation persists. Alternatively, when the distribution of violence capacities is relatively symmetric or when the gains from cooperation are high, predation may cease to be profitable, allowing for the natural evolution of property rights and peaceful economic exchange. In all cases, the natural incentives of the elite will help to fortify the society against *external* threats.

Chapter 2

Bidding Behavior in Pay-to-Bid Auctions

This chapter provides experimental evidence of significant over-bidding in pay-to-bid auctions relative to the risk neutral Nash equilibrium prediction. Several theories that explain the excess revenues are tested and a behavioral bias is identified as a key cause of the over-bidding. Over-bidding decreases with experience, and strategic sophistication is highly correlated with the earnings of individual participants. Many of the least successful subjects (i.e., persistent money losers) cease auction participation altogether. This, along with the observed learning effect, suggests that the pay-to-bid auction mechanism can only sustain revenues above the value of the prize as long as new participants can be attracted.

2.1 Introduction

In recent years, the "pay-to-bid" auction (or "penny" auction), has become an increasingly popular mechanism among online retailers. While this still fledgling industry has experienced some turnover, most notably the 2011 shutdown of Swoopo.com, the industry as a whole is growing quickly.[1] The auction format boasts average revenues well in excess of the value of the prize (Thaler *New York Times* 2009,[2] Augenblick 2012, Platt, Price & Tappen 2013) and yet tens if not hundreds of individuals willingly participate in each new auction. This is in sharp contrast to the standard game theoretic prediction (Hinnosaar 2010, Augenblick 2012, Platt et al. 2013) that average revenues will equal at most the value of the prize. An apparent key to the auction format's success is its ability to effectively exploit behavioral biases while still enticing new participants to join. This enticement stems from the generally great deal offered to the high bidder (e.g., a High Definition Television for $24.36, a $200 gift card for $17.45, a Blu-ray player for $1.71). It is the participants as a collective whole that lose, as the auctioneer collects a small non-refundable bid fee each time a bid is placed. The three auctions in the above example, taken from a popular pay-to-bid auction website, raised $2654.72 and the retail price of the items was $779.98. With such lopsided outcomes it is no wonder the auction has been called "the evil stepchild of game theory and behavioral economics" (Gimein *Washington Post* 2009).[3]

Several theories have been proposed to explain the discrepancy between empirical revenue estimates and the game theoretic prediction (Byers, Mitzenmacher & Zervas 2010, Augenblick 2012, Platt et al. 2013), but no systematic attempt has been made to empirically assess the plausibility of these theories. I seek to bridge this gap with the experimental method. By examining bidding behavior in pay-to-bid auctions in a controlled laboratory setting, I can observe auction outcomes in an environment where many of the confounding factors present in field data have been removed (e.g., imperfect information, asymmetries, unmeasured risk preferences, censored data). I provide perfect information, common values and a symmetric

[1] The industry experienced a 126 percent year over year increase in unique visitors per month in 2011, and a 645 percent increase in 2012. Currently, the industry's web traffic is equivalent to 1.2 percent of the web traffic generated by Ebay. See Table A.3 for details.

[2] Thaler, Richard H. "Paying a Price for the Thrill of the Hunt." The New York Times [New York] 15 Nov. 2009, New York Edition: BU5.

[3] Gimein, Mark. "The Big Money: The Pennies Add Up at Swoopo.com." *The Washington Post* [Washington, D.C.] 12 Jul. 2009.

playing field to allow for a closer comparison to the theory. Bidding behavior is tested in both the continuous time setting experienced in pay-to-bid auctions on the internet, and in the discrete time, simultaneous decision setting modeled in the theory. In addition, the number of auction participants (3 or 5) is varied across treatments to aid in distinguishing among competing explanations for the over-bidding.

A main contribution of this chapter is that the experimental design provides the ability to distinguish between purely behavioral explanations of overbidding (sunk cost fallacy, escalation of commitment, mistakes) and other environmental factors (information asymmetries, risk preferences, shill bidding). On the one hand, if behavioral biases are the primary cause of the empirically observed over-bidding, the pay-to-bid auction mechanism could be a useful case study for researchers interested in studying behavioral economics in general. For example, if the pay-to-bid mechanism serves to magnify known behavioral biases relative to alternative auction formats, it could be used as a tool to explicitly study these biases. On the other hand, if behavioral biases cannot explain the excess revenues, this could be cause for concern for regulators such as the Federal Trade Commission who have expressed concern that pay-to-bid auctions may be a form of gambling and that some websites may secretly use computerized shill bidders to inflate prices.[4] Either way, distinguishing between these theories is a logical first step to help inform future research.

In all treatments, I find persistent over-bidding relative to the risk neutral Nash equilibrium prediction and this over-bidding tapers off with time. Auction revenues in the treatments with 5 auction participants are significantly greater than auction revenues in the treatments with 3 auction participants, consistent with the hypothesis that participants do not adequately adjust bidding probabilities to account for the number of other participants. I find no difference between the auction revenues under the discrete rounds treatment and the auction revenues under the continuous time treatment, suggesting that the game modeled in the

[4]see e.g., www.consumer.ftc.gov/articles/0037-online-penny-auctions.

baseline theory is capturing many of the key strategic considerations present in a pay-to-bid auction on the internet. Risk aversion (measured using the Eckel & Grossman 2008 method) has the predicted effect on number of bids placed relative to risk neutral participants, but the results for risk seeking participants are inconclusive. I find strategic sophistication to be a strong determinant of individual auction outcomes. I observe the use of signaling strategies, but these strategies were not generally successful. This ineffectiveness may be in part due to the shortened auction lengths and symmetries built in to the experiment design. Another key finding is that many less successful participants cease participation entirely over the course of the session. The findings that over-bidding decreases with experience and that less successful participants learn to leave supports the hypothesis of Wang & Xu (2013) that pay-to-bid auction websites profit from a "revolving door of new bidders."[5] For a fixed pool of participants, it does not appear that the pay-to-bid auction format can generate sustained revenues in excess of the value of the prize.

2.2 Background

The pay-to-bid auction, which has similarities to the dollar auction (Shubik 1971), the bucket auction (Carpenter, Holmes & Matthews 2011, 2012), the war of attrition (Smith 1974, Krishna & Morgan 1997) and the all-pay auction (Baye, Kovenock & de Vries 1996, Krishna & Morgan 1997), is distinguished by the following features: Participants choose between bidding and not bidding. When a participant bids, the auction price increases by a small fixed increment and that participant is charged a non-refundable fee. A participant may bid multiple times during the auction, but must pay the non-refundable fee *each* time she bids. When there is no time left on the countdown clock, the *last* participant to place a bid

[5] Wang & Xu (2013) show that newer bidders tend to lose money and and experienced bidders tend to make money, but it is difficult to assess *why* this is the case using field data due to the presence of confounding factors. The experimental study in this chapter avoids the problem by controlling for these confounding factors.

wins the prize and pays the end auction price. Like an English auction, whenever a bid is placed in the last few seconds, a small amount of time (i.e., 15 or 20 seconds) is added to the countdown clock. Thus, the auction can be in the "last few seconds" indefinitely.

It is clear from the structure of this auction that it is rife with the opportunity to succumb to behavioral biases and to make mistakes. The bid fees are a sunk cost, the monetary commitment of the active participants grows in small increments as the auction progresses, decisions must be made in a short window of time, and each decision to not bid risks ending the auction. Because the fees from past bids are lost forever, participants might fall victim to the sunk cost fallacy. If we try to model these sunk costs with prospect theory and loss aversion (Kahneman & Tversky 1979, 1991, 1992, Thaler 1980), excess revenues can be explained by risk seeking behavior in small losses (bid fees) from the reference point (the participant's initial wealth).[6] The growing monetary commitment of active participants may lead to further escalation of commitment as these participants try to dig themselves out of a hole. For instance, participants may be better able to rationalize unsuccessful bids through further bidding (Staw 1981, Wong, Kwong & Ng 2008). Additionally, the negative emotions that participants may encounter in this adversarial setting have been shown to increase the likelihood of escalation (Wong, Yik & Kwong 2006, Tsai & Young 2010), although if the participants are able to predict future regret (Wong & Kwong 2007, Ku 2008*a*), learn not to escalate from prior experience (Ku 2008*b*) or set a mental budget to limit participation (Heath 1995), this effect may be mitigated. Lastly, the short time between decisions may lead to mistakes due to limited cognition (e.g., Simon 1976).

While it is easy to develop a list of behavioral phenomena such as the sunk cost fallacy, escalation of commitment, bounded rationality, and loss aversion, that *could* explain the excess revenues of pay-to-bid auction websites, it is difficult to determine theoretically whether these biases *are* driving excess revenues, and if so, which of these biases are the most salient.

[6] Alternatively, Augenblick (2012) shows a similar result by explicitly modeling naive sunk cost fallacy in pay-to-bid auctions.

A variety of alternative theories have been proposed that can also explain excess revenues, including: Information asymmetries and imperfect information (Byers et al. 2010), risk loving preferences (Platt et al. 2013), shill bidding (Byers et al. 2010, Platt et al. 2013), and signaling strategies (Byers et al. 2010, Augenblick 2012). Furthermore, evidence from field data highlights the important influence that heterogeneities between participants have on individual outcomes. Both strategic sophistication and experience (Wang & Xu 2013),[7] and reputation and signaling behavior (Goodman 2012) have been shown to play a role in individual auction outcomes. While all of the above findings are plausible explanations for the excess revenues of pay-to-bid auction websites, this analysis is complicated by the complexity of the pay-to-bid auction format. For the sake of simplicity and tractability, the baseline pay-to-bid auction theory of Hinnosaar (2010), Augenblick (2012), and Platt et al. (2013) makes many abstractions. But until the theoretical model is better understood, we have no way of knowing whether the game modeled by the baseline theory bears any resemblance to the game being played by the participants on the pay-to-bid auction websites.

While this is the first manuscript to experimentally study the pay-to-bid auction format and the first to compare it to the discrete time version modeled in the baseline theory, this work fits into a large body of experimental work on related auction formats. Lam (2011) studies the impact of an exit option on revenues in Augenblick's (2012) discrete time version of the pay-to-bid auction format. He finds revenues slightly below the value of the prize in the no exit treatments and well below the value of the prize in the treatments with exit.[8]

Carpenter et al. (2011, 2012) find revenues well above the value of the prize in the procedu-

[7] Wang & Xu use an adaptation of Camerer, Ho & Chong (2004) that is especially suited for measuring sophistication in pay-to-bid auctions.

[8] The discrete time pay-to-bid auction format analyzed in this chapter differs from the format used in Lam (2011) in two key ways. First, each bidder pays the bid fee regardless of whether that bidder is selected to be the next round's high bidder. This allows for a closer comparison to the continuous time pay-to-bid auction, in which all bids are accepted. Second, to discourage the use of history dependent strategies such as signaling, the identity of the high bidder is not revealed. This allows for a closer comparison to the theory, as the equilibrium concepts used in the theory do not allow for history dependent strategies. The absence of these design elements in Lam (2011) could serve to lower average revenues relative to the format analyzed in this chapter.

rally similar bucket auction, but the proceeds from these auctions go to charity and so it is difficult to disentangle the revenue generated by selfish motives from the revenue generated by charitable motives.

Much of the contest literature has found evidence of over-bidding in lottery and all-pay contests. As the pay-to-bid auction is a type of dynamic contest, we might expect to observe some over-bidding as well. Of key interest is the *cause* of any over-bidding, and whether this over-bidding persists as subjects gain experience. For example, Sheremeta (2011) finds evidence that mistakes play a role in over-bidding, Amaldoss & Rapoport (2009) attribute this over-bidding to strategic sophistication, Gneezy & Smorodinsky (2006) find that over-bidding depends on the number of players, and Herrmann & Orzen (2008) link over-bidding with spiteful preferences. It is also possible that the dynamic nature of the pay-to-bid auction may turn this result on its head, as Hrisch & Kirchkamp (2010) do find over-bidding in the static all-pay auction, but find under-bidding in the dynamic war of attrition, despite the theoretical similarities, and Lam (2011) finds under-bidding in the discrete time pay-to-bid auction. In addition, there is evidence that experience diminishes over-bidding in the all-pay auction (Davis & Reilly 1998), and so we might expect any over-bidding to diminish with time. It is unclear whether the sunk cost fallacy will play an important role in this setting. The fee structure of pay-to-bid auctions will clearly generate sunk costs, and so we might expect to observe sunk cost fallacy as has been found in more general settings (e.g., Arkes & Blumer 1985) but other studies find the effect to be small (e.g., Friedman, Pommerenke, Lukose, Milam & Huberman 2007) and the results of Sheremeta (2010) for multi-stage contests are inconsistent with sunk cost fallacy all together.

2.3 Baseline Theory

2.3.1 The Model

I first summarize the theoretical models presented in Hinnosaar (2010), Augenblick (2012), and Platt et al. (2013). In this summary, I devote particular attention to the case where decisions are simultaneous and multiple bids can be accepted (i.e., the bid fee is charged for every bid no matter what the outcome), both key features of pay-to-bid auctions on the internet. The main abstraction of this model is that decisions are made in a series of discrete rounds rather than in continuous time, but Augenblick (2012) shows that this abstraction should not change any of the payoff-relevant outcomes.

The baseline model is as follows. A single item is put up for auction. A non-participating auctioneer conducts the auction and a non-participating seller commits to sell the item at the end auction price. Any revenues generated by the auction will be divided in some manner between the auctioneer and the seller. There are $n \geq 2$ participants in the auction, indexed by $i \in \{1, 2, ..., n\}$, who share a known common value Π for the item. The number of participants n is fixed throughout the auction and n is known to all participants. The auction is conducted over a series of discrete rounds, indexed by $t \in \{0, 1, 2, ...\}$. In all rounds $t > 0$, exactly one of the n participants is designated as the "high bidder," a title that is awarded based on the outcome of the prior round, and the remaining $n - 1$ participants are designated as "non-leaders." The auction starts in round $t = 0$ at initial price $P_0 \geq 0$ with no high bidder.

In a given round t, non-leaders simultaneously choose between bidding and not bidding. The high bidder does not participate in the round. If a participant chooses to bid, the auction price increases by the price increment $\varepsilon > 0$, the participant pays a non-refundable bid fee of C ($\varepsilon < C < \Pi - \varepsilon$), and the participant gains a chance to become next round's high bidder. The high bidder is chosen at random from the set of all bidders and so the chance that this participant will become next round's high bidder is $1/k$ where k is the number of bids placed in the round. If a participant chooses not to bid, this participant does not pay the bid fee, but also does not have a chance of becoming the next round's high bidder. If all non-leaders pass in round $t > 0$, then the auction ends and the high bidder wins the item at price P_t. If all non-leaders pass in round $t = 0$, then the auction ends and the seller keeps the item. Hinnosaar (2010) shows that the risk neutral symmetric Markov perfect equilibrium is recursively characterized by Equations A.1 and A.2 **(see Appendix)**.[9] Using these equations, Hinnosaar (2010) provides bounds on expected revenue in the risk neutral symmetric Markov perfect equilibrium. However, many asymmetric and non-Markovian equilibria also exist. I generalize these propositions to include all equilibria in Theorem 2.1 below.

Theorem 2.1. *Let $E(R)$ denote the expected auction revenue conditional on sale in any risk neutral equilibrium characterized by (A.1) and (A.2). Then $E(R)$ has the following properties:*

i $E(R) \leq \Pi$.

ii *For all values of the auction parameters, there exists an equilibrium such that* $E(R) = \Pi$. *This equilibrium may not be Markov perfect.*

Proof. See Appendix A.3. □

This theorem suggests that expected revenue in the risk neutral equilibrium should equal *at most* the value of the item.[10] Because there always exist equilibria such that $E(R) = \Pi$

[9] Markov perfection (Maskin & Tirole 2001) requires that a participant's probability of bidding depend only on payoff relevant-past events. In this context, all payoff relevant information is captured by the current price p_t. History dependent strategies such as signaling or cooperation are not considered.

[10] This is a weaker statement than can be made in the framework of Augenblick (2012) and Platt et al. (2013), in which expected revenue equals the value of the item, and is due to the possibility of multiple bids being accepted in a round.

and never exist equilibria such that $E(R) > \Pi$ under risk neutrality, I use $E(R) = \Pi$ as a conservative baseline prediction to compare against average revenue in the experiment sessions. If instead the assumed equilibrium is such that $E(R) < \Pi$ then this would only serve to exaggerate any difference between the theoretical prediction and the empirical result in the case of over-bidding.

Another important consideration is the effect of risk preferences on expected revenue. Platt et al. (2013) show in their model of pay-to-bid auctions that when participants are symmetrically risk averse expected revenue decreases, and when the participants are symmetrically risk loving expected revenue increases. The following theorem also holds for this model.

Theorem 2.2. *(Platt et al. 2013) Let $E(R)$ denote the expected auction revenue conditional on sale in the symmetric equilibrium characterized by (A.1) and (A.2). Then $E(R)$ has the following properties:*

i $E(R) < \Pi$ *when participants are symmetrically risk averse.*

ii $E(R) > \Pi$ *when participants are symmetrically risk loving.*

To summarize the predictions from the baseline theory: given any full information pay-to-bid auction where participants make bidding decisions simultaneously in a series of discrete rounds by maximizing expected utility, expected revenue is no greater than the value of the item unless the participants are sufficiently risk loving in aggregate.

2.3.2 Reasons Why The Baseline Model Might Fail

Because the profit margins of pay-to-bid auction websites are empirically estimated to be much greater than zero (e.g., Augenblick 2012, Platt et al. 2013), it is clear that the baseline model fails to capture an important factor that is relevant in a real pay-to-bid auction.

Several additional theories have been proposed in an attempt to explain this discrepancy between the theoretical prediction of zero profits and the empirical estimate of positive profits. While each of these explanations is plausible, we must turn to empirical evidence to determine the true cause of this discrepancy. Field data is of limited use for these purposes as the baseline theory makes many abstractions from the strategic environment of a real pay-to-bid auction on the internet.[11] For example, the number of participants is not generally known, participants may have different valuations for the item and may face different bid costs, the auctions are conducted in continuous time rather than discrete rounds, and because the user ID of the high bidder is made public knowledge, bidder reputation and signaling strategies may be a factor. The presence of these confounding factors makes direct comparison to the theory difficult. On the other hand, in a controlled laboratory experiment I can exactly replicate the strategic environment developed in the baseline theory to better differentiate between these explanations. A key contribution of this chapter is its ability to isolate the proximate cause of seller profits in pay-to-bid auctions. I summarize the theories below.

1. Information asymmetries and imperfect information

The baseline model assumes perfect information about the auction parameters (n, Π, C, ε), and that all participants face the same auction parameters. However, only the bid increment ε is known with certainty. In a pay-to-bid auction on the internet, only "bid" actions are observed. The number of participants actively choosing "no bid" is unknown, and thus the number of participants n is unknown. The value Π is also unknown because individual values are private and need not be common to all participants. Lastly, because many pay-to-bid auction websites regularly have auctions for "bid packs" (packs containing a quantity of pre-paid bids), it is possible for some bidders to face a lower cost of bidding than others.

[11]This is not to say that field data is not useful for other purposes. For instance, the analyses of Wang & Xu (2013) and Goodman (2012) shed light on important factors that can affect participant outcomes such as experience, sophistication, reputation, and signaling behavior.

Suppose that $n-1$ participants believe that all n participants face the auction parameters (n, Π, C, ε), but that 1 participant knows she faces a higher value Π' or a lower cost C'. Byers et al. (2010) show that this will increase expected revenue $E(R)$ so long as the $n-1$ other participants do not realize the asymmetry. In addition, Byers et al. (2010) show that when participants systematically underestimate the number of other participants n this will increase expected revenue $E(R)$. To control for these possibilities, I provide full information about the auction parameters in the experiment.

2. Risk loving preferences

The baseline model also assumes that participants are risk neutral. If instead the participants are risk loving, this would increase expected revenue. As pay-to-bid auction participants are faced with relatively low winning probabilities and seemingly random auction outcomes with high rewards, these auctions do have the feel of a lottery. Thus it is plausible that risk lovers would be drawn to pay-to-bid auction websites by the thrill of winning, and that risk averse and risk neutral individuals would learn to avoid the websites due to the losses generally incurred by participants. This selection story, suggested by Platt et al. (2013), may help to explain the large profits of pay-to-bid auction websites. I control for the possibility of a selection effect by drawing a random sample of participants from a large and diverse subject pool. I also measure risk preference to assess whether risk loving participants are a driving force behind the large seller profits.

3. Shill bidding

Another possible explanation for the large profits is the presence of shill bidders. A shill bidder is a participant who is employed by the seller to bid at key points in the auction to keep the auction alive and increase average revenues. Platt et al. (2013) show that when a shill bidder is present and attempts to blend in by playing the optimal bidding strategy, expected revenue is unchanged. There will be more bids when the shill bidder is present, but because the shill bidder is not charged bid fees and cannot actually buy the item this increase in bids does not affect revenue. However, as Byers et al. (2010) point out, if the shill bidder deviates from the optimal strategy by over-bidding at key points, this will increase expected revenue so long as the auction participants cannot detect the shill bidder's behavior. There are no shill bidders in this experiment and so this factor is controlled for.

4. Signaling strategies

Both Augenblick (2012) and Byers et al. (2010) note how signaling behavior may increase expected revenue. Suppose participant i employs the strategy: bid with certainty no matter what the price. Then the best response of all participants $j \neq i$ is to never bid. This asymmetric equilibrium, which exists for all choices of the auction parameters (n, Π, C, ε), results in participant i winning the item after a single bid for a profit of $\Pi - C - \varepsilon$. As this sort of equilibrium is quite favorable for participant i, participant i might try to signal this strategy (out of equilibrium) by bidding with high frequency and by immediately outbidding any opponents. If participant i is successful in implementing this strategy, then expected revenue will decrease substantially. However, if multiple participants attempting to play this strategy enter into a game of chicken, the resulting bidding war may serve to increase auction revenue. To assess the impact of signaling, I conduct some treatments where the ability to signal is limited (i.e. bid decisions are revealed simultaneously and the identity of bidders is not known) and in the other treatments, where signaling is not limited, I measure the amount of signaling and its effectiveness.

5. Behavioral Biases

Augenblick (2012) also suggests that sunk cost fallacy may help to explain the large profits of pay-to-bid auction websites. He shows that when participants experience regret from placing bids that do not win the auction and when these participants naively underestimate their future regret, this will increase auction revenue. Alternatively, the participants we observe in internet pay-to-bid auctions may be boundedly rational or computationally limited. The optimal bidding strategy is complex and must be solved numerically via backwards induction. For this reason, it is unlikely that a new participant would arrive knowing the optimal strategy. The more likely scenario is that a population of symmetric agents participating in a series of identical pay-to-bid auctions might learn over time to bid as if the symmetric equilibrium were being played. If this were the case and if new participants had a tendency to bid more frequently than is optimal given the number of other participants, expected revenue would increase in the short run, until behavior converged to the equilibrium level.

2.4 Experiment Design

2.4.1 Basics

I conducted 6 experimental sessions with a total of 192 subjects. Treatment variables are the number of participants per auction, the strategy space, and the number of auctions in a session. The remaining parameters are held constant across auctions. Each session had between 20 and 39 subjects who were randomly matched into groups of 3 or 5 at the beginning of every auction. Subjects participated in 1 practice auction and 8 real auctions in each the five standard length sessions and 1 practice auction and 20 real auctions in the long session (denoted by "L"). Thus, between 32 and 120 real auctions were conducted in each session and 456 auctions were conducted in total. The sessions took place during November of 2011 and January and April of 2012 in a computer laboratory at a large public university. The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007). Screenshots of the experiment are shown in Figure A.1. Subjects were students who had previously registered to be in the laboratory subject pool, and were recruited by E-mail solicitation. Subjects were not allowed to participate in more than one session. At the end of each session, subjects chose their preferred show-up payment and then answered a short questionnaire. The questionnaire asked a variety of demographic questions (i.e., age, academic status, courses taken, gender, and major), and asked how the subject decided when to bid and when not to bid. This information is summarized in Table 2.1.

I utilize a between subjects design rather than a within subjects design to ensure an adequate number of trials for behavior to converge in case of learning effects. In the first session, denoted by "Discrete3," bid decisions were made simultaneously in a series of discrete rounds and there were 3 participants per auction. In the second and third sessions, denoted by "Discrete5-1" and "Discrete5-2," bid decisions were made simultaneously in a series of discrete rounds and there were 5 participants per auction. In the fourth session, denoted by "Continuous3," bid decisions were made in continuous time and there were 3 participants per auction. In the fifth and sixth sessions, denoted by "Continuous5-1" and "Continuous5-2 (L)," bid decisions were made in continuous time and there were 5 participants per auction. The Discrete5 and Continuous5 treatments were split over two sessions because the computer laboratory could not accommodate the desired number of groups in a single session. As can be seen in Table 2.1, there were a large number of groups for each treatment condition and so it is unlikely that my results are driven by sampling variation. Consistent with this claim, I cannot reject that average revenue was identical across the two Discrete5 sessions or across the first 8 rounds of the Continuous5 sessions.

Subjects earned $1.00 for each 1.00 of experimental currency earned during the session. In

Table 2.1: Session Information

Session	Number of subjects	Number of groups	Percent male*	Percent with 1+ statistics courses*	Major*: Percent Business or Economics	Percent Psychology or Cognitive Science	Percent Other Social Science, Sociology, Criminology	Percent Physical Science / Engineering	Percent Life Science / Biology / Public Health	Percent Humanities or Undecided/undeclared	Average take-home earnings**
All	192	48	50%	60%	21%	7%	15%	28%	23%	6%	$22.66
Discrete3	33	11	39%	61%	30%	6%	18%	21%	18%	6%	$20.57
Discrete5-1	35	7	34%	60%	29%	11%	11%	31%	17%	0%	$19.43
Discrete5-2	20	4	60%	75%	25%	15%	30%	10%	20%	0%	$19.03
Continuous3	39	13	56%	41%	21%	0%	8%	36%	23%	13%	$21.46
Continuous5-1	35	7	49%	57%	9%	9%	14%	26%	34%	9%	$19.71
Continuous5-2 (L)	30	6	67%	80%	17%	3%	13%	37%	27%	3%	$36.16

Notes: * denotes information obtained from the questionnaire. ** includes the show up payment

addition, subjects received \$2.00 at the beginning of the session to prevent bankruptcy, as well as a payment for showing up. The amount of this show-up payment ranged from \$0.25 to \$12.25 and was determined based on the subject's preference over three lotteries and the outcome of the selected lottery. The average take-home amount was approximately \$20 for 80 minutes of participation in the standard length sessions and approximately \$36 for 120 minutes of participation in the long session.

2.4.2 A Single Auction

Subjects are randomly assigned to a group of 3 or 5 participants at the beginning of *each* auction and shown a "start" screen (Figure A.1 (a)), with information on the auction parameters. This group is not fixed and participants are randomly assigned to a new group at the beginning of the next auction. The auction is for a commodity with common value $\Pi = 2.00$. The auction starting price P_0, bid increment ε, and bid fee C are 0.00, 0.10 , and 0.15 respectively. Subjects begin each auction with an endowment (referred to in the software as "bank account") of 1.50, which is sufficient to ensure the subjects are never budget constrained during the auction.[12] These parameters are held constant across all sessions and treatments. After each subject indicates she is ready by pressing the "start" button, the auction begins. In an auction round, the subjects choose between bidding and not bidding. The auction continues until no bids are placed in a round or until $\Pi - P_t \leq 0$, because winning the auction ceases to be profitable in this range. At the conclusion of the auction, the high bidder receives the common value of the commodity less the end auction price. The resulting profit for the high bidder is given by $\Pi - P_t \equiv 2.00 - 0.10 \times (total\ \#\ of\ bids)$. In addition, each participant (including the high bidder) receives the residual value of her bank account as profit at the end of the auction. To test whether the losing bidders cease to participate over time, the subjects are paid for *each* auction, rather than a randomly selected auction as is commonly done when trying to induce one-shot game behavior.

2.4.3 A Single Auction Round

Each round begins with 15 seconds on the countdown clock and the countdown begins immediately. During a round, the subjects choose between bidding and not bidding. Subjects can place a bid at any time by pressing the "Bid" button. If a subject does not press the "Bid" button, then no bid is placed. Placing a bid costs 0.15 and increments the auction price by 0.10. The round ends when the countdown clock reaches zero and, in the continuous time treatments, whenever a bid is placed. At the end of a round, both the choices that were made and the outcome of the round are revealed. If a subject places a bid in a round, that subject becomes the next round's high bidder. The high bidder is not allowed to bid while she retains this title, but receives the commodity at the end auction price if no more bids are placed. If multiple bids are placed in a round, then the high bidder is chosen at random from the set of participants who placed a bid in this round.

[12] The endowment was given in pieces, rather than in a single lump sum payment at the beginning, so as not to encourage over-bidding. For example, Price & Sheremeta (2011) find more over-bidding in lottery contests when the endowment is given in a single lump sum payment.

2.4.4 Number of Auction Participants

One key treatment variable is the number of auction participants. The baseline theory suggests that expected revenue will not depend on the number of auction participants, as each individual will adjust her probability of bidding to account for the number of other participants. However, this may not be the case in practice. For example, Gneezy & Smorodinsky (2006) find that over-bidding is increasing in the number of participants in all-pay auctions. If participants are bad at adjusting the probability of bidding to account for the number of other participants, this could be one reason why pay-to-bid auction websites have revenue in excess of the value of the prize. The two auction group sizes used in this experiment are 3 and 5. These group sizes were chosen to ensure that a large amount of data could be collected (smaller groups yield more auction level data for a given number of participants), while also providing sufficient variation in the treatment variable. Because the high bidder may not place a bid, there is one less active participant after the first round. Thus, 3 is the smallest group size that can accommodate multiple participants in the later rounds of the auction. The group size 5 doubles the number of participants in the later rounds of the auction relative to 3.

2.4.5 Strategy Space

Another key treatment variable is the strategy space. The baseline theory assumes that participants act simultaneously in a series of discrete rounds, but pay-to-bid auctions on the internet are conducted in continuous time. Augenblick (2012) shows that the abstraction from continuous time to simultaneous rounds should not theoretically affect auction revenue because any equilibrium in the continuous time model can be converted into a payoff and outcome equivalent equilibrium in the discrete time model. However, this result still relies on the assumption that there is no additional informational content in the timing of a bid. Anyone watching a live pay-to-bid auction would see almost immediately that this assumption may be violated. Bids are clustered towards the beginning and end of the countdown clock, and the average timing of bids varies substantially across individuals. Wang & Xu (2013) find evidence in auction data that different levels of bidder sophistication play a role in this behavior. In addition, Goodman (2012) finds evidence of signaling behavior in auction data, suggesting that non-Markovian strategies are played. These deviations from equilibrium behavior in the baseline theory may have an effect on the revenue and so it is possible that differences in the strategy space can explain why pay-to-bid auction sites have revenues in excess of the value of the prize.

To test for this possibility I vary the strategy space between sessions. In one treatment, the subjects make decisions simultaneously in *discrete rounds* as modeled in the baseline theory, and in the other treatment, the subjects make decisions in *continuous time* as they would in a pay-to-bid auction on the internet. In the discrete rounds treatment, subjects have 15 seconds to decide between bidding and not bidding and the outcome of these decisions is not revealed until the countdown clock reaches 0. Because the timing of a bid has no strategic value in this context and bids are revealed simultaneously, this treatment is equivalent to the game modeled in the baseline theory. In addition, no identifiers are given to the subjects in this treatment so as to not encourage the use of Markov strategies (signaling strategies in particular). Omitting the identity of the bidder(s) each round makes it more difficult to use these Markov strategies effectively, and so the strategies played are more likely to conform with the strategies analyzed in the baseline theory. In the continuous time treatment, subjects also have 15 seconds to decide between bidding and not bidding, but now the outcome of the bid decision is revealed and the clock is reset as soon as a bid is placed. Here the timing of a bid does have strategic value as information is revealed instantaneously. The subjects are given identifiers in this treatment to allow for Markov strategies to be played.[13]

[13]Restricting the timing of decisions and withholding the identity of the bidder(s) both effectively reduce the set of strategies that can be played. Thus, the combination of these restrictions represents a single change in the strategy space.

2.4.6 Measuring Risk Preferences

Risk preferences may affect the bid probabilities used in equilibrium, and thus may affect expected revenue. For instance, Platt et al. (2013) show that when the pay-to-bid auction participants are risk-seeking, expected revenue may be greater than the value of the prize. To account for the effect these risk preferences have on expected revenue, I elicit the subjects risk preferences at the end of the experiment session. With this measure of risk aversion, I then categorize the subjects as risk averse, risk neutral, or risk loving. The lotteries are designed with simplicity in mind, in accordance with the method of Eckel & Grossman (2008). In particular the lotteries involve two outcomes, each with 50% probability, and the lotteries are increasing linearly in standard deviation.

Each subject is asked to select her preferred show-up payment from a set of three possible lotteries. The lotteries are presented in a random order to ensure there is no ordering bias to the selections. One lottery gives \$7 with probability 100%, the second lottery gives \$10.25 with probability 50% and \$4.25 with probability 50%, and the third lottery gives \$12.25 with probability 50% and \$0.25 with probability 50%. This decision is framed in the context of a choice of show up payment so that subjects will view the lottery payoffs separately from prior earnings in the experiment. If this is the case, then a subject with constant relative risk aversion (CRRA) coefficient greater than 0.4 should prefer the first lottery, a subject with CRRA coefficient less than -0.4 should prefer the third lottery, and a subject with CRRA between -0.4 and 0.4 should prefer the second lottery. Thus, the first lottery will generally appeal to risk averse subjects, the second lottery will generally appeal to risk neutral subjects, and the third lottery will generally appeal to risk loving subjects. If a subject does take her prior earnings into account when making her selection, this will serve to increase the range of coefficients for which she prefers the "risk neutral" lottery, making the measure cruder. Regardless of which case occurs, so long as the subjects are consistent in how they evaluate these lotteries (i.e., they all generally do not account for prior earnings or they all generally do account for prior earnings), this approach will appropriately categorize subjects. The most risk averse subjects will be categorized as risk averse, the most risk loving subjects will be categorized as risk loving, and the subjects in between will be categorized as risk neutral.

2.4.7 Hypotheses

I develop three testable hypotheses based on the predictions of the baseline theory.[14]

Hypothesis 1:

> For all treatments, average revenue conditional on sale $\bar{R}$ will be less than or equal to the value of the item Π ($\bar{R} \leq \Pi$).

As shown in Theorem 2.1 and Corollary A.1, $E(R)$ should be no more than the value of the item Π. This result does not depend on the number of participants n and it does not depend on the strategy space.

Hypothesis 2:

> For both group sizes ($n = 3, 5$), average revenue conditional on sale in the discrete round sessions $\bar{R}_d$ will equal average revenue conditional on sale in the continuous time sessions $\bar{R}_c$ ($\bar{R}_d = \bar{R}_c$).

The baseline theory suggests that equilibria in the continuous time setting are equivalent to equilibria in the discrete time setting in terms of expected payoffs. Thus average revenue conditional on sale should be the same across formats.

[14]In formulating Hypotheses 1 and 2, I assume that the participants are not risk loving in aggregate. As will be seen later, most subjects measure as either risk neutral or risk averse, and so this assumption appears to be satisfied.

Hypothesis 3:

> For both group sizes ($n = 3, 5$), risk averse individuals will bid the least on average, risk loving individuals will bid the most on average, and risk neutral individuals will be somewhere in between.

The baseline theory suggests that expected revenue will be lower when participants are risk averse and higher when participants are risk loving. Ultimately this manifests itself as a decrease in the aggregate number of bids when participants are risk averse and an increase in the number of bids when participants are risk loving. Thus, we would expect risk averse individuals to bid the least, risk loving individuals to bid the most, and risk neutral individuals to bid a middling amount.

2.5 Results

All hypothesis tests reported in this section are also summarized in Tables A.1 and A.2. Standard errors used for hypothesis testing are clustered at the session-auction level to account for possible dependencies within auction periods. The random re-matching procedure ensures no group level dependencies across auction periods.

2.5.1 Hypothesis 1

For each auction in which at least one bid was placed, I record the revenue of this auction in dollars. Auctions in which no bids are placed are excluded as the item was not sold.[15] This yields average revenue conditional on sale, which can be directly compared to the revenue predictions of the baseline theory. Average revenue is reported by session and auction number

[15] In practice, items that are not sold on a pay-to-bid auction website can be re-listed at negligible cost on a later date. Thus, it is easier to directly compute gross profits from revenues when non-sales are excluded.

in Table 2.2. Average revenues in the Discrete3, Discrete5-1, Discrete5-2, Continuous3, Continuous5-1, and Continuous5-2 (L) sessions are \$2.22, \$3.10, \$3.02, \$2.32, \$2.94 and \$2.90 (\$3.06 for the first 8 auctions) respectively. As the Discrete5-1 and Discrete5-2 sessions and the Continuous5-1 and Continuous5-2 (L) each experienced the same treatment, I first test whether average revenue is the same between these sessions. When comparing the average revenue over the first 8 auctions from each of these sessions, I cannot reject the null hypothesis of equal means (p-values 0.838 and 0.782 respectively). Thus I pool the data from these sessions when possible in the remaining hypothesis tests. The resulting averages for the first 8 auctions of the Discrete5 and Continuous5 treatments are \$3.07 and \$3.00 respectively. For all treatments, I test whether average revenue is less than or equal to the \$2 value of the item. I reject this hypothesis at the 0.01 level for the Discrete5 and Continuous5 treatments, and at the 0.05 level for the Continuous3 treatment. I cannot reject that average revenue in the Discrete3 treatment is less than or equal to \$2 (p-value 0.140). The above analysis casts doubt on Hypothesis 1, but it does not take into account the

Table 2.2: Average Revenue by Session and Auction Trial

	Auction Trial										
Session	1	2	3	4	5	6	7	8	9	10	Average
All	$ 3.10	2.47	2.98	2.47	2.93	2.05	2.52	2.22	--	--	$ 2.68
Discrete3	$ 2.25	2.17	3.00	1.63	3.04	2.41	1.75	1.78	--	--	$ 2.22
Discrete5-1	$ 4.14	2.70	3.38	3.42	3.39	1.83	2.96	2.60	--	--	$ 3.10
Discrete5-2	$ 1.58	3.19	3.25	2.83	4.31	2.42	2.92	3.17	--	--	$ 3.02
Continuous3	$ 2.88	2.13	2.96	1.67	2.50	1.96	2.54	1.94	--	--	$ 2.32
Continuous5-1	$ 4.29	2.93	2.54	3.82	2.32	1.86	3.11	2.68	--	--	$ 2.94
Continuous5-2 (L)	$ 3.50	2.25	3.95	2.38	4.33	3.79	1.85	2.38	2.96	3.40	$ 2.90
Continuous5-2 (L) (11-20)	$ 3.29	3.04	2.17	3.00	3.08	2.88	3.30	1.63	3.60	1.50	

Notes: Auctions in which the item was not sold are excluded.

possibility of learning over time. As can be seen in Figure 2.1, there is substantial variability in the average revenue across auction trials but average revenue tends to lie above the $2 value of the item, especially in early auctions. I find a significant downward trend in average revenues, controlling for possible serial correlation, in all but the Discrete5-2 session (p-value 0.588). This trend is significant at the 0.01 level for the Discrete5-1 session, the 0.05 level for the Continuous5-1 and Discrete3 sessions, and the 0.1 level for the Continuous5-2 (L) and Continuous3 sessions. To accommodate this possibility, I also test Hypothesis 1 for the last auction from each treatment using the estimated regression coefficients. Assuming average revenues of $2.00 in the prior round, predicted revenues in the last auction of the Discrete3, Discrete5-1, Discrete5-2, Continuous3, and Continuous5-1 and Continuous5-2 (L) sessions are $1.41, $2.57, $2.86, $2.19, $2.64, and $2.53 respectively. For all sessions, I test whether these predicted values less than or equal to the $2 value of the item. I reject this hypothesis at the 0.05 level for the Discrete5-1 session and at the 0.1 level for the other 5 participant sessions. I cannot reject this hypothesis for the Discrete3 or Continuous 3 sessions (p-values 1.000 and 0.156). Although I cannot reject that the predicted values in the last period of the 3 participant sessions are less than or equal to $2, it is worth noting that in all but one case the predicted value lies above the theoretical prediction. This is even the case in the 20 round Continuous5-2 (L) session. That revenues have yet to converge to the $2 theoretical prediction even after such a lengthy period is indicative of a broader behavioral phenomenon. One possible explanation for this behavior, which is consistent with the data and prior research on contests, is that subjects have difficulty learning to bid at an optimal level. Once involved in the bidding, these subjects exhibit a persistent bias towards over-bidding. However, as I will show later, individual subjects may learn to avoid participation altogether, especially after incurring large losses. The consistent drop in revenue in the later rounds of a session also happens to coincide with the largest drop in observed participation. Thus it appears that subjects are eventually learning not to bid at all, rather than to bid less.

As average revenues in the treatments with 5 participants are significantly greater than

Figure 2.1: Average Revenue by Auction Trial

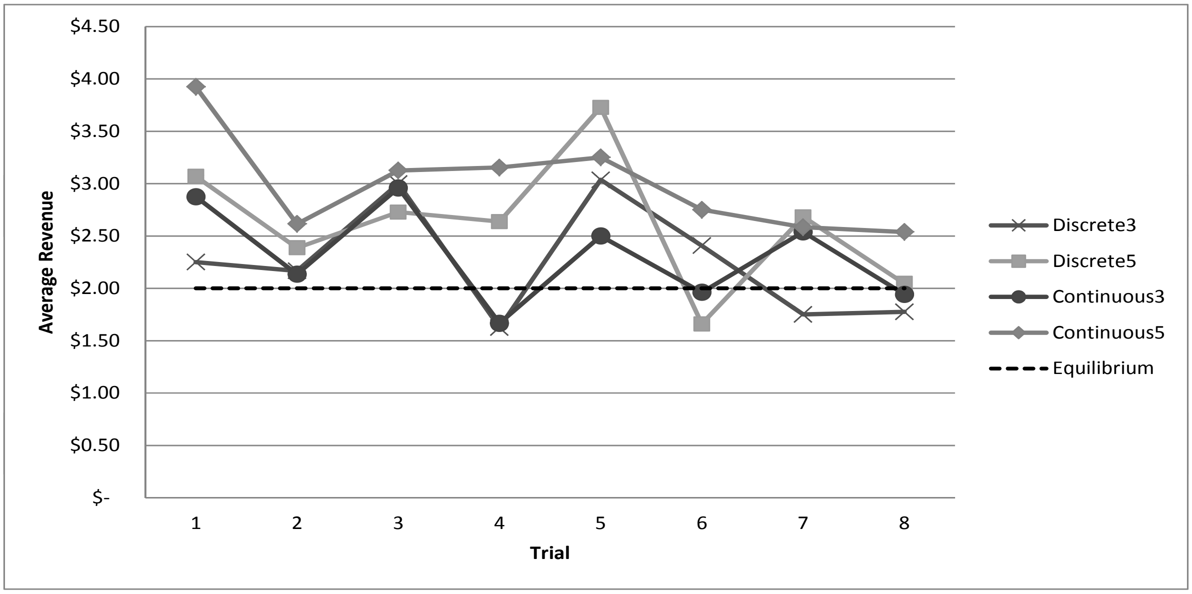

(a) All Treatments

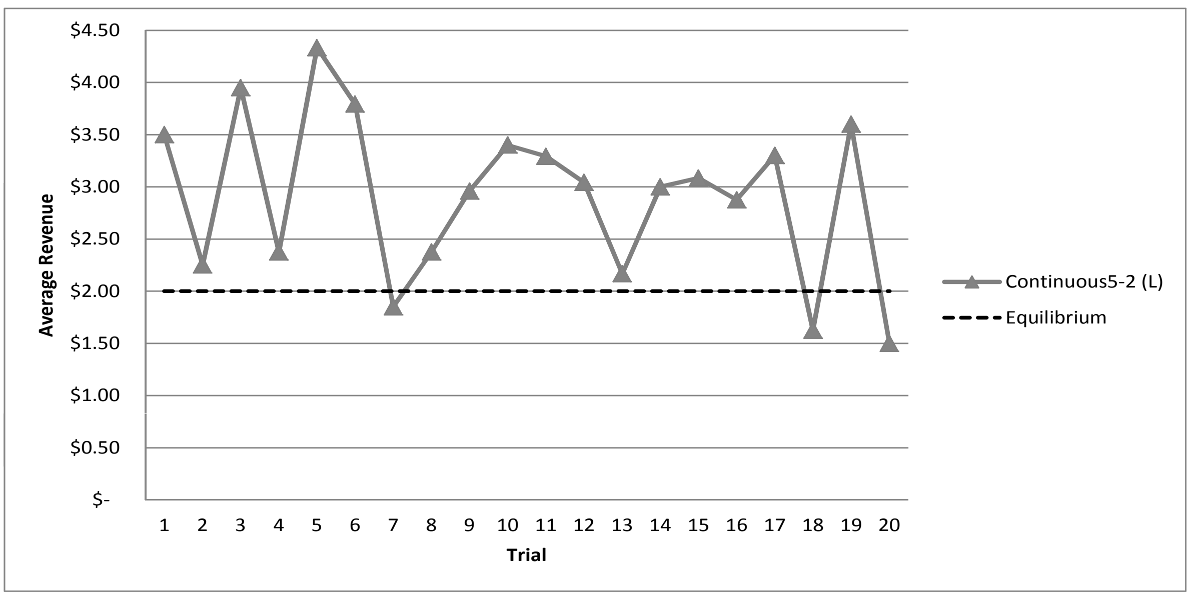

(b) Continuous5-2 (L) Session

the value of the item and this difference remains in the later auctions, Hypothesis 1 is *rejected*. Subjects exhibit a persistent bias towards over-bidding. This bias becomes less severe as subjects gain experience, but it does not disappear even with moderate levels of experience. Also, in the discrete round treatments, the aggregate effect of this bias depends on the number of participants n. I find that average revenue is greater when $n = 5$ than when $n = 3$ (significant at the 0.01 level). This finding may help to explain why pay-to-bid auction websites have such large average profits. These auctions generally have a large number of bidders and so even a small individual bias towards over-bidding may lead to a lot of over-bidding in aggregate. However, if bidders learn to stop participation altogether, this can lead to steep declines in average revenue as is seen in the later auctions of each treatment.

It is also worth noting that these revenue estimates are a lower bound on the true bias towards over-bidding. In this experiment, participants were not allowed to bid once the auction price reached \$2 as this could potentially lead to bankruptcy.[16] This limit was reached in 57 of the 456 auctions I conducted (15 of which occurred sometime after the first five auctions) even though any bid placed after the price reached \$1.80 was guaranteed to lose money. This suggests that participants may have been willing to bid past \$2 if given the opportunity. While this behavior is seemingly irrational, it is consistent with findings in other experimental settings (e.g., Gneezy & Smorodinsky 2006, Herrmann & Orzen 2008).

2.5.2 Hypothesis 2

For both levels of the number of participants n, I compare average revenue in the discrete round treatment to average revenue in the continuous time treatment. In both cases, I cannot reject the null hypothesis that average revenue is the same across the two treatments (p-values 0.680 for $n = 3$ and 0.507 for $n = 5$). I interpret this as *confirming* Hypothesis 2. This finding is encouraging as it suggests that the abstraction from continuous time to discrete rounds does not fundamentally change the auction dynamics (at least when using the model where multiple bids can be accepted). Thus, further efforts to develop the theory

[16] The obvious disadvantage of this restriction is that it places an upper bound on the observed over-bidding. However, with this restriction in place, more data can be collected in the allotted time (more rounds = longer auctions). This restriction also ensures that participants are never budget constrained given the initial endowments and allows for a larger proportion of earnings to be tied to decisions rather than given in the form of a flat payment to prevent bankruptcy.

of pay-to-bid auctions can still make headway using this simpler framework.

2.5.3 Hypothesis 3

I code each participant as risk averse, risk neutral, or risk loving based upon the participant's preferred lottery. In total, 84 subjects measured as risk averse, 85 subjects measured as risk neutral, and 23 subjects measured as risk loving. I then compute the average number of bids placed per auction by the participants in each class of risk preferences. In order for average revenue to equal the $2 prize, the average number of bids placed by each participant per auction would need to be 2.67 in the $n = 3$ treatments and 1.6 in the $n = 5$ treatments. As can be seen in Figure 2.2, the average number of bids placed per auction by risk averse, risk neutral, and risk loving participants in the $n = 3$ treatments is 2.4, 3.1, and 2.4 respectively. I reject the null hypothesis that all 3 populations are equal at the 0.05 level. Risk neutral is significantly greater than risk averse (0.01 level) but risk loving is not significantly greater than risk neutral (p-value 0.936). The average number of bids placed by risk averse, risk neutral, and risk loving participants in the $n = 5$ treatments is 2.0, 2.4, and 2.3 respectively. In this case, I cannot reject the null hypothesis that all 3 populations are equal (p-value 0.428). Thus, I find weak evidence that risk preferences do affect the frequency of bids. As expected, risk averse participants bid less than risk neutral participants, but I cannot reject that risk seeking participants bid less than or equal to risk neutral participants. It is difficult to draw broad conclusions from the behavior of the risk seeking participants in particular due to the limited number of observations in this category. Only 7 participants measure as risk neutral for the $n = 3$ treatments and 16 measure as risk neutral for the $n = 5$ treatments. In this limited sample, I do not find evidence that risk seeking participants bid systematically more than risk neutral participants.[17] As the hypothesized relationship holds in the $n = 3$

[17] I see no theoretical justification as to why risk seeking participants would bid less than risk neutral participants. One possibility is that the mean of this relatively small sample of risk lovers is not indicative of the population mean. Another possibility is that the risk loving category is capturing some unobservable factor in addition to risk loving preferences. For example, the 12 risk loving subjects from the continuous

Figure 2.2: Bids by Risk Preference

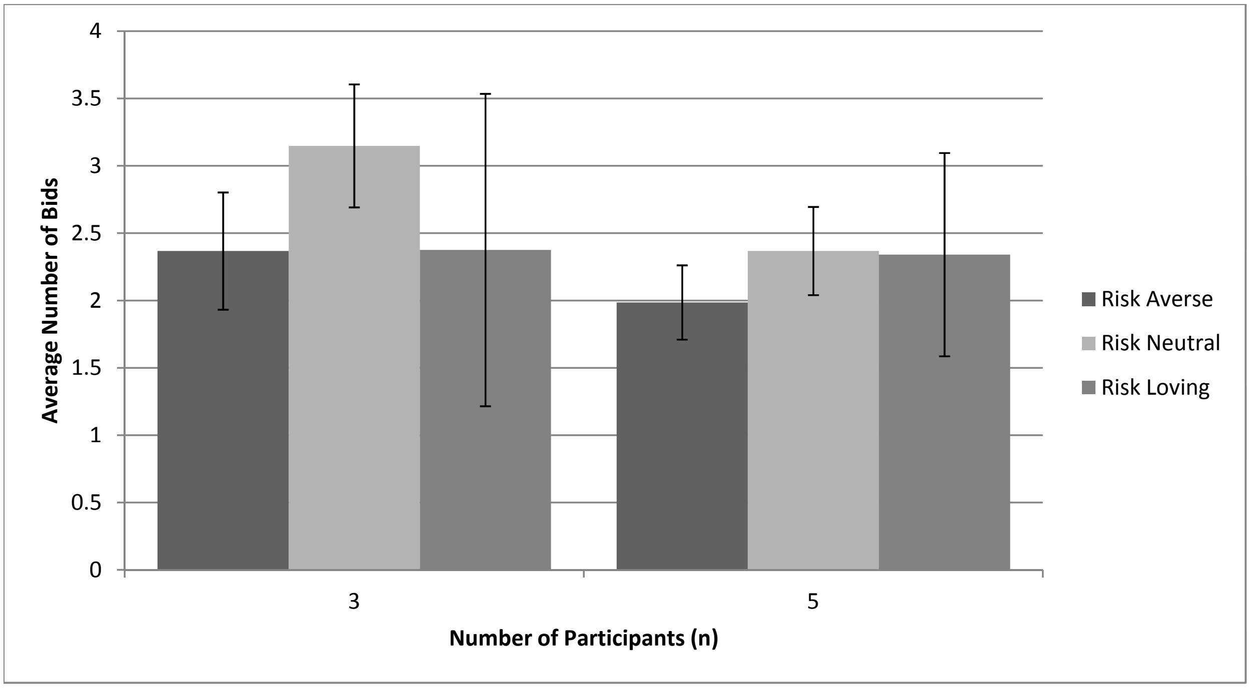

treatments for risk averse and risk neutral subjects, but not for risk loving subjects, and there is not a significant order effect in the $n = 5$ treatments my finding for Hypothesis 3 is *inconclusive*. I do however find strong evidence that risk seekers are not *necessary* for pay-to-bid auctions to generate revenues well above the value of the prize. In the $n = 5$ treatments, even risk averse participants average more than 1.6 bids/auction it takes for the seller to break even. In the $n = 3$ treatments, risk neutral participants bid more than the 2.67 bids/auction it takes for the seller to break even, and while risk averse participants bid slightly less than this threshold, the average number of bids per item sold is 2.63 yielding average losses of a mere $0.03/auction.

time treatments measure as less sophisticated on average than the other subjects (although not significantly so).

2.5.4 Individual Outcomes

Wang & Xu (2013) and Goodman (2012) use the timing of bids to show that factors such as sophistication and signaling behavior can have a large impact on individual auction outcomes. Key to both analyses is the strategic value of the timing of a bid in the continuous time setting.

Consider a single sophisticated participant, who is cognizant of the strategic value of the timing of a bid, and $n - 1$ unsophisticated participants, who act at random times in the countdown clock.[18] Suppose that the sophisticated participant is currently a non-leader and the expected continuation value for a non-leader at the current auction price is greater than or equal to 0 ($V(p_t) \geq 0$). First consider the case where the expected continuation value for a non-leader at p_{t+1} is greater than or equal to the expected continuation value for the high bidder at p_{t+1} less the cost of bidding ($V(p_{t+1}) \geq V^*(p_{t+1}) - C$). In this case, if another participant will place a bid, the sophisticated participant prefers not to place a bid. Thus, if the sophisticated participant knows the other participants are bidding at random times, that participant should wait until the last second to place a bid to ensure that she only places a bid when no other participant will bid. Next consider the case where $V(p_{t+1}) < V^*(p_{t+1}) - C$. In this case, the sophisticated participant prefers to act first to ensure that she is the high bidder at price p_{t+1}. Thus, the sophisticated participant is able to improve her own individual outcomes by strategically bidding either early or late in the countdown clock depending on which role, high bidder or non-leader, is more advantageous in the next round.

While in practice not all opponents will be unsophisticated, the presence of some less sophisticated participants gives an advantage to those who use the timing of bids strategically. In particular, bidding in the middle of the clock is strategically inferior when the population

[18]While this example differs from that of Wang & Xu (2013), it is easy to understand and captures the gist of their argument.

contains a sufficient number of unsophisticated participants, and so Wang & Xu propose using the proportion of middle bids as a proxy for strategic sophistication. They show that a lower proportion of middle bids is associated with higher profits, suggesting that strategic sophistication plays a role in individual outcomes. In addition, participants may be able to use the timing of a bid as a signal. Bidding early in the countdown clock and bidding frequently could signal strength (i.e., a high value, lower costs, a larger endowment) and aggression. Goodman finds that participants who effectively employ signaling strategies tend to fare better than participants who do not signal in pay-to-bid auctions on the internet. This suggests that a participant's perceived strength and aggression may play a role in individual auction outcomes. In my setting, because each participant is symmetric, signals of strength are not credible. However, signals of aggression may still be credible, and so I test for the effect of these signals. Following the method of Wang & Xu, I use proportion of

Figure 2.3: Bid Time Histogram

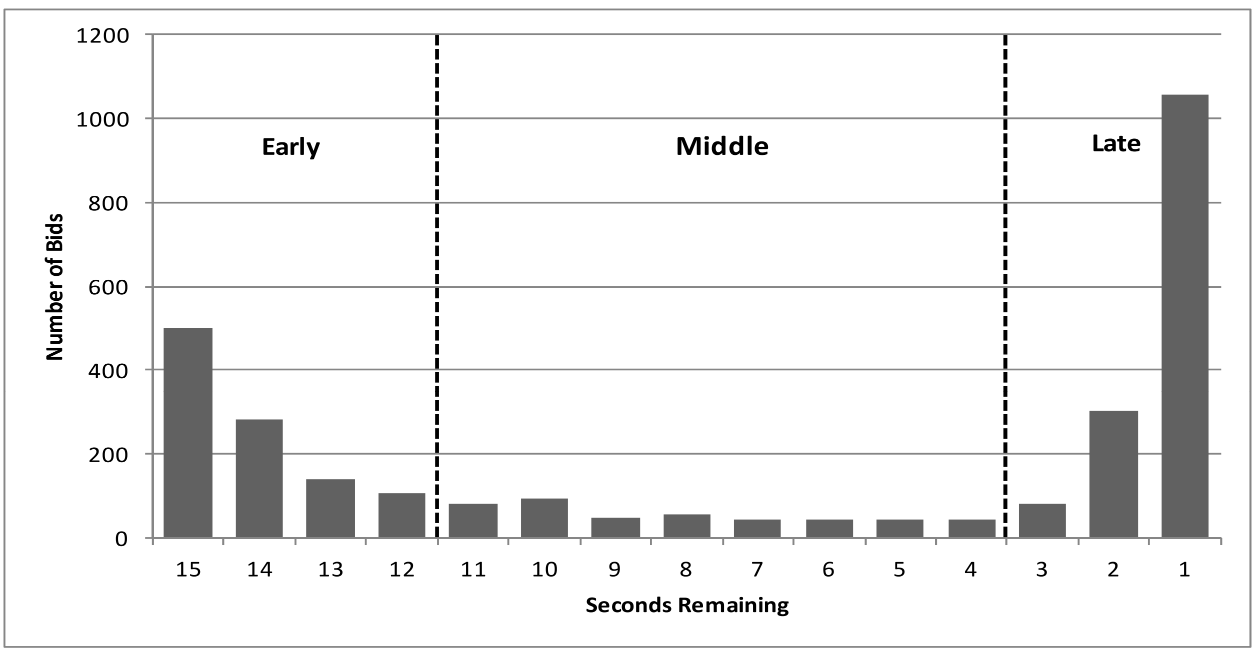

middle bids as a proxy for sophistication. I classify a bid as a "middle bid" if it is placed from the 5th second to the 12th second of the 15 second countdown clock. I chose this range conservatively based on the bid timing histogram illustrated in Figure 2.3. As reported in

Table 2.3 and illustrated in Figure 2.4, I find strong evidence supporting the relationship between strategic sophistication and individual outcomes.[19] Proportion of middle bids has a strong negative correlation with the net earnings of subjects, as defined by total earnings less initial endowment across all auctions. This relationship is significant at the 0.01 level for the $n = 3$ treatment, and at the 0.05 level for the $n = 5$ treatment. The least squares estimates suggest that a 1% increase in the proportion of middle bids is accompanied by an average \$0.04 to \$0.08 reduction in net earnings. Furthermore, I find no statistically or economically significant relationship between the proportion of middle bids and number of bids for the $n = 3$ treatment (p-value 0.96) and only a slight (0.4%/bid) relationship for the $n = 5$ treatment, so this finding is robust to the frequency of bids. As described by Good-

Figure 2.4: Net Earnings by Strategic Sophistication

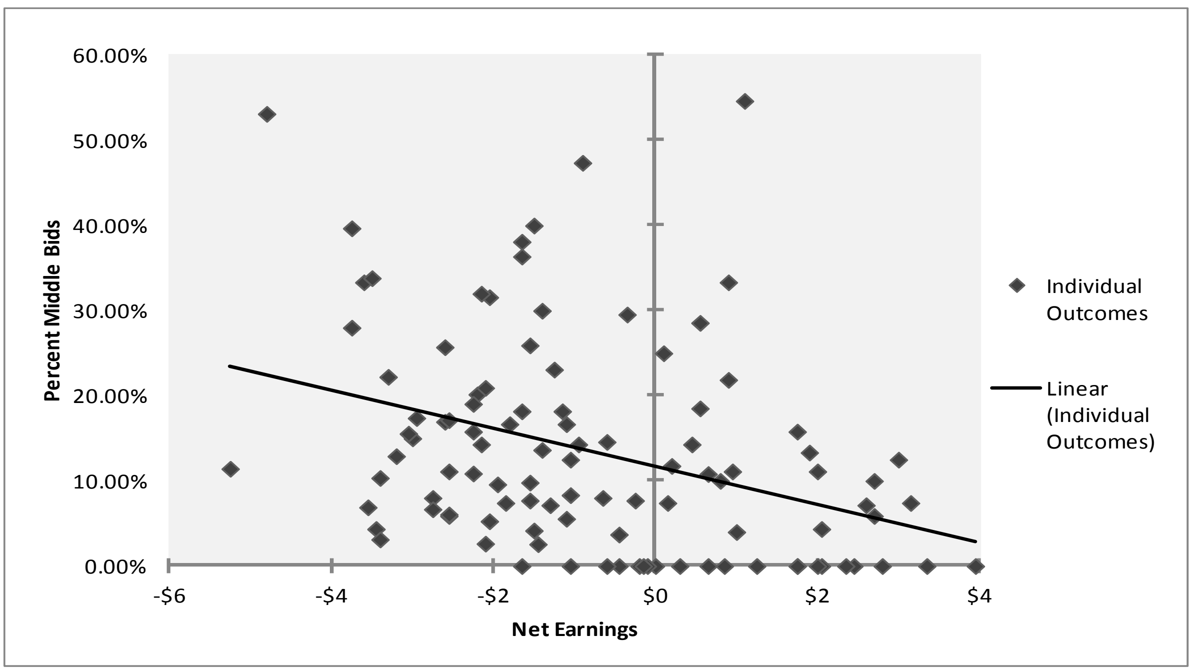

man (2012), "bidding runs" and "bid speed" are two potential ways in which a participant could try to signal aggression. A bidding run is when a participant places multiple bids in a row. For the purposes of this analysis, I consider any bidding run of length 3 or greater to be a potential signal. At most, 20 bids will be placed in any given auction, and so a run of

[19]Values from the Continuous5-2 (L) session are taken from the end of the 8th auction for ease of comparison. Using the values from the 20th auction does not substantively change the results.

Table 2.3: Strategic Sophistication Regressions

	Dependent Variable: Net Earnings				Dep Var: % Middle Bids	
	n=3		n=5		n=3	n=5
	(1)	(2)	(3)	(4)	(5)	(6)
% Middle Bids	-8.421***	-8.476***	-3.620**	-2.026		
	(2.680)	(2.483)	(1.696)	(1.632)		
# of Bids		-0.087**		-0.079***	-0.0001	0.004**
		(0.033)		(0.023)	(0.002)	(0.002)
Constant	0.953**	3.033***	-0.691*	0.591	0.121**	0.069*
	(0.451)	(0.884)	(0.327)	(0.480)	(0.055)	(0.036)

Notes: *significant at 0.1 level, **significant at 0.05 level, ***significant at 0.01 level.

length 3 makes up a large proportion of the total bids in an auction. In my data set, runs are most commonly of length 3 or 4 but sometimes are as long as 10. One complication is that a bidding run may not always be intended as a signal. We might also observe a bidding run if two participants enter into a bidding war or if a participant succumbs to sunk cost fallacy. To be conservative, I only consider a participant as having used a signaling strategy if that participant has at least one bidding run in the session and also uses bid speed to signal. Bid speed is the timing of a bid on the countdown clock. Earlier bids generally send a signal of aggression, although as Goodman notes, when a bid is recorded at the 15 second mark it is difficult to determine whether this is a fast bid or a result of nearly simultaneous bids in the prior round. For this reason, and due to the relatively high proportion of bids that are placed at the 1 second mark and the 15 second mark, I only code a participant as using the fast bid signal if that participant places more early bids (first 4 seconds) than late bids (last 3 seconds).

In total, 8 out of 39 subjects in the Continuous3 treatment and 10 out of 65 subjects in the Continuous5 treatment engaged in signaling behavior. As the goal of signaling in this context is to discourage the participation of other bidders, and there are fewer participants per auction in the Continuous3 treatment, we might expect that participants who signaled in the Continuous3 treatment fared better than participants who signaled in the Continuous5 treatment. I find this to be the case. Participants who signaled in the Continous3 treatment averaged net earnings of $-\$0.37$ which is significantly greater (0.05 level) than the $-\$2.10$ average net earnings of participants who signaled in the Continuous5 treatment. Relative to their counterparts who did not signal, participants who signaled did not fare very well. Participants who did not signal in the Continuous5 treatment averaged net earnings of $-\$1.06$ which is significantly greater (0.05 level) than the average net earnings of the participants who signaled in this treatment. Participants who did not signal in the Continuous3 treatment averaged net earnings of \$0.04 which is greater, but not significantly so (p-value 0.32), than the average net earnings of the participants who signaled in this treatment. As reported in Table 2.4, the actual effect of signaling is unclear. When controlling for the number of bids placed, signaling has a positive but insignificant effect on net earnings in both treatments. However, participants who signaled placed 10 more bids on average than participants who did not signal, and net earnings are decreasing in the number of bids placed. If these excess bids can be attributed to the cost of signaling, it appears that signaling was not an effective strategy. The apparent ineffectiveness of signaling in this setting is not entirely surprising.

Table 2.4: Signaling Regressions

	Dependent Variable: Net Earnings				**Dep Var: # of Bids**	
	n=3		**n=5**		**n=3**	**n=5**
	(1)	(2)	(3)	(4)	(5)	(6)
Signal	-0.412	0.357	-1.039*	0.024	8.343**	12.096***
	(0.893)	(0.911)	(0.597)	(0.633)	(3.438)	(2.718)
# of Bids		-0.092**		-0.088***		
		(0.040)		(0.026)		
Constant	0.044	2.076**	-1.056***	0.457	22.032***	17.204***
	(0.404)	(0.970)	(0.236)	(0.494)	(1.557)	(1.074)

Notes: *significant at 0.1 level, **significant at 0.05 level, ***significant at 0.01 level.

In designing this experiment, I have abstracted away from three important factors that could make signaling effective in a pay-to-bid auction on the internet. First, in an effort to generate more data, I chose the auction parameters so that an auction would last no more than 20 rounds. In many pay-to-bid auctions on the internet, the auction could last for thousands of rounds, giving participants much more leeway to signal through bidding runs. For example, bidding run of length 50 sends a much more effective signal than a bidding run of length 3. Thus, as a consequence of the auction parameters I have chosen, it is more difficult to signal through bidding runs. In addition, in this setting all participants have equal endowments and are aware that they have equal endowments. This rules out the possibility that the signaling party is a strong participant with a large endowment that will muscle out the competition, a possibility that likely adds to the effectiveness of signaling in pay-to-bid auctions on the internet. Even if participants have different intrinsic levels of aggression, the signaling strategies that can be employed in this setting lend limited credibility. It is not very costly for an intrinsically passive participant to signal aggression with fast bids and a bidding run of length 3 or 4. Lastly, participants are assigned new identifiers at the beginning of every auction. This precludes the possibility of a participant building a reputation across auctions, a possibility that may aid in the effectiveness of signaling in pay-to-bid auctions on the internet. So, while I find that signaling is not effective in the setting of this experiment, this does not rule out the possibility that signaling is effective in pay-to-bid auctions on the internet. In fact, evidence suggests that signaling is quite effective in the internet setting (Goodman 2012).

2.5.5 Attrition

As suggested by Wang & Xu (2013), persistently unsuccessful participants may learn to stop participating all together. To test for this possibility, I analyze each subject's participation decisions. If a subject places at least one bid within an auction, then it is clear that this subject participated in the auction. On the other hand, if a subject does not bid in an auction, this may or may not mean the subject participated in the auction. We only observe when a subject *does* bid and not when a subject *would* bid under the right circumstances. The latter is the criterion of interest in this analysis. To avoid over-estimating the level of attrition, I consider three conservative measures.[20] The measures described below are ordered from least to most conservative. In all measures I consider the decision not to participate to be irreversible. In the first measure of attrition, I code any subject who has bid at least once during the session (only 4 out of 192 subjects never placed a bid) as participating in every auction up until the point where that subject never bids again. So if a subject bids once in the first auction, and once in the eighth auction, then I consider this subject to have participated in every auction in the session. On the other hand, if a subject does not place any bids in the seventh and eighth auctions, then I consider this subject to have participated in the first 6 auctions. This measure will limit the bias towards over-estimating attrition, although it is still subject to error in the later auctions of the session because we do not observe what would have happened had the session continued. As can be seen in Figure 2.5 (a), the first instance of attrition occurred in the second auction, and by the last auction 53 out of the 188 initially participating subjects had ceased participation by this measure.

One potential problem with this measure is that some auctions end at low prices (even $0) and we may not have enough evidence to plausibly claim that a subject attrited in an auction for which we have observed so few decisions. To account for this, in the second measure of attrition, I code any subject who has bid at least once during the session as participating in every auction up until the point where that subject never bids again AND at least one auction has elapsed in which n or more bids were placed. This restricts the attrition analysis to cases where each subject had an opportunity to bid at least once on average. As can be seen in Figure 2.5(b), by this measure, the first instance of attrition occurs in the third auction, and by the last auction 37 out of the 188 initially participating subjects had ceased participation. The third measure of attrition is the same as the second measure of attrition,

[20] I use the term attrition synonymously with permanent non-participation given the play of opponents. Thus, I consider any subject who finds her persistent best response to be "never bid" to have attrited.

Figure 2.5: Attrition Over Time

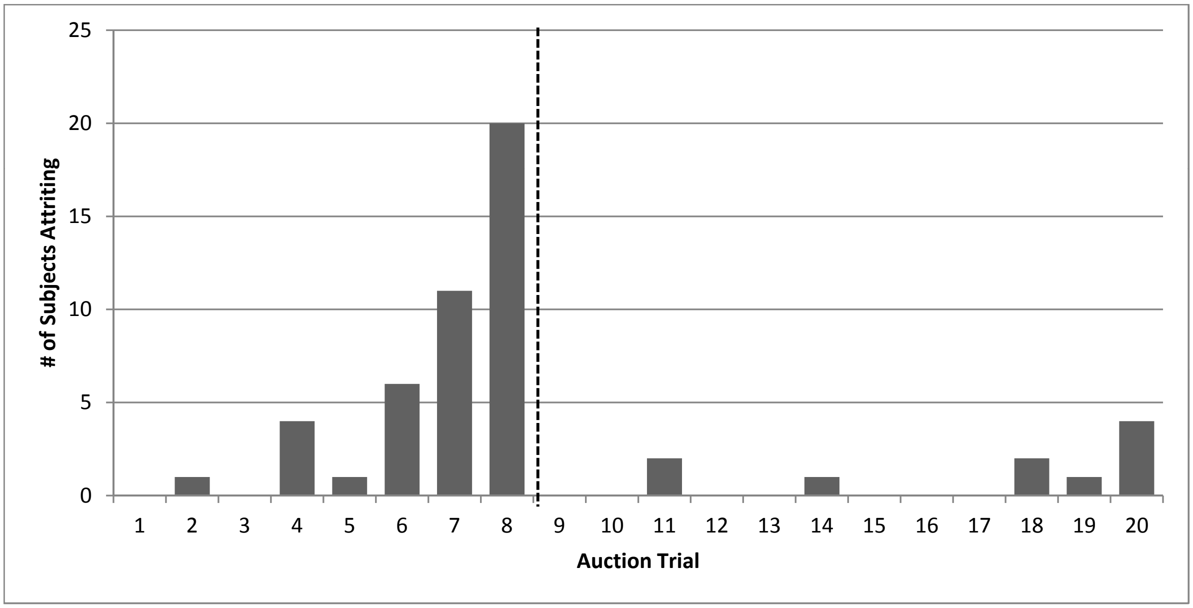

(a) Measure 1

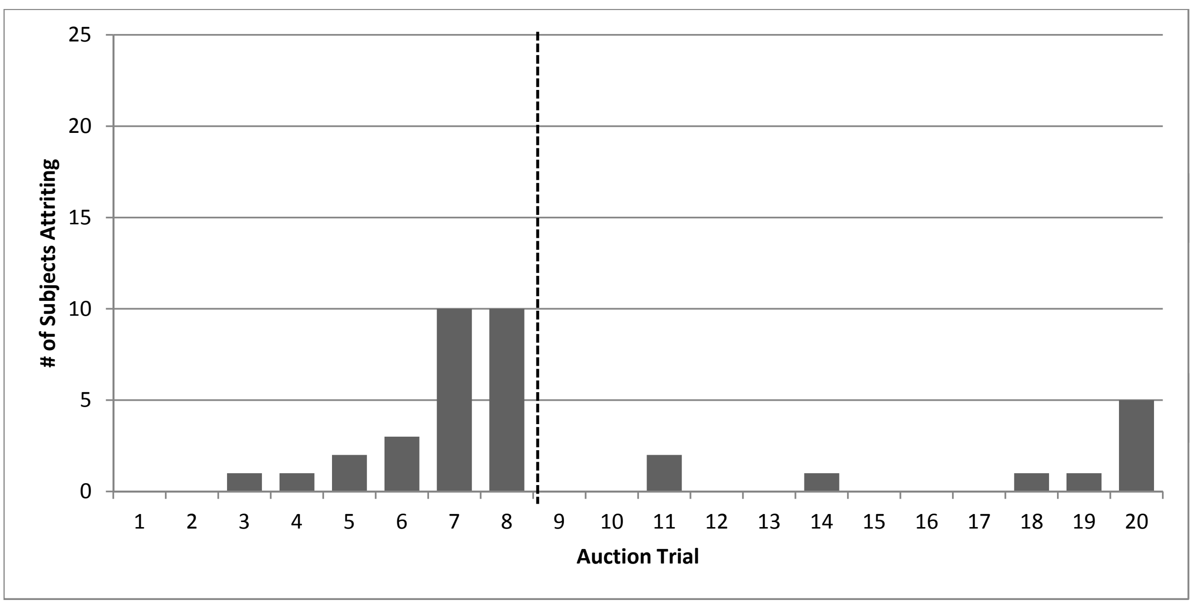

(b) Measure 2

but with the additional assumption that everyone who participated in the second to last auction also participated in the last auction. Due to the uncertainty about what would have happened if the session had continued, the last auction is the most susceptible to overestimating attrition. By dropping the last auction from consideration, these questionable cases are no longer considered. The third measure of attrition cuts the number of subjects who ceased participation down to 22 out of 188. It is unclear whether the second and third measures of attrition are over-estimates or under-estimates of the true level of attrition, but these measures seem to strike a balance between the different trade-offs.

All 53 subjects who ceased to participate in at least one of the three measures had cumulative losses at the time that they stopped participating. For each of the three measures, I compare the net earnings of the subjects who ceased to participate to the net earnings of the subjects who participated throughout the session, adjusted to account for the aggregate trend in net earnings. In all cases, the average net earnings of the subjects who ceased participation is lower than the average net earnings of the subjects who participated throughout. This difference is significant at the 0.05 level for the first and second measures, and significant at the 0.1 level for the third measure. Across all auctions, subjects who ceased participation averaged \$0.58 - \$0.72 less in net earnings than the subjects who participated throughout. This finding is in strong support of Wang & Xu's (2013) hypothesis. I observe that persistent money losers do cease to participate and cease participation at a relatively high rate.

2.6 Conclusion

The main take away from this study is that the large profits of pay-to-bid auction websites can be explained simply through mistakes, and inexperience. While it is certainly possible that information asymmetries, imperfect information, risk preferences, and signaling strategies may further increase revenues, none of these dynamics are needed to observe over-bidding. Thus, the simple setting used in the baseline theory of Hinnosaar (2010), Augenblick (2012), and Platt et al. (2013) captures the essence of the strategic environment of a pay-to-bid auction. Consistent with Wang & Xu (2013), I find that the symmetric equilibrium concept is the main deficiency in the current theory. There is substantial heterogeneity in the strategies employed by participants and these participants vary greatly in terms of strategic sophistication. While the population as a whole may slowly approach the equilibrium outcome where average revenue equals the value of the item, the path it takes to get there in no way resembles symmetric equilibrium play. Thus even with symmetric agents, the symmetric equilibrium is no more meaningful empirically than any of the many asymmetric equilibria (other than analytical convenience). Furthermore, a combination of learning and the long run attrition of less successful participants threatens the sustainability of the pay-to-bid auction as a mechanism to generate revenues above the value of the auctioned item. These excess revenues will only last as long as the pay-to-bid auction websites can attract new, inexperienced bidders.

As behavioral biases appear to drive the excess revenues of pay-to-bid auction websites, the pay-to-bid auction mechanism may be utilized as a tool in future research to study behavioral biases more generally. In addition, the highly competitive pay-to-bid auction industry may present an interesting test bed for researchers interested in behavioral game theory, as these websites have developed a variety of related auction mechanism that may further magnify behavioral biases.

Chapter 3

An Experimental Study of Network Formation with Limited Observation

Many social and economic networks emerge among actors that only partially observe the network when forming network ties. We ask: what types of network architectures form when actors have limited observation, and does limited observation lead to less efficient structures? In this chapter, co-written with Michael McBride, we report numerous results from a laboratory experiment that varies both network observation and the cost of forming links. Overall, we find that limited network observation does not inevitably lead to highly inefficient networks but instead might actually inhibit inefficient positional jockeying among actors.

3.1 Introduction

Amy is not employed and looking for work. She asks her acquaintances to relay information they learn about job openings in the hopes of finding a job that matches her skills and experience. She also actively seeks to form new acquaintances outside her current social network to broaden her awareness of job openings. Extending her social network is costly and only yields benefits if the new acquaintances are well-connected in social groups outside her existing network. However, Amy only partially observes her own network: she knows her acquaintances and her acquaintances' acquaintances, but not her acquaintances' acquaintances' acquaintances and so on. Because the observational range of her network is limited, she may unknowingly devote a substantial amount of her scarce time forming ties with people already in her network but outside her observational range. This inefficient job search results in chronic unemployment.

Amy's situation illustrates two important features of economic life. The first is that many economic activities, from the exchange of goods and services to the transmission of valuable information,[1] rely on underlying social networks that form endogenously through the decentralized forming and severing of network ties. Recognition of this fact has led a large theoretical literature on the formation of networks in various settings,[2] and, more recently, a small but growing experimental literature.[3] The second feature is that an individual generally observes only a limited and local portion of her social network, a fact established from field data.[4] However, only a small theoretical literature examines how the presence of limited observation affects network formation,[5] and there is no experimental literature on the topic.

A number of questions thus remain unanswered: What types of network architectures emerge

[1]There are many examples of which a few are listed here: Conley & Udry (2001) study learning about production techniques through informal networks in developing countries; Calv-Armengol & Jackson (2004) examine learning about job openings through social contacts; Kranton & Minehart (2001) model exchange in buyer-seller networks. For a larger survey of the literature on the economics of social networks, see Jackson (2010).

[2]See e.g., Bala & Goyal (2000) for a model of network formation in which ties are formed unilaterally and Jackson & Watts (2002) for a model of network formation in which ties are formed by mutual consent.

[3]Examples of the experimental literature include Falk & Kosfeld (2012), Callander & Plott (2005), Goeree, Riedl & Ule (2009), Pantz (2006), Mantovani, Kirchsteiger, Mauleon & Vannetelbosch (2013), Rong & Houser (2013), van Leeuwen, Offerman & Schram (2013), Carrillo & Gaduh (2012).

[4]Laumann (see 1969), Friedkin (see 1983), Kumbasar, Romney & Batchelder (see 1994), Bondonio (see 1998), Casciaro (see 1998).

[5]McBride (see 2006*a*,*b*, 2008), Francetich & Troyan (see 2010), Song & van der Schaar (see 2013). There is a separate theoretical literature on learning in networks, for example Acemoglu, Dahleh & Lobel (2011).

when actors have limited observation while forming and severing ties? Does the presence of this limited observation lead to less efficient network structures than would form with full observation? If so, what form do those inefficiencies take, and how large are the inefficiencies? Answers to these questions will provide a better understanding of the effectiveness of social networks in fostering or inhibiting economic activity and provide insights into potential efficiency gains. For example, an intervention that allows Amy to increase the observational range of her own network may help her allocate her time more effectively when forming new ties.[6]

This chapter addresses these questions with the first experimental study of social network formation in which decision makers have limited observation of the network. We choose for our setting the Connections Model first introduced by Jackson & Wolinsky (1996). Actors decide with whom to form costly ties, with each tie providing direct and indirect access to benefits. Mutual consent is required for a tie to form, but a tie can be terminated unilaterally. This model does not represent a specific real-life network setting but rather is a flexible model that captures features of many real-life social networks. It has also been subsequently theoretically and experimentally studied in later work (Jackson & Watts 2002, Watts 2001, Pantz 2006, Calv-Armengol & Ilkili 2009, Mantovani et al. 2013, Carrillo & Gaduh 2012), including in a limited observation context (McBride 2006*b*, Francetich & Troyan 2010).

Our experiment randomly matches subjects into groups of twelve for fifteen rounds of interaction. At the end of each round, a participant receives one point for each other participant with whom she is directly or indirectly connected (i.e., a point for each other actor in her component) and pays cost c for each of her direct ties. We exogenously vary two treatment variables. One is *the level of observation*; under *full observation* the whole network is observed in each round, and under *neighbor observation* the actor only observes her own ties and the ties of her network neighbors. The other is *the cost of forming ties*; the cost c is

[6] The business development tool RelSci (www.relsci.com) does precisely this by helping individuals to discover potential clients and investors that they are indirectly connected to through mutual acquaintances.

either low ($c = 0.7$) or high ($c = 1.5$). This design differs from previous network formation experiments (see footnote 3) in two important ways. First, we allow for both full and limited observation of the network, whereas prior studies allow only full observation. Second, our twelve-person groups are much larger than the four or six-person groups used in prior experiments, allowing for a richer set of possible network structures to emerge and increasing the possibility of networks in which actors can only observe a limited portion of the entire network structure in the *neighbor observation* setting.

We rely on the prior theoretical and experimental literature on network formation to generate testable hypotheses. Jackson & Wolinsky (1996) define and apply the *Pairwise Stability* (PS) concept to identify stable networks in the Connections Model, and McBride's (2006*b*) *Conjectural Pairwise Stability* (CPS) concept generalizes PS to identify stable networks when actors in the Connections Model have limited observation. Analyses that more directly consider formation dynamics vary in what is assumed about the forward-lookingness of actors, though all prior work assumes full observation. In Bala & Goyal (2000) and Jackson & Watts (2002), for example, actors form ties myopically by looking only for best responses to immediately past or current period play; in other work, such as Herings, Mauleon & Vannetelbosch (2009), Grandjean, Mauleon & Vince (2011) and Mantovani et al. (2013) in which is defined the *Pairwise Far-sighted Stability* (PFS), the actors are forward-looking in anticipating the future network structure.

The static analysis and the myopic-actor dynamic analysis yield clearer predictions, yet other experiments reveal that subjects in network formation settings exhibit far-sighted behavior (Mantovani et al. 2013, Carrillo & Gaduh 2012, van Leeuwen et al. 2013), a fact we account for here. We synthesize the prior theoretical and experimental work and identify seven main hypotheses about convergence and efficiency to test in our experiment. 1. Networks will be more likely to converge under low cost than high cost due to the existence of (near) efficient (C)PS networks. 2. Networks will be more likely to converge under neighbor observation than full observation because the set of CPS networks is a proper superset of the set of PS networks. 3. Efficiency will be higher under low cost due to stronger incentives to maintain connectivity. 4. Cycles (redundant ties) will be more likely under neighbor observation because they may exist outside of observational range. 5. Stem actors (one tie) will be dropped more often under full observation than neighbor observation because they are more likely to reconnect with others. 6. Stem actors will be dropped more often under high cost because the cost of these ties exceeds the immediate benefit of maintaining them. 7. Risk averse individuals will be less likely to choose stem positions to avoid the risk of being dropped from the network.

Our main findings match the hypotheses in some cases but not others. With respect to convergence, we find evidence in favor of Hypothesis 1 but reject Hypothesis 2: convergence (defined here as five rounds of no change in the network) is more likely with low cost than high cost but is less likely under neighbor observation than full observation, contrary to the prediction. With respect to efficiency, the evidence largely supports Hypotheses 3,4, and 6: efficiency is higher on average under low cost, cycles are more prevalent with neighbor observation, and stems are more often removed under high cost. However, we do not find that stems are more often removed under full observation, contrary to the prediction of Hypothesis 5. We do not find a significant relationship between risk aversion and position in the network, calling hypothesis 7 into question.

We also report other findings of interest. Holding both cost and observation fixed, the networks that converge are more efficient than those that do not, thus suggesting that subjects are less likely to disrupt a well-functioning network even if it is not a PS network. Moreover, holding cost fixed, network efficiency in the last five rounds is higher, on average, under neighbor observation than full observation, and this is apparently due to less jockeying for advantageous positions in the later periods under neighbor observation. Overall, though convergence to inefficient networks is more common with neighbor observation, the lower incidence of jockeying led, surprisingly, to higher overall efficiency under neighbor observation than under full observation. We do find evidence of far-sighted behavior, as expected, but this behavior leads to even richer outcomes than those predicted by the existing solution concepts.

We draw a number of important lessons from our findings. Limited observation does not inevitably lead to inefficient networks as might be expected. To the contrary, limited observation may actually help efficiency because it leads to a reduction in jockeying among actors attempting to gain advantages associated with certain positions in the network. This reduction in jockeying keeps the network more connected than otherwise which, in our setting, outweighs the inefficiencies in the form of redundant ties that we might expect under limited observation. Whether or not limited observation is a help or hindrance to network formation in the real world will likely depend on the relative importance of under-connection versus over-connection. We also learn that far-sighted behavior may be a source of stability, instability, efficiency or inefficiency, all depending on the context. Inefficient jockeying may destabilize otherwise stable networks in the short run. Alternatively, when it promotes efficiency, the anticipation of future jockeying may stabilize otherwise *unstable* networks in the short run. In settings where stable networks are highly inefficient, far-sighted behavior may even lead to efficient instability.

3.2 Theory

3.2.1 The Connections Model

Consider a set of actors $N = \{1, ..., n\}$, each of whom is a *node* in graph g. Let $ij \in g$ denote that there is a tie between i and j in g, while $ij \notin g$ means there is no tie. We assume that the ties are symmetric, i.e., $ij \in g \Leftrightarrow ji \in g$. Let $g + ij$ denote a graph that results from adding tie ij to g holding all else fixed, while $g - ij$ denotes the removal of the ij tie from g. We say that there is a *path* between i and j in g if either $ij \in g$ or there exists m distinct players i_1, i_2, ..., i_m, such that $\{ii_1, i_1i_2, ..., i_mj\} \subset g$. Let $N_i(g) \subseteq N$ denote i's component, which is the set of all $j \in N$ with a (finite) path to i. Let $L_i(g) \subseteq N$ denote the set of i's ties, i.e., the set of all $j \in N$, $j \neq i$, such that $ij \in g$. A network is *connected* if $N_i(g) = N$, and we say that a network is *minimally connected* if there is exactly one single path between any two i and j. Being minimally connected implies no cycles. Figure 3.1

Figure 3.1: Network Examples

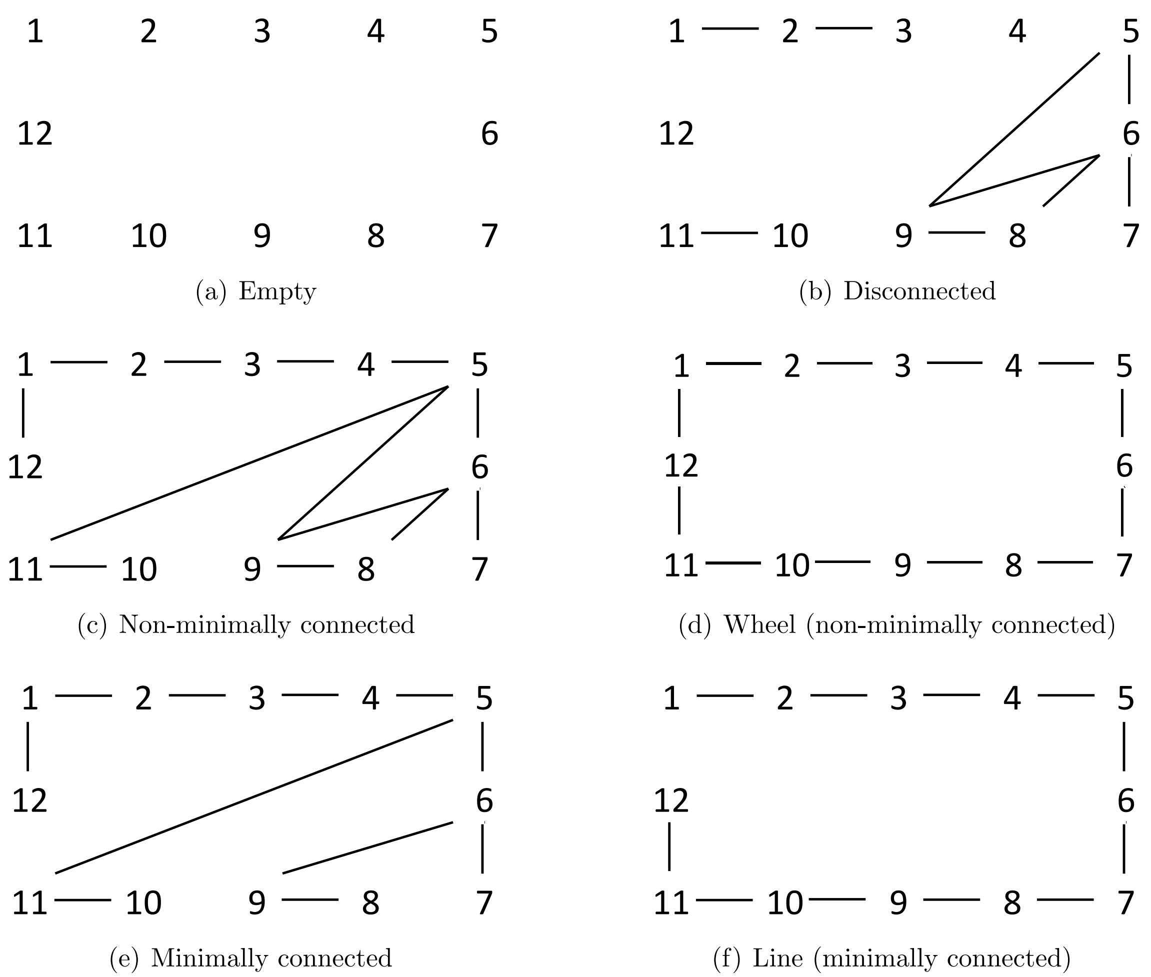

(a) Empty

(b) Disconnected

(c) Non-minimally connected

(d) Wheel (non-minimally connected)

(e) Minimally connected

(f) Line (minimally connected)

provides examples of networks and relevant network characteristics: (a) the empty network which has no ties; (b) a network that is not empty but also not connected; i.e., 4 and 12 are isolated; (c) a connected network that is not minimal, i.e., it has redundant ties that creates cycles of size three (569, 689) and four (5689); (d) a wheel network that is connected; (e) a minimally connected network, created by removing the 5-9 and 6-8 ties from (c); and (f) another minimally connected network, called the line network.

The simplified version of the Connections Model examined here assumes each actor has payoff function

$$u_i(g) = \sum_{j \in N_i(g)} 1 - \sum_{j \in L_i(g)} c.$$

The value of being connected, either directly or indirectly, to another actor is 1, and the cost of forming a direct tie is c.

Our experiment uses $c \in \{0.7, 1.5\}$ and $n = 12$. It is straightforward to show that at each cost level, an efficient network, defined here as a network that generates the maximum sum of utilities, must be minimally connected, and any minimally connected network is efficient. The intuition comes from recognizing the positive externality in ties. With low cost, adding ties up to minimal connectedness adds consecutively more individual and social benefits than cost, while adding ties beyond minimal connectedness produces additional cost but no benefits. With high cost, minimally connected networks still produce the highest total net benefits (though, if the cost were sufficiently large or if n were sufficiently small, then the only efficient network is the empty network in which $L_i(g) = \varnothing$ for all $i \in N$).

3.2.2 Full Observation

Multiple stability concepts have been used to identify the types of network structures that may emerge in the Connections Model. Here we mention two prominent ones, Pairwise Stability (PS) and Pairwise Farsighted Stability (PFS), as well as the dynamic implications of each. The first examination of the Connections Model in Jackson & Wolinsky (1996) used the Pairwise Stability (PS) concept to identify stable networks. A graph g is *Pairwise Stable* (PS) if (i) for all $ij \in g$, it is true that $u_i(g) \geq u_i(g - ij)$ and $u_j(g) \geq u_j(g - ij)$ and (ii) for any $ij \neq g$, it is true that if $u_i(g + ij) > u_i(g)$ then $u_j(g) > u_j(g + ij)$. The first condition requires both parties to prefer keeping a tie to severing a tie, and the second implies that if one party strictly prefers a new tie then the other party to that tie must strictly prefer against that tie. Notice that the PS concept implicitly assumes that, when deciding to form or sever ties, each actor fully observes the entire network g, i.e., each actor perfectly knows the utility of $g + ij$ and $g - ij$.

It is straightforward to show that, with our experiment's low cost $c = 0.7$, the set of PS networks is the set of minimally connected networks, while the empty network is the unique PS network with our high cost $c = 1.5$. With low cost, adding a tie to an actor not in your own component is individually rational, as is removing redundant ties. Only for minimally connected networks is it the case that no actor wants to add or remove ties. Under high cost, it is again true that removing redundant ties is individually rational, so any PS network must not have redundant ties. However, a minimal network that is not empty must necessarily have at least one stem, i.e., an actor has that exactly one tie (e.g., nodes 1 and 12 in Figure 3.1(f), and an actor directly connected to a stem is strictly better off removing the tie to that stem. The only network without redundant ties and stems is the empty network.

Watts (2001) and Jackson & Watts (2002) further examine the Connections Model in a dynamic framework, again implicitly assuming full observation. In these models, some initial network evolves according to the following *inertia* process: In each period, one (i, j)-pair is selected at random, and those actors decide myopically whether to keep the *status quo*, form a new tie, or sever an existing tie. As with the static PS concept, mutual consent is required to add a tie but ties can be removed unilaterally. Networks that this dynamic might enter but never leave are considered stable. Applying this dynamic in our setting yields a similar prediction as the PS concept: the system converges to a minimally connected network when the cost is low, and it converges to the empty network when the cost is high.

Other work considers network formation with far-sighted actors. One far-sighted concept used in the experimental literature on network formation is Pairwise Farsighted Stability (PFS) (Herings et al. 2009, Grandjean et al. 2011, Mantovani et al. 2013). This relies on the concept of a *far-sighted improving path*. As before, only one tie is changed at a time, but best responses are evaluated relative to the end network rather than the immediately adjacent network. A network g is PFS if all deviations from g to $g \pm ij$ are deterred by the threat of ending in some PFS network g' with no greater payoff for the deviator(s). The PFS concept refines the set of PS networks (Herings et al. 2009) although it is motivated by far-sighted rather than myopic actors. Here, the set of PFS networks is the same as the set of PS networks in both cost environments.[7] Far-sighted improving paths do exist between PS networks under the low cost (e.g., when an actor with multiple ties removes a tie and some other actor re-connects), but all such deviations could be credibly deterred by the reverse path, which would result in no improvement. The empty network is the unique PS, and thus PFS, network under high cost.

However, while PFS is a good descriptor of long-run group behavior in some settings (see e.g., Carrillo & Gaduh 2012, Mantovani et al. 2013), we note two additional but relevant complexities in our environment that may generate different far-sighted behavior than implied by the prior theory. First, our paired inertia and fixed time horizon hinder the ability of the group to deter deviations. It is improbable that inertia will select the correct pairs to reverse any given deviation in the short-run, and thus these paths are not credible deterrents. An actor may still choose to deviate in hope that it will lead to a nework with a greater individual payoff (we refer to this as "jockeying for position" in later sections).[8] Second, in

[7]Grandjean et al. (2011) and Mantovani et al. (2013) study environments in which some PS networks yield different aggregate payoffs. Thus certain network structures Pareto dominate the others. In our environment, all PS networks yield the same aggregate payoff.

[8]In a recent study, van Leeuwen et al. (2013) also find evidence of jockeying behavior in a setting where ties are formed unilaterally and investments in a public good determine the benefits.

our high cost environment, stability and efficiency are directly at odds. Far-sighted behavior could also lead the actors towards a recurrence class of (near) efficient networks and so far-sighted actors may still rationally achieve high levels of efficiency without ever converging to a PFS network.[9] Lastly, we note the main departure of our experiment from this previous theoretical work is that multiple ties can change at a time (see Section 3.3.3 for a detailed explanation). This aspect of our experimental design may introduce potential coordination problems, as situations could arise in which two distinct pairs collectively would like to add a single tie. However, this will only affect the path taken to the equilibrium networks rather than the set of equilibrium networks itself. Redundant links and missed connections that occur due to coordination failures will ultimately be remedied on the path to a PS network.

Bringing the theory and these issues together suggests the following predictions:

Claim 3.1. *Assume $n = 12$ and full observation.*

(a) With low cost $c = 0.7$, the process should converge to a minimally connected network.

(b) With high cost $c = 1.5$ and myopic actors, the process should converge to the empty network.

(c) With high cost $c = 1.5$ and far-sighted actors, the process may or may not converge. It will either converge to the empty network or it will transition between several high efficiency networks.

3.2.3 Neighbor Observation

There is less theoretical work on the Connections Model with neighbor observation. McBride (2006*b*) defines a Conjectural Pairwise Stable (CPS) network as a generalization of the PS

[9] Due to the finite length of our experiment, not all improving paths will yield a positive expected payoff, especially in the later rounds. Thus we might observe some convergence to (near) efficient networks even with far-sighted actors in the high cost setting.

concept to a limited observation setting. Each actor has a belief about the network, which takes the form of a probability distribution over possible networks. A network is CPS if each actor in that network believes she is better off keeping her existing ties and not forming any new ties and if her her beliefs are not contradicted by what she observes in the network. The key difference between CPS and PS is that in CPS an actor may have incorrect beliefs about the network as long as she has no information to contradict those beliefs. With neighbor observation, this means that, in equilibrium: the actor must have correct beliefs about her own ties and her neighbors' ties; with common knowledge of the payoff structure and knowledge of her own payoff, each actor's beliefs must assign probability 0 to any network in which i's component differs in size from the actual network; but the actor's beliefs about the existence and location of other ties in the network may be incorrect. Any PS network in which actors have correct beliefs is also CPS, but a few examples illustrate how some networks may be CPS but not PS. One particular case is to have a connected but not minimal network. Suppose the actual network is the wheel network in Figure 3.2(a), but suppose actor 6 believes with probability 1 that the actual network is the line in Figure 3.2(b). This belief is consistent with what actor 6 observes and with her payoff, and she does not want to remove or add ties given those beliefs. If every other actor similarly believes that she is in the middle spot in a line network, then this network is CPS. Notice that if an actor could see the cycle, then she would want to remove a tie. However, the cycle is outside of each actor's observational range, thereby allowing her beliefs to be incorrect about the network's exact structure. CPS networks can have smaller cycles, but the cycles must be of size 5 or more to be outside the observational range under neighbor observation.

Some disconnected networks may also be CPS networks. Figure 3.2(c) contains a disconnected network made of two 6-actor wheels. Redundant ties are again out of observational range. Yet, each actor knows from her payoff that her own component is of size 6. Whether or not an actor i wants to initiate a tie to some j not in her observational range depends on her belief about j's ties. If actor 6 believes that her network is that shown in Figure 3.2(d),

Figure 3.2: Network Examples: Neighbor Observation

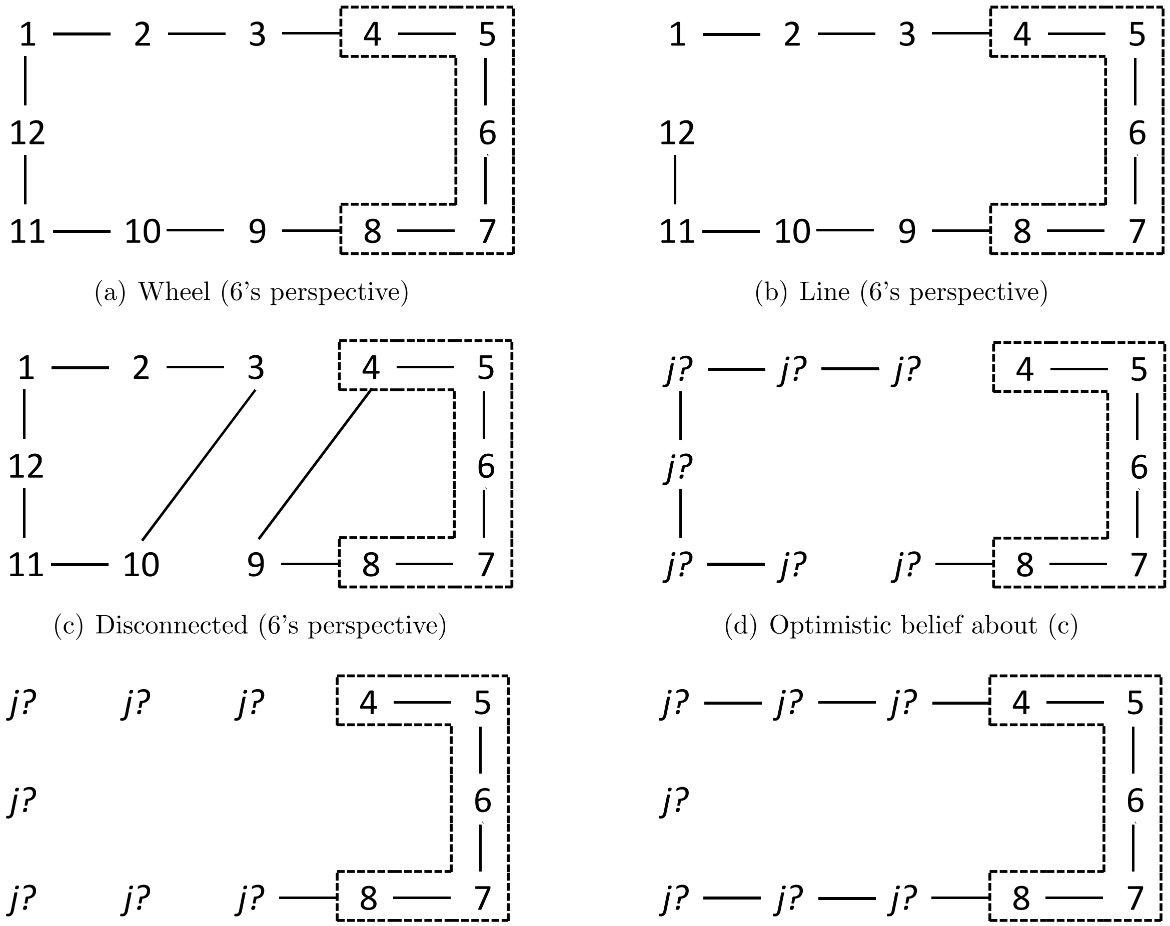

(a) Wheel (6's perspective)

(b) Line (6's perspective)

(c) Disconnected (6's perspective)

(d) Optimistic belief about (c)

(e) Pessimistic belief about (c)

(f) Uncertainty impedes connection

with $j \in \{1, 2, 3, 9, 10, 11, 12\}$ equally likely to be connected to actor 8, then her expected net benefit from forming a tie is $\frac{6}{7}(6) + \frac{1}{7}(0) - c = \frac{36}{7} - c$, which is strictly greater than 0 for $c \in \{0.7, 1.5\}$. Thus, if she optimistically believes, without contradicting her observation, that unobserved parts of the network are highly connected, then she will form the tie. If instead she believes that the network is that shown in Figure 3.2(e), again with each unobserved actor equally likely to be connected with 8 and consistent with her observation, then her expected net benefit from forming a tie is $\frac{6}{7}(1) + \frac{1}{7}(0) - c = \frac{6}{7} - c$. So, if she pessimistically believes that unobserved parts of the network are highly disconnected, then she will want to form a tie if the cost is low $c = 0.7$ but not if the cost is high $c = 1.5$. Even when the size of all components is known, as in Figure 3.2(f), uncertainty about the position of each actor in the network may inhibit connection. Actor 6's expected net benefit from linking to j with $j \in \{1, 2, 3, 9, 10, 11, 12\}$ equally likely to be isolated is $\frac{1}{7}(1) + \frac{6}{7}(0) - c = \frac{1}{7} - c$ regardless of her level of optimism or pessimism. She will not want to form this connection in either cost setting, even though she would want to form this connection with low cost $c = 0.7$ if she knew the isolated actor's identity with certainty.

Despite the potential for many networks to be CPS, CPS networks share several commonalities that limit the scope of what is possible. First, in both cost environments, cycles of size less than 5 are not CPS. Each actor in such a cycle can see the whole cycle and thus will want to remove a tie. Second, with low cost $c = 0.7$, the network must have a component of at least size 8 before even the most pessimistic of actors would not benefit from connecting to an unobserved actor. For example, when all other actors are isolated, the expected benefit from connecting to an unobserved actor for the center actor in a line of size 7 is $\frac{5}{7}(1) + \frac{2}{7}(0) > 0.7$ and for the center actor in a line of size 8 is $\frac{4}{7}(1) + \frac{3}{7}(0) < 0.7$. Third, with high cost $c = 1.5$, any network with a stem is not CPS because the actor tied to this stem would prefer to remove this tie. However, connected networks with large cycles and no stems such as the wheel are CPS with high cost. An additional restriction on beliefs, the *rationalizability criterion*, which requires that all actors in the conjectured network are best responding to this network, further limits the scope of possible CPS networks. For example, for high cost, the line network shown in Figure 3.2(b) is not rationalizable because actors 2 and 11 would benefit from removing the ties 1-2 and 11-12 respectively.

The CPS concept, like the PS concept, is static, and no theory to date has examined the connections model with limited observation in the dynamic framework we consider here.[10] In order to adapt the findings of Watts (2001) and Jackson & Watts (2002) to the case of

[10] Song & van der Schaar (2013) consider dynamic formation in a different limited observation setting. They assume that upon being paired in a period, i observes j's component (not the entire network) and vice versa before deciding whether to form or sever a tie. Our information setting does not allow actors to have such information. A larger difference, however, is that, similar to McBride (2008), they allow actors to be of different type. This incomplete information setting introduces additional complexities not present in our setting.

limited observation, we must specify the initial beliefs of all actors and the process by which these beliefs are updated as the network evolves. Here we consider two simple cases, at opposite extremes of the belief spectrum, to help pin down the set of possible CPS networks to which the system might converge. In both cases, we assume that beliefs satisfy *anonymity*, *stationarity*, and *memorylessness*, in addition to the typical CPS requirement that beliefs do not contradict the observed information.[11] Anonymity requires that all unobserved actors are assigned equal likelihood of being in each unobserved position in the conjectured network. Thus, decision makers only need to consider the structure of the conjectured network rather than all isomorphisms. Stationarity requires that beliefs be fixed unless the observed portion of the network changes. Thus decision makers do not attempt to forecast the dynamic path by which the network will evolve. Memorylessness requires an actor's belief about the existing network be formed independently of information observed in prior periods. Thus, once an actor leaves the decision maker's observational range, all information about this actor's position in the network is forgotten.

First consider the case of optimistic beliefs, e.g., Figure 3.2(d), in which actors assume that there are at most 2 disconnected components, both of known size. With low cost, actors will begin by connecting and continue connecting to unobserved actors until the network reaches at least size 10 ($\frac{3}{7}(3) + \frac{4}{7}(0) > 0.7$, but $\frac{2}{7}(2) + \frac{4}{7}(0) < 0.7$). Networks of size 10-12 are all possible depending on the order in which (i, j)-pairs are selected. Cycles of size 5 or greater are possible as (i, j)-pairs in the same component who cannot see each other may be randomly selected before the network reaches size 10+. With high cost, actors will begin by connecting and continue connecting to unobserved actors until the network reaches at least size 9 ($\frac{4}{7}(4) + \frac{4}{7}(0) > 1.5$, but $\frac{3}{7}(3) + \frac{4}{7}(0) < 1.5$). Networks of size 9-12 are possible and, as before, there may be cycles of size 5 or greater. However, whenever an actor is paired with a known stem she will remove the tie to that stem, and so the network would never converge

[11] These assumptions are admittedly quite strict and we do not expect them to hold in all cases. However, they do provide a simple and natural starting point for prediction. Furthermore, these assumptions are consistent with extreme forms of myopia.

in this case.

Now consider the case of pessimistic beliefs, e.g., Figure 3.2(e), in which actors assume all individuals outside of the component are isolated. With low cost, actors will begin by connecting and continue connecting to unobserved actors until the network reaches size 8 ($\frac{5}{7}(1) + \frac{2}{7}(0) > 0.7$, but $\frac{4}{7}(1) + \frac{3}{7}(0) < 0.7$). Networks of size 8-12 are possible and there still may be cycles of size 5 or greater. With high cost, actors will never connect, and the network will remain empty. If instead, actors are assumed to be far-sighted, they may take myopically detrimental actions as part of a far-sighted improving path. Such actions might include adding ties to unobserved individuals beyond what is myopically rational in an effort to achieve a maximal component (accidental redundant ties could be removed later) or adding ties with pessimistic beliefs and high cost as part of a path to a more efficient network (the dynamic would still not converge in this case).

As before, allowing multiple ties to change at a time introduces potential coordination problems. This may lead to visible cycles or missed connections, but we expect these known inefficiencies to be remedied by the actors at the next opportunity. Coordination failures may also lead to unobserved cycles that are never removed, but most of these cycles could occur even when one tie changes at a time.

We summarize the above logic with this prediction:

Claim 3.2. *Assume $n = 12$ and neighbor observation.*

(a) With low cost $c = 0.7$, the process should converge to a network with no more than two components (one at least size 8) and no visible cycles. All CPS networks can be reached.

(b) With high cost $c = 1.5$, the process may or may not converge, but if it does, it will converge to the empty network.

3.2.4 Jockeying for Position and Positional Risk

Individual connection decisions become more complicated when actors jockey for position. Consider, for example, the stem position in a connected network (e.g., actor 1 is a stem in Figure 3.2(b)). The stem position is the best possible position in a component for an actor because she receives the maximal benefit at the minimum cost. This is no longer necessarily the case when actors jockey for position in the short-run. Recall that jockeying for position involves removing a tie with the expectation that some other actor will re-connect. The stem position is particularly vulnerable to this behavior because of the low cost of removal. Thus, with random paired inertia, the stem position is less attractive to a risk averse actor due to the small chance that her partner will remove this tie in subsequent periods. We expect risk averse actors to be more likely to avoid this position, when possible.

3.3 Experiment Design

3.3.1 Basics

We conducted 8 experiment sessions at a large public university with a total of 240 subjects.[12] Our treatment variables are the cost of forming ties ($c = 0.7$ and $c = 1.5$) and the level of observation (neighbor or full). The treatment was chosen exogenously prior to the beginning of each session and remained fixed throughout the session. A total of 2 sessions were conducted for each of the 4 treatments in our between-subjects 2x2 factorial design. Subjects were students who had previously registered to be in the laboratory subject pool,

[12] In addition, we conducted a 2 group, 24 subject pilot prior to the first session to inform our choice of control variables. The two main conclusions from the pilot were that 20 periods per match was unnecessarily long (we reduced it to 15) and that assessing one relationship at a time was more manageable for the subjects than assessing 5 relationships at a time (we listened to this feedback from the questionnaire and restricted the decision to assessing one relationship at a time). We do not include the pilot data here due to the differences in the control variables and the absence of random rematching.

Table 3.1: Session Information

	# of Subjects	% Male	Major						Avg Earn ($)
			% Bus/ Econ	% Psy/ Cog Sci	% Other Soc Sci	% Eng/ Phys Sci	% Bio/ Life Sci	% Hum/ Undec	
All	240	42.1	13.3	7.9	15.0	27.9	25.8	10.0	33.09
NO – LC	60	36.7	10.0	10.0	6.7	33.3	28.3	11.7	37.00
NO – HC	72	38.9	13.9	6.9	18.1	18.1	33.3	9.7	31.16
FO – LC	60	48.3	16.7	5.0	21.7	25.0	21.7	10.0	32.41
FO – HC	48	45.8	12.5	10.4	12.5	39.6	16.7	8.3	31.94

Demographic information obtained from questionnaire. Avg earnings include $7 show-up payment.

and were recruited at random via E-mail. Subjects were allowed to participate in at most one session. Each session had either 24 or 36 subjects who were randomly matched into groups of 12. The sessions consisted of 10 matches[13] and subjects were randomly assigned to a new group at the beginning of each match. Each match consisted of 15 periods and the groups were held constant across periods. The experiment was programmed and conducted using the z-Tree software package (Fischbacher 2007). Screenshots of the experiment can be seen in Figure 3.3. At the end of the session, each subject's risk preference was elicited using the Eckel & Grossman (2008) method and the subjects were asked to complete a short questionnaire. Subjects received $1.00 for every 50 units of experimental currency earned during the session and were given a lump sum payment ($2.00 for low cost and $4.50 for high cost) at the beginning of the session to prevent bankruptcy. Additionally, subjects received a $7.00 show-up payment and could earn between $0.00 and $14.00 in the risk preference elicitation, depending on that subject's choice of lottery and the outcome of the lottery. In total, subjects took home an average of approximately $33.00 for 2 hours of participation (Table 3.1).

[13]One session of full observation with low cost was stopped a match early due to the cumulative effects of a slight, unanticipated, software lag at the end of each round that prevented all matches from being completed in the allotted time. We reduced the number of groups per full observation treatment to reduce the lag in future sessions. This adjustment fixed the problem.

Figure 3.3: Software Screenshots

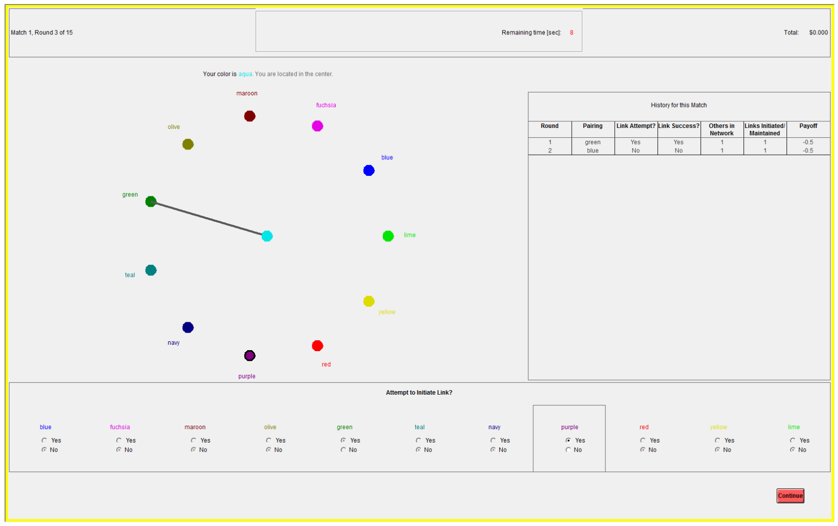

(a) Decision Screen

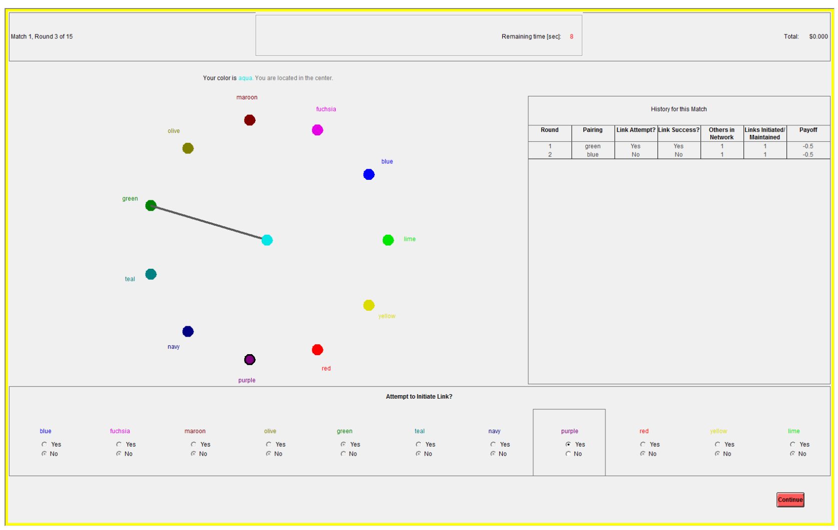

(b) Results Screen

3.3.2 A Single Match

At the beginning of a match, each participant is randomly assigned a group and a color and is shown an image of the initial (empty) network. The image (see Figure 3.3) contains 12 colored circles or "nodes;" one node corresponding to each participant in the group. The color represent's the subject's experimental identity and is used instead of assigning each subject an ID number. A participant's own node is located at the center of her image surrounded by a ring consisting of the other 11 nodes. The order in which the nodes are patterned remains fixed throughout the match, but this order is determined randomly for each participant at the start of the match to avoid any position specific framing effect. Ties within a participant's observational range are depicted by a dark grey line connecting two nodes, whereas ties outside of a participant's observational range are not depicted. A match consists of 15 rounds and the image is updated at the end of every round to display all visible changes in the network. Participants are paid for every round, rather than the end network, to ensure that all decisions are incentivized.[14] The participants are shown the round number, the match number, their most recent earnings, their cumulative earnings, and a match specific history at all times throughout the match.

3.3.3 A Single Round

At the beginning of a round, each participant is randomly paired with one other participant in her group. These pairings are restricted to be non-overlapping and so there are 6 distinct pairs in each round.[15] Participants are first shown the Decision Screen (Figure 3.3 (a)) and have 15 seconds to decide whether or not to initiate (or maintain) a tie to their partner. The participant's partner is highlighted on the network image to aid in the decision-making process. A participant indicates her choice by selecting yes or no in the decision prompt and then pressing a submit button. The default option, which is initially selected and which is submitted on this participant's behalf if time runs out, is set so as to maintain the status quo. In the vast majority of cases (98%), the submit button was pressed before time ran out and so this design element did not substantively affect our results. After all decisions are

[14] These incentives are particularly relevant for decisions that involve jockeying for position and/or decisions that will change the visible portion of the network.

[15] The main difference between our inertia process and the inertia process of Jackson & Watts (2002) is that we allow six pairs to evaluate their relationship at a time rather than one pair at a time. Allowing six pairs to evaluate their relationship at a time substantially increases the speed of convergence, allowing us to collect more data in the allotted time. This decision comes at the cost of introducing potential coordination problems. However, this will only affect the path taken to the equilibrium networks rather than the set of equilibrium networks itself.

submitted (or after time runs out), all participants are then shown the result screen (Figure 3.3 (b)), in which they learn the outcome of their decision, their payoff for the round, and are shown the visible portion of the resultant network. Participant i's payoff for network g is given by

$$u_i(g) = 1 \times (\# \text{ of direct ties}) - c \times (\# \text{ of direct ties} + \# \text{ of unsuccessful tie attempts}) \, .$$

Thus a participant pays the cost of any tie attempt she makes, whether or not it is successful, as well as the cost of any direct tie she formed in a prior round, whether or not she was able to change that tie in this round. This payment scheme matches the one used in the Connections Model (Jackson & Wolinsky 1996; no decay of benefits) and in McBride (2006*b*).

3.3.4 Treatment Variables

One first treatment variable is the level of observation. The participants either experience *full observation* or *neighbor observation* depending on the session. Under full observation, each participant can see the ties of every individual in her group, whether she is connected to that individual or not. This is the level of observation that has been used in prior experimental analyses, albeit with different control variables (e.g., Falk & Kosfeld 2012, Callander & Plott 2005, Berninghaus, Ehrhart & Ott 2006, Goeree et al. 2009, Pantz 2006, Mantovani et al. 2013, Carrillo & Gaduh 2012). Under neighbor observation (called 2-link observation in McBride 2006*a*), each participant sees only her direct ties and the ties of her neighbors. Ties beyond this distance are not displayed. In addition, a participant will not see the ties of any other participant that is not in her component. Individuals that are far away in one's own component are indistinguishable from individuals in a separate component and no information is provided as to the size of this component. Participants do, however, learn their own payoff, and thus can infer how many individuals are in their own component.

Our other treatment variable is the cost of forming ties. The participants either experience *low cost* or *high cost* depending on the session. Under low cost, an individual's cost of initiating (or maintaining) a direct tie is $c = 0.7$. Because the benefit of adding a person exceeds the cost of adding a person in this treatment, the myopic best response is to add a tie to any individual who is known not to be in the same component. Under high cost, an individual's cost of initiating (or maintaining) a direct tie is $c = 1.5$. Because the benefit of adding a person is less than the cost of adding a person in this treatment, the myopic best response to never initiate (or maintain) ties to isolated individuals (or stems).

3.4 Hypotheses

Given these parameters, we formulate the following testable hypotheses based upon the predictions from Section 3.2.

3.4.1 Convergence

1. Holding the observation level fixed, the frequency of convergence is higher under low cost than under high cost.

2. Holding the cost fixed, the frequency of convergence is higher under neighbor observation than under full observation.

3.4.2 Efficiency

3. Holding the observation level fixed, the average efficiency is higher under low cost than under high cost.

4. Holding the cost fixed, the frequency of cycles is higher under neighbor observation than under full observation.

5. Holding the cost fixed, stems are dropped more frequently under full observation than under neighbor observation.

6. Holding the observation level fixed, stems are dropped more frequently under high cost than under low cost.

7. A higher degree of risk aversion is associated with a lower likelihood of being in a stem position.

3.5 Results

Our unit of observation for hypothesis tests is a single match. When the relevant data is generated by period, we compute match level averages. Because we match the same 24-36 participants repeatedly into random groups of 12 and provide regular feedback, our data could feasibly depend on the participants, group, and match. For all participant level hypotheses, we correct the standard-errors for two-way clustering at the participant and group level (Cameron, Gelbach & Miller 2011, Thompson 2011; compared to other approaches in Peterson 2009) and include indicator variables to control for the match. For all group level hypotheses, we assume in the main discussion that the observations do not depend on the participants or match. We then relax this assumption in Table B.3 and find that the p-values are practically identical, if not smaller. This is not surprising in our setting, because the large group size, random paired inertia, and required mutual consent to form ties limit the idiosyncratic effect of individual participants on group outcomes. Furthermore, while we do find significant trends in some group outcomes across matches (see Tables B.1 and B.2), these trends do not generally depend on the treatment. In the one case where the trend does depend on the treatment, our standard errors are biased upward, if anything.

For each match, we record whether the network converged, the efficiency of the convergent networks, and the end of match efficiency of the non-convergent networks. As was first noted by Callander & Plott (2005), measuring convergence in network formation experiments involves "trading off empirical certainty with experimental constraints" such as boredom and end of match effects (p. 1478). A stricter measure has the advantage of reducing the occurrence of false positives, where a non-convergent network is categorized as convergent, but it comes at the cost of increasing the likelihood of false negatives, where a convergent network is categorized as non-convergent due to non-strategic changes that may occur after a time of stationarity. For this reason, we consider a network to have converged if it remains unchanged for 5 consecutive rounds. This measure is relatively strict as 5 rounds represents 1/3 of the total match and requires between 30 and 60 independent decisions to not change the network. We compute the efficiency of a network by dividing the aggregate group payoff by the maximum possible group payoff (116.6 with low cost and 99 with high cost). This represents a measure of both absolute and relative efficiency as the aggregate group payoff of the empty network is 0. To allow for closer comparison to the convergent networks, we consider the efficiency of a non-convergent network to be its average efficiency over the last 5 rounds of the match.

At the end of the session, we measure the participants' risk preferences using the Eckel & Grossman (2008) method. A participant is given the choice between 5 lotteries and this choice maps to a range of possible risk coefficients under the assumption of constant relative risk aversion (the range of possible ρ is reported in parenthesis).[16] The lotteries are: \$4 with probability 100% ($2.63 < \rho$), \$3 with probability 50% or \$7 with probability 50% ($0.90 < \rho < 2.63$), \$2 with probability 50% or \$10 with probability 50% ($0.53 < \rho < 0.90$), \$1 with

[16] The CRRA utility function is given by $U(m) = m^{(1-\rho)}/(1-\rho)$, where m is money and ρ is the risk aversion parameter. $\rho = 0$ corresponds to risk neutral, $\rho > 0$ corresponds to risk averse (greater values indicate more severe risk aversion), and $\rho < 0$ corresponds to risk seeking (lesser values indicate more severe risk seeking).

probability 50% or \$13 with probability 50% ($0.00 < \rho < 0.53$), and \$0 with probability 50% or \$14 with probability 50% ($\rho < 0.00$). Here we list these lotteries from most risk averse to least risk averse, but the order of the lotteries was randomly permuted for each participant to avoid any order effect.

We present our results, ordered by hypothesis, below. These results are summarized in Table B.3.

3.5.1 Convergence

In total, 83 out of 197 networks converged. Our convergence results, broken down by treatment, are shown in Table 3.2. Figure 3.4 shows the median number of network changes by

Table 3.2: Frequency of Convergence by Treatment

	Neighbor Observation	Full Observation
Low Cost	50.0% (25/50)	42.6% (20/47)
High Cost	23.3% (14/60)	60.0% (24/40)

period for each treatment. The typical network converges in period 7 for Neighbor Observation - Low Cost, in period 10 for Full Observation - Low Cost, and in period 5 for Full Observation - High Cost. The typical network for Neighbor Observation - Full Cost does not converge. To test for a match level trend in the frequency of convergence, we estimate a pooled logistic regression of the form

$$\Pr(\text{converge}_i = 1) = \text{logit}^{-1}(\mathbf{T}_i\vec{\alpha} + \beta\text{match}_i + \text{match}_i\mathbf{T}_i\vec{\gamma}),$$

where $\mathbf{T}_i$ is a vector of treatment dummy variables and match_i is a linear trend. We then compute the average marginal effect for each treatment. On average there was a 5.5% increase in the frequency of convergence between matches in the Neighbor Observation - Low Cost treatment, a 3.2% increase in the Neighbor Observation - High Cost treatment,

Figure 3.4: Median Number of Network Changes by Period

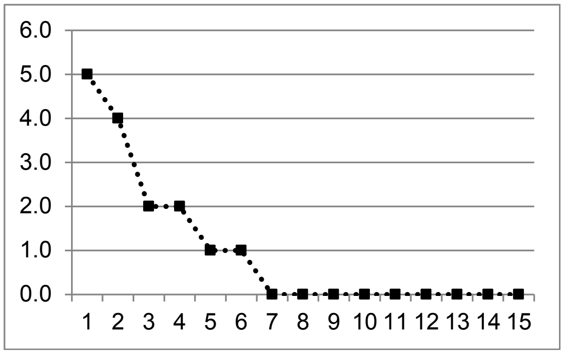

(a) Neighbor Observation - Low Cost

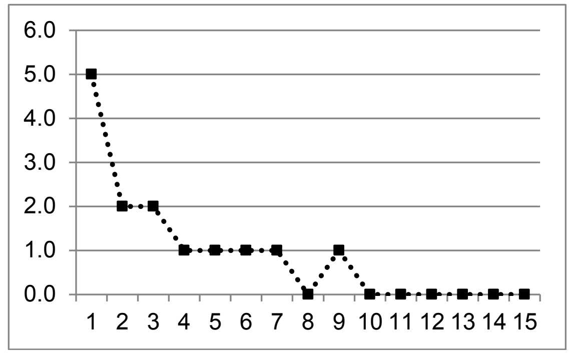

(b) Full Observation - Low Cost

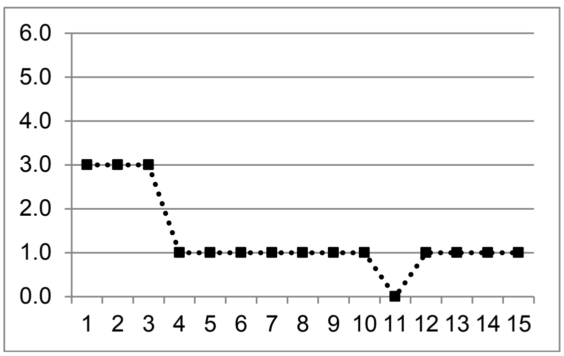

(c) Neighbor Observation - High Cost

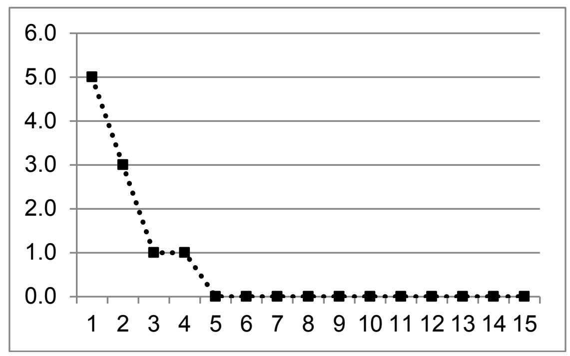

(d) Full Observation - High Cost

a 2.8% increase in the Full Observation - Low Cost treatment, and a 5.0% increase in the Full Observation - High Cost treatment. We cannot reject the null hypothesis that these trends are jointly equal (p-value 0.760).Overall, there was a 4.0% increase in the frequency of convergence between matches, significant at the 0.01 level.

Hypothesis 1: Holding the observation level fixed, the frequency of convergence is higher under low cost than under high cost.

We test the null hypothesis that the frequency of convergence under low cost is less than or equal to the frequency of convergence under high cost using Barnard's Exact Test. We reject this hypothesis at the 0.01 level for neighbor observation but cannot reject this hypothesis for full observation (p-value 0.938). For neighbor observation, networks were significantly more likely to converge with low cost (50.0%) than with high cost (23.3%) but for full observation there was no statistically significant difference in convergence. We interpret this as weakly confirming Hypothesis 1. The existence of (near) efficient (C)PS networks under low cost does seem to increase the likelihood of convergence, but this is clearly not capturing the full picture. We do not find the expected effect for full observation and, if anything, the evidence seems to point towards *less* convergence with low cost than high cost (42.6% vs. 60.0%).

Hypothesis 2: Holding the cost fixed, the frequency of convergence is higher under neighbor observation than under full observation.

We test the null hypothesis that the frequency of convergence under neighbor observation is less than or equal to the frequency of convergence under full observation using Barnard's Exact Test. We cannot reject this hypothesis under low cost (p-value 0.279) or high cost (p-value $>$ 0.999). For high cost, networks were significantly less likely to converge with neighbor observation (23.3%) than with full observation (60.0%) and for low cost there was no statistically significant difference in convergence (50.0% vs. 42.6%). We interpret this as rejecting Hypothesis 2. The existence of CPS networks that are not PS networks is not sufficient to ensure a higher frequency of convergence under neighbor observation.

3.5.2 Efficiency

We record the final efficiency of each network pairing. To allow for direct comparisons between convergent networks (C) and non-convergent networks (NC), we consider the final efficiency of a non-convergent network to be its average efficiency over the last 5 periods. Our efficiency results, broken down by treatment, are shown in Table 3.3 (sample size in parentheses). Our primary hypotheses do not explicitly differentiate between convergent and non-convergent networks, but it is clear from Table 3.3 that convergent networks are more efficient. We reject the null hypothesis of equal medians (Mann-Whitney U Test) at

Table 3.3: Median Efficiency by Treatment

	Neighbor Observation	Full Observation
Low Cost	93.2% (50)	89.5% (38*)
C | NC	97.6% (25) | 88.4% (25)	99.4% (16) | 84.4% (22)
High Cost	80.4% (60)	75.8% (40)
C | NC	80.8% (14) | 78.0% (46)	88.9% (24) | 55.3% (16)

*Excludes 9 networks containing a subject who never connected.

the 0.05 level for Neighbor Observation - High Cost and at the 0.01 level for the other three treatments.

When we instead plot the median efficiency by period, as shown in figure 3.5, we see that efficiency for the typical network increases over time and then eventually plateaus, with a slight efficiency dropoff at the end for the high cost treatments. To test for a match level

Figure 3.5: Median Efficiency by Period

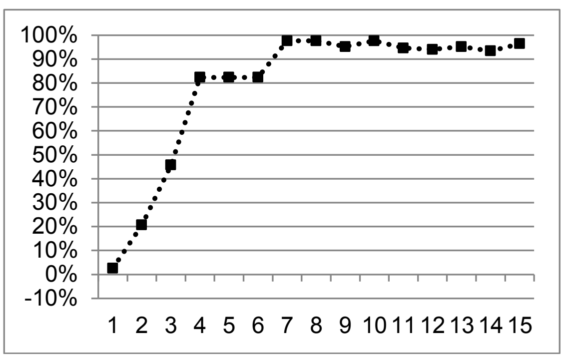

(a) Neighbor Observation - Low Cost

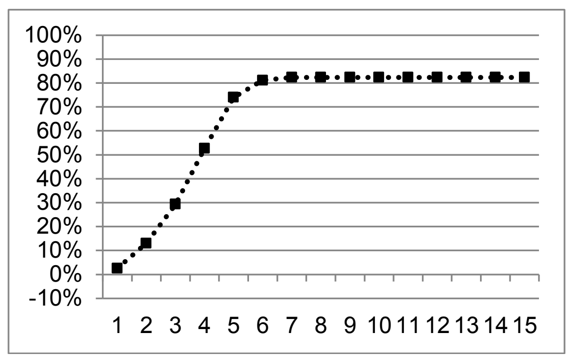

(b) Full Observation - Low Cost

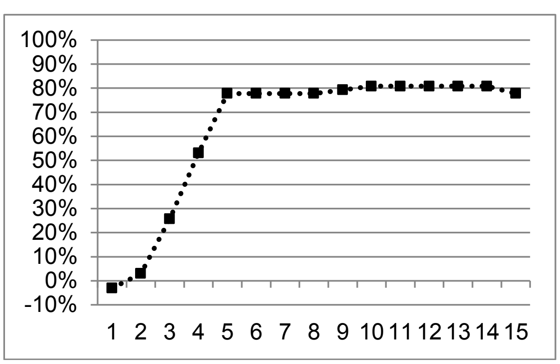

(c) Neighbor Observation - High Cost

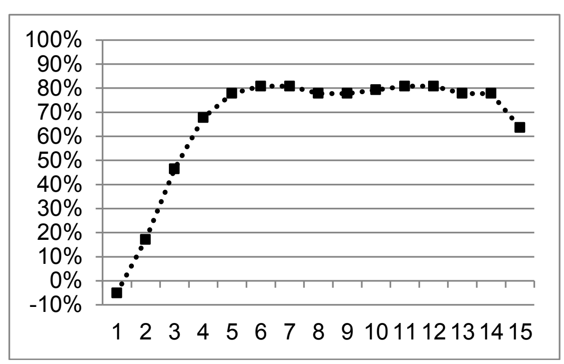

(d) Full Observation - High Cost

trend in the final efficiency of the network, we estimate a pooled regression of the form

$$\text{efficiency}_i = \mathbf{T}_i\vec{\alpha} + \beta\text{match}_i + \text{match}_i\mathbf{T}_i\vec{\gamma},$$

where $\mathbf{T}_i$ is a vector of treatment dummy variables and match_i is a linear trend. We then compute the average marginal effect for each treatment. On average there was a 0.6% decrease in efficiency between matches in the Neighbor Observation - Low Cost treatment, a 0.7% decrease in the Neighbor Observation - High Cost treatment, a 0.9% increase in the Full Observation - Low Cost treatment, and a 3.7% increase in the Full Observation - High Cost treatment. We reject the null hypothesis that these trends are jointly equal at the 0.01 level, the trend for Full Observation - High Cost is significantly different from the other treatments and zero. The overall trend for the remaining three treatments (0.2% decrease) is not significantly different from zero (p-value 0.537).

Hypothesis 3: Holding the observation level fixed, the average efficiency is higher under low cost than under high cost.

We test the null hypothesis that the efficiency under low cost is less than or equal to the efficiency under high cost using the Mann-Whitney U Test. We reject this hypothesis at the 0.01 level for both neighbor observation and full observation.[17] For both observation levels, median efficiency is greater with low cost than with high cost (93.2% vs. 80.4% for neighbor observation and 89.5% vs 75.8% for full observation). We interpret this as confirming Hypothesis 3. The stronger incentives to maintain connectivity with low cost do seem to yield higher average efficiencies.

In total, 33/83 convergent networks and 69/114 non-convergent networks had at least one cycle. We further break this down by treatment in Table 3.4. To test for a match level trend

[17] Controlling for the different match level trends under full observation and conducting a Wald test does not change the significance level. The p-value drops from 0.008 to < 0.001.

Table 3.4: Frequency of Cycles by Treatment

	Neighbor Observation	Full Observation
Low Cost	72.0% (50)	40.4% (47)
C | NC	60.0% (25) | 84.0% (25)	30.0% (20) | 48.1% (27)
High Cost	61.7% (60)	25.0% (40)
C | NC	57.1% (14) | 63.0% (46)	16.7% (24) | 37.5% (16)

in the frequency of cycles, we estimate a pooled logistic regression of the form

$$\Pr(\text{cycle}_i = 1) = \text{logit}^{-1}(\mathbf{T}_i\vec{\alpha} + \beta\text{match}_i + \text{match}_i\mathbf{T}_i\vec{\gamma}),$$

where $\mathbf{T}_i$ is a vector of treatment dummy variables and match_i is a linear trend. We then compute the average marginal effect for each treatment. On average there was a 5.6% decrease in the frequency of cycles between matches in the Neighbor Observation - Low Cost treatment, a 8.2% decrease in the Neighbor Observation - High Cost treatment, a 8.6% decrease in the Full Observation - Low Cost treatment, and a 4.2% decrease in the Full Observation - High Cost treatment. We cannot reject the null hypothesis that these trends are jointly equal (p-value 0.189).Overall, there was a 6.9% decrease in the frequency of convergence between matches, significant at the 0.01 level.

Hypothesis 4: Holding the cost fixed, the frequency of cycles is higher under neighbor observation than under full observation.

We test the null hypothesis that the frequency of cycles under neighbor observation is less than or equal to the frequency of cycles under full observation using Barnard's Exact Test. We reject this hypothesis at the 0.01 level for both low cost and high cost. For both cost environments, cycles were significantly more likely to occur in convergent networks with neighbor observation than with full observation (60.0% vs. 30.0% for low cost and 57.1% vs 16.7% for high cost). We interpret this as confirming Hypothesis 4. Nearly 60.0% of convergent networks had cycles with neighbor observation, but only Only 15.4% (6/39) of these networks had *visible* cycles. Thus the lack of full observation *does* increase the likelihood with which inefficient cycles form.

Jockeying for position may result in periodic disconnection; particularly when a stem is paired with its neighbor. For each individual and match, we record the number of times the individual was paired with her own stem and the number of times the individual elected to remove the tie. From this we compute this individual's frequency of stem removal for the match. The overall frequency of stem removal by treatment is shown in Table 3.5.

Table 3.5: Frequency of Stem Removal by Treatment

	Neighbor Observation	Full Observation
Low Cost	27.3% (65/238)	23.9% (62/259)
High Cost	47.5% (135/284)	34.6% (63/182)

Hypothesis 5: Holding the cost fixed, stems are dropped more frequently under full observation than under neighbor observation.

To test the null hypothesis that the frequency of stem removal under neighbor observation is greater than or equal to the frequency of stem removal under full observation, we we estimate a regression of the form

$$\text{droppedstems}_i = \mathbf{M}_i\vec{\alpha} + \beta\text{neighbor}_i$$

for each cost setting, where $\mathbf{M}_i$ is a vector of match dummy variables and neighbor_i equals 1 for neighbor observation, 0 otherwise. We then conduct a t-test of whether β is greater than or equal to 0, correcting for two-way clustering at the group and participant levels. We cannot reject this hypothesis under low cost (p-value 0.666) or high cost (p-value 0.914). For high cost, stems were significantly more likely to be dropped with neighbor observation (47.5%) than with full observation (34.6%) but for low cost there was no statistically significant difference in stem removal (27.3% vs 23.9%). We interpret this as rejecting Hypothesis 5. Stem actors are more likely to be dropped under neighbor observation even though it may be harder for these actors to reconnect.

In an effort to better understand this rejection, we conducted a follow-up analysis. Our hypothesis was premised on the notion that reconnection would be more difficult under neighbor observation than full observation due to uncertainty about one's own connectivity to an unobserved actor; however, an unobserved actor may also be viewed *more* favorably than an observed isolate when actors have optimistic beliefs (i.e., that the actors outside one's component are connected, not isolated). This effect is especially large in early periods, when components are much smaller than the maximum of 12, and with high cost, where connecting to an isolated actor is myopically detrimental.

We do not elicit the beliefs of our subjects, and so we cannot definitively show that optimistic beliefs are the cause of this rejection. However, there are several checks that we can conduct to assess the plausibility of this explanation. First, we note that the networks we observe under neighbor sight are more consistent with optimistic beliefs than pessimistic beliefs. Only 1/40 convergent networks have less than 10 actors in the maximal component. Further, only 2/71 non-convergent networks had less than 10 actors in the maximal component throughout the last 5 periods. If actors held pessimistic beliefs, we would expect more instances where the maximal component had 8 or 9 actors. Second, we note that under neighbor observation, a much larger percentage of stem removal occurs in the early rounds, before network efficiency plateaus (44.6% vs 27.4% for low cost, and 31.1% vs 14.3% for high cost). When we compare the distribution of stem removals by period, we reject the null hypothesis of equal medians at the 0.05 level for high cost. We cannot reject this hypothesis for low cost (p-value 0.24). This evidence points (weakly) towards a higher percentage of jockeying in the early rounds for neighbor observation than full observation, which is consistent with the incentives of far-sighted actors when beliefs are generally optimistic.

Hypothesis 6: Holding the observation level fixed, stems are dropped more frequently under high cost than under low cost.

To test the null hypothesis that the frequency of stem removal under low cost is greater than or equal to the frequency of stem removal under high cost, we we estimate a regression of the form

$$\text{droppedstems}_i = \mathbf{M}_i\vec{\alpha} + \beta\text{lowcost}_i$$

for each observation level, where $\mathbf{M}_i$ is a vector of match dummy variables and lowcost_i equals 1 for low cost, 0 otherwise. We then conduct a t-test of whether β is greater than or equal to 0, correcting for two-way clustering at the group and participant levels. We reject this hypothesis at the 0.01 level under neighbor observation and at the 0.05 level under full observation. For both observation levels, stems are more likely to be dropped with high cost than with low cost (47.5% vs. 27.3% for neighbor observation and 34.6% vs 23.9% for full observation). We interpret this as confirming Hypothesis 6. Actors do tend to remove ties whose cost exceeds the immediate benefit.

An individual's chosen position in the network may also depend on risk aversion. We measure each subject's risk aversion using the Eckel & Grossman (2008) method and separate these subjects into 5 categories based on the lotteries chosen. For each individual and match, we record the number of rounds the individual spent in the stem position. From this we compute this individual's fraction of time spent in the stem position for the match. For the purpose of this analysis, we exclude the first 5 periods, because networks do not typically begin to stabilize until round 6. The overall fraction of time spent in the stem position by risk category is shown in Table 3.6.

Table 3.6: Fraction of Time Spent in Stem Position by Risk Category

$\rho < 0.00$	33.8%
$0.00 < \rho < 0.53$	33.0%
$0.53 < \rho < 0.90$	34.5%
$0.90 < \rho < 2.63$	29.0%
$2.63 < \rho$	30.2%

Hypothesis 7: A higher degree of risk aversion is associated with a lower likelihood of being in a stem position.

To test the null hypothesis of equal proportions across risk categories against the alternative of different proportions, we first estimate a pooled regression of the form

$$\text{droppedstems}_i = \mathbf{M}_i\vec{\alpha} + \mathbf{R}_i\vec{\beta},$$

where $\mathbf{M}_i$ is a vector of match dummy variables and $\mathbf{R}_i$ is a vector of risk category dummy variables (one category omitted). We then conduct an F-test of whether the coefficients in $\vec{\beta}$ are jointly equal to 0, correcting for two-way clustering at the group and participant levels. We cannot reject this hypothesis (p-value 0.456). While time spent in the stem position does follow the predicted pattern across risk categories, the difference across categories is neither statistically or economically significant. The evidence does not generally support Hypothesis 7. The risk of being dropped from the network does not seem to be a major factor in each actor's decision about where to locate in the network.

3.5.3 Typical Networks

Figure 3.6: Typical Network by Treatment

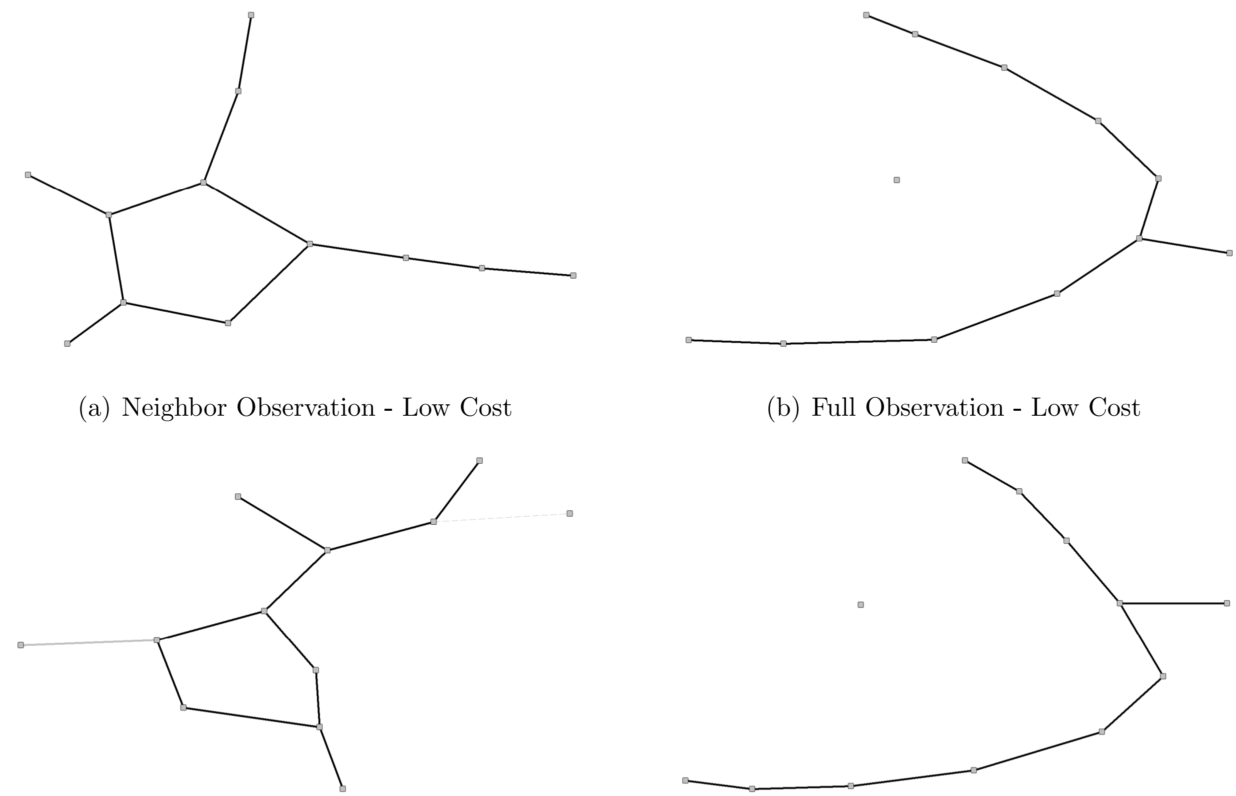

(a) Neighbor Observation - Low Cost

(b) Full Observation - Low Cost

(c) Neighbor Observation - High Cost

(d) Full Observation - High Cost

Realized networks clearly differ in many key characteristics, yet it is instructive to identify what we can be considered typical for each treatment condition. Figures 3.4 and 3.5 characterize the median behavior of the network formation process by treatment. We can infer from these analyses when the median network converges (if at all) and at what level of efficiency. These are just two of many dimensions on which networks can be compared; other dimensions include: number of redundant ties, number of cycles, size of cycles, and size of maximal component. To better understand the the properties of the typical networks in each treatment, we compute the median values for each of these dimensions and then find the one realized network from each treatment that most closely matches the median across dimensions. Visuals of these networks are shown in Figure 3.6. Each network converged within 1 period of the median for that treatment, has an efficiency within a few percentage points of the median, and exactly matches the median in all other dimensions. In these visuals, created using the software package UCINET,[18] actors are denoted by grey squares and ties are denoted by solid black lines. When the network did not converge, we present a visual of that network over the last 5 periods. Ties that were added after period 11 are colored grey, and ties that were removed before period 15 are denoted by grey dashed lines.

The typical network for Neighbor Observation - Low Cost converges to a connected network with one hidden cycle/redundant tie (Figure 3.6(a)), the typical network for Neighbor Observation - High Cost does not converge and has a component of size 11 (one isolate) with one cycle/redundant tie (Figure 3.6(c)), and the typical networks for both Full Observation treatments converge to a component of size 11 (one isolate) with no cycles/redundant ties (Figures 3.6(b) and 3.6(d)).

3.6 Conclusion

We present the first experimental study of network formation when actors have limited observation of the network. Our 2x2 experiment design exogenously varies two treatment variables: the cost of forming links and the level of network observation. Observed behavior is largely, but not always, consistent with our predictions. Convergence is more likely with low link cost than high link cost, but it is also less likely with neighbor observation than with full observation. Efficiency is higher under low cost as expected, with more cycles observed under neighbor observation, and more stems dropped with full observation. Stems are removed more often under neighbor observation and high cost, and we do not find a significant relationship between risk aversion and the stem position.

[18]Borgatti, S.P., Everett, M.G. and Freeman, L.C. 2002. Ucinet for Windows: Software for Social Network Analysis. Harvard, MA: Analytic Technologies.

We also find strong evidence of far-sighted behavior, a double-edged feature in this setting. Far-sighted behavior leads to the rapid formation of connected or nearly connected networks, generates large utility gains. Moreover, the network often converges to a network that is not Pairwise Farsightedly Stable, again to the benefit of actors. However, far-sighted behavior also leads to jockeying for advantageous positions in the networks, a behavior which leads to short-run utility losses as links are severed and reformed. The frequency and characteristic of jockeying appears to be different under neighbor observation than full observation. While there is a higher frequency of inefficient, redundant links (cycles) under neighbor observation, there is also a lower frequency of inefficient jockeying, and this leads, surprisingly, to higher average utilities under neighbor observation. When jockeying occurs under neighbor observation, it appears relatively more often in earlier rounds, unlike under full observation, in which case it is later rounds. This may be evidence of optimistic beliefs.

Future work has many avenues to explore. The incentives to form and sever links vary widely across network settings, and additional work is needed to determine how many of our findings generalize. That the location of a subject in the network may depend on her risk preferences also deserves closer examination, especially in network settings where network position is more consequential for welfare than in our setting. Finally, future theoretical and experimental work can identify just how much network observation is enough to mimic full information, and can even examine the actors' incentives to collect information, perhaps at a cost, about others links. We expect this work to provide even more insights into the role limited observation plays in network formation.

Chapter 4

The Origin of the State as a Stationary Bandit

In this chapter, I develop a model of the origin of the state. Starting from an anarchic environment, where property is insecure and must be defended, I show how a primitive state can emerge out of anarchy. The coalition in power is determined endogenously, and the emergent order can resemble autocracy, oligarchy or meritocracy, depending on the initial distribution of violence capacities. I find that aggregate production increases under a cooperative pact, and that when the coalition in power is predatory, producers are often, but not always, better off in the emergent order than under anarchy. The existence of a powerful coalition that is able to secure payoffs beyond its productive capacity cannot be ruled out, even when individuals are *ex ante* symmetric.

4.1 Introduction

> In a world of roving banditry there is little or no incentive for anyone to produce or accumulate anything that may be stolen and, thus, little for bandits to steal. Bandit rationality, accordingly, induces the bandit leader to seize a given domain, to make himself the ruler of that domain, and to provide a peaceful order and other public goods for its inhabitants, thereby obtaining more in tax theft than he could have obtained from migratory plunder. – Olson (1993)

How did the state emerge out of anarchy? This age-old question has puzzled philosophers, political scientists, and economists alike for centuries. Olson's (1993, 2000) seminal theory that the state originated from a *stationary bandit*, who tempers her theft in the short-run to preserve long-run production incentives, provides an intriguing starting point to study the origin of the state. In particular, the key problem that society must overcome to achieve order in Olson's framework has direct analogs to other prominent theories such as Nozick (1977) and North, Wallis & Weingast (2009). In all of these theories, order is ultimately seen a solution to the problem of endemic conflict. Where these theories differ is in the precise definition of conflict, whether that be theft (Olson 1993), violence (North et al. 2009) or infringement on natural rights (Nozick 1977), and in the specifics of how the problem is solved. Olson and North et al. posit that it is powerful individuals (i.e., the would be conflict initiators) who organize to solve the conflict, whereas Nozick posits that it is the rest of society (i.e., the would be victims) who organize.

In this chapter, I show that, starting with the primitive of a distribution of violence capacities and a simple production economy, many of the natural states that have emerged throughout recorded history (see e.g., Fukuyama 2011, for an overview) can be explained with a framework such as Olson's or North et al.'s and cannot be explained with a framework such as Nozick's.[1] One possible explanation for this discrepancy is that natural law theory (e.g., Grotius 1625, Pufendorf 1672, Hume 1739, Locke 1689) relies on consent (tacit or otherwise) to justify the initial social contract, which, although not typically addressed, is consistent with the idea that concessions might need to be made to powerful individuals when entering into the social contract.[2] Thus, a strict reliance on Locke's "workmanship model" of property as is done by Nozick can lead to conclusions that are not generalizable. If the distribution of violence capacities in society is relatively asymmetric, then concessions might need be made to stronger individuals in the initial social contract.[3] This can lead to the creation and persistence of a powerful coalition that is able to secure supranormal profits beyond its productive capacity. In Olson's framework this is the stationary bandit, and in North et al.'s framework this is the political *elite* in a limited-access order. In this chapter, I do not find conditions where a pure meritocracy (and thus possibly democracy) could emerge out of anarchy. Thus, these forms of government might require additional political infrastructure to be viable (for a thorough analysis of transition to democracy see e.g., Acemoglu & Robinson 2009).

I develop a model of the origin of the state as a stationary bandit, in the spirit of Olson (1993). Starting from an anarchic environment in which property is insecure and must be defended, I assess the potential cooperative pacts that might emerge in equilibrium.[4] The model presented in this chapter differs from prior research on predatory states (e.g., Grossman & Noh 1994, McGuire & Olson 1996, Moselle & Polak 2001, Robinson 2001, Skaperdas

[1] In Nozick's defense, his goal was to show how a legitimate government *could* emerge out of anarchy and not necessarily to describe how a government *would* emerge out of anarchy. The main assumption of Nozick's that is violated by this treatment of the origin of the state is the assumption of *moral agents*. Here, I instead assume that agents only seek to maximize their own economic payoff.

[2] See Appendix C.1 for a detailed discussion.

[3] Experiments on the endogenous formation of institutions have shown that even when participants are equally adept at taking, property will not necessarily be respected (e.g., Kimbrough, Smith & Wilson 2010, Wilson, Jaworski, Schurter & Smyth 2012). However, Wilson et al. (2012) shows that under some conditions a respect for property might emerge out of social custom.

[4] While these pacts are modeled explicitly, they need not be the result of conscious design. For example, Sugden (2004) shows how these sorts of cooperative norms might emerge through long run evolutionary processes.

2014, Konrad & Skaperdas 2012) in two key ways. Firstly, the prior work exogenously assumes the existence of a ruler, whereas the ruler(s) are determined endogenously in my model. By requiring that the ruling coalition emerge out of anarchy, I can more accurately capture the constraints faced by the ruler in any cooperative agreement. Typically, even an absolute monarch needs the support of her fellow elites to retain power, and it is difficult to capture such constraints when abstracting away from the process by which a ruler is selected. Myerson (2008) finds that when a ruler is subject to constraints, she may be able to make more credible promises than an unconstrained ruler could. Also, by restricting attention to coalitions of bandits that could emerge out of anarchy, I am able to draw comparisons between the structure of society in the emergent order and well known forms of government such as autocracy, oligarchy, and meritocracy.[5] Such comparisons are difficult to make when the structure of society is taken as given. Secondly, the prior work assumes symmetric agents (with the possible exception of the ruler), whereas I allow for a very general form of asymmetry between agents. By allowing the distribution of violence capacities to be asymmetric I can consider many cases that have not yet been studied, in addition to the cases studied in prior work. This model uses similar production and defense technologies to Konrad & Skaperdas (2012) and Skaperdas (2014), but I also explicitly model conflict between (possibly asymmetric) bandits under anarchy and this feature is not present in their model.

In the interest of simplicity, I too have made many abstractions, and so this work should be viewed as a complement rather than a supplement to the prior work on predatory states. I do not, for example, consider the case of a finite and endogenous time horizon, as do Grossman & Noh (1994), nor do I consider collective protection technologies, as do Konrad

[5]I use the term "meritocracy" to describe an order in which property rights are respected and each individual keeps the fruits of her own labor, less perhaps a small tax to support various public goods such as defense. In a meritocracy, no party is able to secure supranormal profits through the use of force. I prefer to use the term "meritocracy" rather than a term such as "democracy" because it does not have the additional context of a voting system, which is beyond the scope of this chapter.

& Skaperdas (2012) and Skaperdas (2014).[6] The public goods considered in this analysis are primarily respect for property and resistance to banditry. I do not consider productivity enhancing public goods, as is done in other studies (e.g., Grossman & Noh 1994, McGuire & Olson 1996, Moselle & Polak 2001, Robinson 2001) with mixed results for the welfare of producers.

In addition to the fundamental importance of this topic to political science and philosophy, a better understanding of the process of state formation can be applied to many modern day economic phenomena. In particular, the incentives faced by economic agents in an anarchic environment are similar to the incentives faced by economic agents in states that lack the economic and political institutions to protect property and/or limit the use of coercive force. For example, the observed behavior of urban squatters in Peru matches the predicted behavior of producers under anarchy in this model. Field (2007) finds that prior to an intervention by the Peruvian government in 1996 that assigned property rights to these squatters, they spent an average of 13.4 hours per week defending their residences. Also, the behavior of local government officials in Uganda when administering public expenditures is not unlike the behavior of bandits in this model. Reinikka & Svensson (2004) find that in the mid-1990s, Ugandan schools received only 13% of the grant money that was allocated to education by the government. The majority of the grant was instead captured by local government officials. Easterly (2001) details the many well-intentioned, but ultimately unsuccessful aid attempts by the World Bank over the years at spurring growth in developing countries. A recurring theme in the unsuccessful aid attempts is an incomplete understanding of the underlying incentives of the elite. Thus, an improved understanding of the incentives and constraints that led to the initial formation of the state will also help to improve our understanding of developing countries and the feasibility of various development aid interventions.

[6] Konrad & Skaperdas (2012) and Skaperdas (2014) find that while there are gains to producers organizing for collective protection, these forms of governance do not fare well against warlords and predatory states. This model shares the feature that "offense trumps defense." For an alternative analysis of state formation in which "defense trumps offense" see e.g., Dow & Reed (2013).

4.2 Static Model

The model society consists of $N \geq 2$ infinitely lived, risk neutral agents, indexed by ($i = 1, 2, ..., N$). Each agent i is endowed with a violence capacity $\theta_i \in \mathbb{R}_+$. For analytic convenience, I index the agents in order of violence capacity such that agent 1 has the largest violence capacity and agent N has the smallest violence capacity. In addition, I normalize the set of violence capacities $\theta = \{\theta_1, \theta_2, ..., \theta_N\}$ such that $\sum_{j=1}^{N} \theta_j = 1$. I refer to this normalized set θ as the *distribution of violence capacities*. This distribution can range from the uniform distribution, in which all agents are equally capable of violence, to a degenerate distribution, in which one agent has all of the violence capacity. Each agent knows the entire distribution of violence capacities and knows which violence capacity belongs to whom. Let

$$p(g_i, g_{-i}) = \frac{\theta_i g_i}{\theta_i g_i + \sum_{j \neq i} \theta_j g_j}, \tag{4.1}$$

denote agent i's probability of winning an N player asymmetric lottery contest (e.g., Gradstein 1995, Stein 2002), where g_i is the effort of the ith agent and g_{-i} is the effort of all other agents. Then θ_i can be interpreted as agent i's *ex ante* probability of winning the contest.

Agents have one inalienable unit of time to divide between labor, l, defense, x, and expropriation, g. Each agent i has access to the same production technology $F(l_i)$ that converts labor into output, with maximum feasible production $F(1) = 1$. In this anarchic environment, output is not secure and only a fraction $\alpha(x_i)$, $0 \leq \alpha(\cdot) \leq 1$, can be successfully defended from theft.[7] The remaining fraction, $1 - \alpha(x_i)$, is insecure. I assume $F(0) = 0$, $F'(l) > 0$, $F''(l) \leq 0$, $\alpha(0) = 0$, $\alpha(1) = 1$, $\alpha'(x) > 0$ and $\alpha''(x) \leq 0$. In keeping with the banditry motivation, I also assume labor and expropriation are incompatible activities. Labor requires settling down and working the land, whereas expropriation requires traveling

[7]While I opt to leave $\alpha(\cdot)$ as a general function, common specifications in related work include linear returns to defensive effort, $\alpha(x_i) = x_i$ (e.g., Skaperdas 2014), and a generalized lottery contest success function, $\alpha(x_i, g) = \frac{x_i}{x_i + \gamma g}$ for some $\gamma > 0$ (e.g., (e.g., Gonzalez 2007).

from farm to farm in order to take insecure output.[8] Thus agents may engage in labor or expropriation, but never both. The agents can be naturally divided into two groups: agents who spend time on labor, henceforth referred to as producers, and agents who spend time on expropriation, henceforth referred to as bandits. Let P denote the set of producers and B denote the set of bandits. Then $N = |P| + |B|$.

The returns to expropriation are determined through an N player asymmetric lottery contest with prize

$$V = \sum_{i \in P} \left[(1 - \alpha(x_i)) F(l_i) \right], \tag{4.2}$$

equal to the sum of all insecure output. When $B \neq \emptyset$, agent i's probability of winning the prize (or equivalently, her share of the loot) is given by (4.1). When, $B = \emptyset$ then each producer retains the insecure portion of her output.

4.2.1 Summary of Static Model

Ex Ante: Each agent i is endowed with a violence capacity θ_i and one inalienable unit of time.

Ad Interim: Agents simultaneously choose l, x, g such that $l_i + x_i + g_i \leq 1$ and either $l_i = 0$ or $g_i = 0$ for all i. V, B, and P are determined.

Ex Post: If $B \neq \emptyset$, the insecure output is taken and agents compete in an asymmetric lottery contest with prize V. Each agent i's contest effort is given by g_i. Otherwise, each producer retains the insecure portion of her output.

[8] I focus primarily on agrarian societies as the sedentary nature of farmers magnifies the social dilemma presented by a lack of property rights. Unlike a hunter or a shepherd, there are high costs for a farmer to relocate, and so it is more likely that a farmer would be forced to interact with the same bandit repeatedly in the long run. Anthropological evidence suggests that early agricultural societies made much larger investments in defense than did hunter gather societies (see e.g., Seabright 2008).

4.2.2 Anarchic Equilibrium

I first consider the static, non-cooperative Nash equilibrium, henceforth referred to as the *anarchic equilibrium.* This equilibrium consists of a vector of choices $s_i = (\hat{l}_i, \hat{x}_i, \hat{g}_i), \forall i$, such that no agent i has an incentive to unilaterally deviate given the actions of others.

The equilibrium choices completely determine the equilibrium P, B, and V, denoted $\hat{P}$, $\hat{B}$ and $\hat{V}$, as well as the equilibrium payoff to producers, Π_P^A, and the equilibrium payoff to each bandit $i \in \hat{B}$, $\Pi_{B,i}^A$.

Producers maximize the payoff $\Pi_P(l, x) = \alpha(x)F(l)$ such that $l + x \leq 1$.[9] The earlier conditions on $F(\cdot)$ and $\alpha(\cdot)$ are sufficient to ensure a unique optimum. Thus the anarchic equilibrium payoff to producers is given by

$$\Pi_P^A(\hat{l}) = \alpha(1 - \hat{l})F(\hat{l}). \tag{4.3}$$

This decision does not depend on the number of bandits $|B|$, and so (4.3) also represents the opportunity cost of banditry.[10]

Bandits maximize the payoff $\Pi_{B,i}(g_i, g_{-i}) = p(g_i, g_{-i})V$. The above constraints on l, x , and g, imply $\hat{g}_i = 1, \forall i \in B$. Thus, each bandit's *ex post* probability of winning the contest is given by

$$p(\hat{g}_i, \hat{g}_{-i}) = \frac{\theta_i}{\sum\limits_{j \in B} \theta_j} \tag{4.4}$$

[9]I do not consider the case where $\hat{B} = \emptyset$ as this will never be an equilibrium outcome for N sufficiently large (it requires $\hat{V} \leq 2$). In the case where $a(x_i) = x_i$, $N \geq 2$ will suffice.

[10]The fraction of output that can be secured from theft may, in practice, depend on the number of bandits $|B|$. However, the number of bandits is fixed local to a particular equilibrium and so this added generality does not meaningfully affect the results. Thus, I maintain the independence assumption for the sake of simplicity. This observation does highlight that caution should be excercised when making comparisons *across* equilibria in settings where the independence assumption is implausible.

and the anarchic equilibrium payoff to bandit i is given by

$$\Pi_{B,i}^{A}(\hat{g}_i, \hat{g}_{-i}) = p(\hat{g}_i, \hat{g}_{-i})\hat{V} = \frac{\theta_i}{\sum\limits_{j\in\hat{B}} \theta_j} \left|\hat{P}\right| (1-\alpha(\hat{x}))F(\hat{l}). \tag{4.5}$$

This payoff structure is similar to the one analyzed in Gradstein (1995) and Stein (2002), with the main departures being the endogeneity of the prize V and the additional constraints on expropriative effort. When an agent forgoes production and enters into the contest among bandits, she not only incurs an opportunity cost of Π_P^A, but she also reduces the prize for all bandits by $(1-\alpha(\hat{x}))F(\hat{l})$. This has several implications for the sets $\hat{B}$ that can arise in equilibrium.

Proposition 4.1. *Let $\mathbb{B}$ denote the set of all possible equilibrium $\hat{B}$ for a given set of parameters. Then the following claims are true:*

i For all $\hat{B} \in \mathbb{B}$, $i \in \hat{B}$, $\Pi_{B,i}^A \geq \Pi_P^A$.

ii There exists a $\hat{B} \in \mathbb{B}$, such that $\hat{B}$ consists of the $\left|\hat{B}\right|$ agents with the most violence capacity. This set, denoted $\hat{B}_{viol}$, is unique up to a permutation of identities for agents with identical violence capacities.

iii For some values of the parameters, there exist $\hat{B} \in \mathbb{B}$ such that $\hat{B} \neq \hat{B}_{viol}$.

Proof. (i) Suppose $i \in \hat{B}$ and $\Pi_{B,i}^A < \Pi_P^A$. Then agent i could profitably deviate to production by choosing $(\hat{l}, \hat{x}, 0)$ which implies $i \notin \hat{B}$. Thus $\Pi_{B,i}^A \geq \Pi_P^A$.

(ii) Suppose not. Pick some $\hat{B} \in \mathbb{B}$ such that $\hat{B}$ does not consist of the $\left|\hat{B}\right|$ agents with the most violence capacity. Then, by definition, there exists at least one agent $i \in \hat{B}$ with less violence capacity than some agent $j \notin \hat{B}$ and by Proposition 4.1 (i), $\Pi_{B,i}^A \geq \Pi_P^A$. Furthermore, we can see from (4.5) that $\Pi_{B,i}^A$ is strictly increasing in θ_i. Pick the agent $j \notin \hat{B}$ with the most violence capacity. By assumption, $\theta_j > \theta_i$. This implies that if agent j

were to replace agent i in the set of bandits (i.e., if $B = \hat{B} \cup \{j\} \backslash \{i\}$), agent j would also receive a payoff greater than Π_P^A (i.e., $\Pi_{B,j}^A > \Pi_P^A$). As agent j does not have an incentive to deviate from $B = \hat{B} \cup \{j\} \backslash \{i\}$, either $B = \hat{B} \cup \{j\} \backslash \{i\} \in \mathbb{B}$ or some other agent wants to switch to production. Suppose some other agent k wants to switch to production. Then k must be receiving a payoff below Π_P^A which implies $\theta_k < \theta_j$. Once k switches to production, either $\hat{B} \cup \{j\} \backslash \{i, k\} \in \mathbb{B}$, or some agent $l \in \hat{B} \cup \{j\} \backslash \{i, k\}$ ($\theta_l < \theta_j$) also wants to switch to production, or some agent $m \notin \hat{B} \cup \{j\} \backslash \{i, k\}$ wants to switch to banditry ($\theta_m < \theta_j$ as j was the agent with the most violence capacity not already in the set $\hat{B}$). Regardless of the precise response to the original switch to $B = \hat{B} \cup \{j\} \backslash \{i\}$, this response will only involve adding and removing agents with less violence capacity than j and thus it will never be a best response for j to switch back to production in this chain of events. Eventually the process will settle on some $\hat{B}' \in \mathbb{B}$ that contains j. Either $\hat{B}'$ consists of the $\left|\hat{B}'\right|$ agents with the most violence capacity or not. But if not, then by similar logic we can replace at least one agent in $\hat{B}'$ with some agent not in $\hat{B}'$ that has more violence capacity and again we will arrive at some new equilibrium $\hat{B}'' \in \mathbb{B}$ that contains both j and this new agent. However many times this process must be repeated, eventually we will arrive at some equilibrium in $\mathbb{B}$ that consists of the agents with the most violence capacity, which contradicts our original assumption.

(iii) For example, when $N = 3$, $\theta_1 > \theta_2 = \theta_3$, $F(\hat{l}) = 1/2$, and $\alpha(\hat{x}) = 1/3$, there exists an equilibrium such that $\hat{B} = \{2, 3\}$. Here $\hat{B} \neq \hat{B}_{viol}$ because $1 \notin B$. $\square$

Part (i) of Proposition 4.1 states that all bandits receive a weakly greater payoff than producers, part (ii) states that an equilibrium exists where the strongest agents all choose banditry, and part (iii) states that this need not be the unique equilibrium. This multiplicity of equilibria presents an equilibrium selection problem. Of the many possible equilibria that exist for a given set of parameters, we would like to know which equilibrium is the most plausible prediction of behavior. As detailed in Proposition 4.2 below, an equilibrium where the strongest agents choose banditry has several notable properties that make it a plausible candidate. Before stating the proposition, I outline the thought process behind this claim in Section 4.2.3.

4.2.3 Evolutionary Model of Role Selection

The model does not, in general, yield a unique anarchic equilibrium. This presents a natural coordination problem. An agent might be a bandit in one equilibrium and a producer in another. Even when every agent intends to play some equilibrium strategy, without a mechanism to coordinate actions there is no guarantee that the realized strategy profile will correspond with one of the equilibria. In addition, different agents may prefer different equilibria due to the adversarial nature of the interaction, and so there is no clear focal equilibrium. Thus, we should not necessarily expect the agents to immediately coordinate on an anarchic equilibrium. However, if roles were not consciously selected and instead evolved over time, the evolutionary process might naturally gravitate toward certain equilibria. In this section, I present an evolutionary model of role selection to assess the properties of various equilibria in the static model.

Consider a population of N agents ($i = 1, 2, ..., N$), interacting over an infinite number of discrete periods ($t = 0, 1, 2, ...$). Each agent i has violence capacity θ_i.

Agents play one of two strategies, p or b. The payoff to strategy p is Π_P^A, given by (4.3), and the payoff to strategy b is $\Pi_{B,i}^A$, given by (4.5), where B is the set of agents currently choosing strategy b and P is the set of agents currently choosing strategy p. In period $t = 0$, each agent plays some initial strategy. In all future periods, one agent is selected at random and given the opportunity to revise her strategy. I assume that agents revise their strategies according to the *logit choice* protocol, given by

$$\rho_{ij} = \frac{\exp(\eta^{-1}\pi_j)}{\exp(\eta^{-1}\pi_p) + \exp(\eta^{-1}\pi_b)}, \tag{4.6}$$

where ρ_{ij} is the likelihood of a revision from strategy i to strategy j, $i, j \in \{p, b\}$, π_i is the payoff to strategy i, and η is the noise level.

Under this protocol, agents myopically best respond to the strategies being played by the rest of the population with some possibility of error (parameterized by $\eta > 0$). Each strategy has a positive probability of being played, but costly errors are penalized. Thus, a best response is always the most likely revision, and the likelihood of a non-best response is decreasing in the cost (in terms of forgone payoff) of the error. As $\eta \to 0$, this revision protocol approaches a pure best response protocol.

This game admits an irreducible Markov process yielding a stationary distribution μ, where μ_s is the proportion of the time that strategy profile $s \in S$ is played in the long run. As $\eta \to 0$ and $t \to \infty$, the proportion of the time spent in each state s by the evolutionary process approaches μ_s. A state s is termed *stochastically stable* if $\mu_s > 0$, i.e., if the strategy profile is played positive proportion of the time in the long run (Foster & Young 1990). By proposition 4.2 below, the anarchic equilibrium where $B = \hat{B}_{viol}$ is the unique stochastically stable state (i.e., $\mu_{B=\hat{B}_{viol}} = 1$). Thus the evolutionary process to spends almost all of the time in this equilibrium as $\eta \to 0$ and $t \to \infty$ (Young 1998).

Proposition 4.2. *$\hat{B}_{viol}$ has the following properties:*

i $\left|\hat{B}_{viol}\right| \leq \left|\hat{B}\right|$ *for all* $\hat{B} \in \mathbb{B}$.

ii $\hat{B}_{viol}$ *is the unique stochastically stable state in the evolutionary game of role selection.*

Proof. (i) By Proposition 4.1 (i) All agents $i \in \hat{B}_{viol}$ earn a payoff greater than or equal to

Π_P^A. Furthermore, by definition, $\hat{B}_{viol}$ consists of the $\left|\hat{B}_{viol}\right|$ agents with the most violence capacity. Construct a new set B such that $|B| < \left|\hat{B}_{viol}\right|$ and such that B consists of the $|B|$ agents with the most violence capacity. Then $B \subset \hat{B}_{viol}$ and any agent $i \in \hat{B}_{viol}$, $i \notin B$ would earn a payoff greater than or equal to Π_P^A by joining B (otherwise $\hat{B}_{viol}$ could not be an equilibrium set). Thus,agent i earns a payoff greater than or equal to Π_P^A by joining even most capable set of bandits of size $|B| < \left|\hat{B}_{viol}\right|$, which implies that this agent would earn a payoff greater than or equal to Π_P^A for by joining any set of size $|B| < \left|\hat{B}_{viol}\right|$. This implies $B \notin \mathbb{B}$. Therefore, there does not exist a $\hat{B} \in \mathbb{B}$ such that$\left|\hat{B}_{viol}\right| > \left|\hat{B}\right|$.

(ii) Consider a transition from some equilibrium $\hat{B} \in \mathbb{B} \diagdown \hat{B}_{viol}$ to $\hat{B}_{viol}$. Under the logit choice protocol, the *cost* of an action is the absolute difference in payoff between this action and the best response to the population strategy profile. Note that the cost of a relatively strong agent choosing banditry when production is a best response is less than the cost of a relatively weak agent choosing banditry when production is a best response. Also note that the cost of a relatively strong agent choosing production when banditry is a best response is greater than the cost of a relatively weak agent choosing production when banditry is a best response. This follows from the payoff to banditry, given by (4.5), which is strictly increasing in θ_i. Recall that the *basin of attraction* of a particular equilibrium is the set of strategy profiles that can reach the equilibrium through a sequence of individual best responses. The *radius* of this basin of attraction is the minimum cost route required to escape the basin of attraction and the *coradius* of this basin of attraction is the maximum over all other equilibrium states of the minimum cost route required to reach the basin of attraction. A transition from some $\hat{B}$ into $\hat{B}_{viol}$ involves adding relatively strong agents to the set of bandits and taking away relatively weak agents, whereas a transition out of $\hat{B}_{viol}$ into some $\hat{B}$ involves adding relatively weak agents to the set of bandits and taking away relatively strong agents. Therefore the radius of a transition from $\hat{B}$ into $\hat{B}_{viol}$ is necessarily less than the radius of a transition out of $\hat{B}_{viol}$ into $\hat{B}$ for all $\hat{B} \in \mathbb{B} \diagdown \hat{B}_{viol}$, which implies that the radius of $\hat{B}_{viol}$ is greater than the coradius of $\hat{B}_{viol}$.[11] By Alos-Ferrer & Netzer (2010) Proposition 3, this implies that $\hat{B}_{viol}$ is the unique stochastically stable state of the evolutionary process. □

Part (i) of Proposition 4.2 states that $\hat{B}_{viol}$ has the smallest number of bandits (or equivalently, the most producers) out of all possible equilibrium $\hat{B}$. Part (ii) states that, in the long run, the set $\hat{B}_{viol}$ will emerge as the steady state of an evolutionary process whereby bandit and producer roles are selected repeatedly over time. Thus starting from any initial condition, if agents are less likely to make very costly errors than minor errors, we should expect $\hat{B}_{viol}$ to emerge as the equilibrium set of bandits.

Proposition 4.2 suggest that an equilibrium where the strongest agents choose banditry has several advantages over equilibria that do not. Firstly, by Proposition 4.2 (i), these equilibria have the highest aggregate production in the class of all anarchic equilibria. Thus, equilibria in which bandits are relatively weak are less socially efficient than equilibria in which bandits are relatively strong. Secondly, evolutionary processes are more likely to select equilibria where strong agents choose banditry. The accidental entry of a strong bandit into a weak bandit equilibrium is much more costly for the weak bandits than for the strong bandit. Thus, weak bandit equilibria are not particularly robust to random perturbations, whereas strong bandit equilibria are more resilient. If society spent some time in an anarchic environment before coordinating on a cooperative outcome, we might expect the bandits to be relatively strong compared to producers in the emergent steady state.

[11]The obvious exception to this statement is when $\hat{B}_{viol}$ itself is not unique (then its radius would equal its coradius), but as these equilibria are identical up to a permutation of identities for agents with identical violence capacities I group these equilibria into the same equivalence class.

4.3 Dynamic Model

Section 4.2 considered anarchy as a static non-cooperative baseline. While the equilibrium where the strongest agents choose banditry stood out as a plausible prediction of play, this equilibrium is by no means ideal. Society can feasibly produce N and instead produces $(N - |\hat{B}|)F(\hat{l})$ under anarchy. Not every member of society contributes to production, and those members that do contribute spend a fraction of their time on non-productive defensive effort. In this fixed environment, where the same individuals encounter each other repeatedly, there may be substantial gains from cooperation. Bandits in particular, whose payoffs are proportional to the sum of all insecure output, could benefit from contracting with the producers to ensure that producer efforts are spent on production rather than defense. By agreeing not to take all the insecure output, bandits enable the producers to spend more time on production, thus increasing the size of the economic rent accruing to the bandits.[12] This is the critical observation behind Olson's (1993) stationary bandit hypothesis. Furthermore, the bandits might also benefit from cooperating with one another if this were to enable a subset of the bandits to become producers.

Consider the following adaptation of the static model defined in Section 4.2. The static game of production, defense, and expropriation is now infinitely repeated ($t = 0, 1, 2, ...$). Agents share a common discount factor $\delta \in (0, 1)$ and seek to maximize the discounted sum of all future payoffs, given by

$$\Pi_i = \sum_{t=0}^{\infty} \delta^t \Pi_{i,t}. \tag{4.7}$$

I assume that agents can freely and costlessly transfer payoffs at the end of each period (after actions are observed), where $\tau_{i,t}^j$ denotes the transfer made from agent i to agent j in period

[12] Here I am implicitly assuming that the bandits can fully observe the effort of the producers. Mayshar, Moav & Neeman (2012) show that when the effort of producers is less transparent, this limits the extent to which bandits can capture the surplus from production.

t. Thus the payoff for a producer in period t is given by

$$\Pi_{P,i,t} = \alpha(x_{i,t})F(l_{i,t}) - \sum_{j \neq i} \tau_{i,t}^{j} + \sum_{k \neq i} \tau_{k,t}^{i} \tag{4.8}$$

and the payoff for a bandit in period t is given by

$$\Pi_{B,i,t} = p(g_{i,t}, g_{-i,t})V_t - \sum_{j \neq i} \tau_{i,t}^{j} + \sum_{k \neq i} \tau_{k,t}^{i}. \tag{4.9}$$

Transfers are assumed to be voluntary and non-negative, and any promise of future transfers can be reneged upon.

4.3.1 Summary of Dynamic Model Period

Ex Ante: Each agent i is endowed with a violence capacity θ_i and one inalienable unit of time.

Ad Interim: Agents simultaneously choose l_t, x_t, g_t such that $l_{i,t} + x_{i,t} + g_{i,t} \leq 1$ and either $l_{i,t} = 0$ or $g_{i,t} = 0$ for all i. V_t, B_t, and P_t are determined. Each agent i also chooses a set of transfers $\{\tau_{i,t}^{j}\}_{j \neq i}$ to make at the end of the period.

Ex Post: If $B_t \neq \emptyset$, the insecure output is taken and agents compete in an asymmetric lottery contest with prize V_t. Each agent i's contest effort is given by $g_{i,t}$. Otherwise, each producer retains the insecure portion of her output.

In the analysis that follows, I consider only *subgame perfect equilibria*, i.e., the class of equilibria in which for all periods $t = 0, 1, 2, ..$, the strategy profile played from period t onward is also an equilibrium of the subgame that begins in period t. This equilibrium refinement ensures that the strategies utilized are *time consistent*. That is, no agent wants to revise her strategy profile any in period $t \neq 0$, given the strategy profiles of the other agents.

Furthermore, I restrict attention to equilibria that satisfy the property of *role stationarity.*

Definition 4.1. *An equilibrium of the dynamic model satisfies "role stationarity" if and only if* $\hat{B}_t = \hat{B}_s = \hat{B}$ *and* $\hat{P}_t = \hat{P}_s = \hat{P}$ *for all periods* t *and* s $(t, s \in \{0, 1, 2, ...\})$*.*

This restriction ensures that each agent has a well defined role, producer or bandit, along the equilibrium path that does not vary over time and thus that the emergent order resembles some form of steady state. However, this restriction does not prevent an agent from switching roles off the equilibrium path (e.g., to deviate from the cooperative equilibrium strategy or to enforce a punishment strategy).

4.3.2 Feasible Payoffs

To characterize the set of feasible cooperative equilibria, it is useful to first consider the range of payoffs that can be supported in equilibrium. In the first best case, all N agents choose $l = 1$, resulting in maximum feasible production N. The equilibrium payoffs are then achieved through a vector of transfers $\tau_{i,t} = (\tau_{i,t}^1, \tau_{i,t}^2, ..., \tau_{i,t}^N)$ for all i, t such that the net payoff to each agent is the agreed upon equilibrium payoff. These payoffs are characterized in Proposition 4.3 below. Before stating the proposition, I outline the credible punishment strategies that could be used to enforce such an equilibrium.

Recall that each agent's *minimax payoff*, i.e. the maximum payoff that an agent could secure if all other agents chose a strategy profile intended to minimize this agent's payoff, is given by Π_P^A. Even if all other agents choose banditry, this agent could secure Π_P^A by choosing the optimum level of labor and defense under anarchy. Furthermore, even the agent with the highest violence capacity could not do better by switching to banditry as this would result in $V = 0$ and thus $\Pi_B = 0$. However, all other agents choosing banditry is not a credible punishment. Namely, to enforce cooperation through non-equilibrium actions, the punishment strategy must be agent-specific. If all other agents were to choose banditry, they would punish the intended agent, but they would punish themselves even more in the process. For an *agent-specific punishment strategy* to exist, there must exist an action profile such that agent i's payoff is below Π_i^A, and all other agents either (1) receive a payoff greater than Π_i^A or (2) are unaffected by the punishment. It is easy to see that no agent's payoff can be reduced below Π_P^A because any agent can guarantee herself this payoff with the appropriate action. A punishment for agent $i \in \hat{B}$ cannot involve alternative choices of l and x by a producer because this would reduce the producer's payoff below Π_P^A without reducing agent i's payoff below Π_P^A. It also cannot involve the removal of any agents $j \neq i$ from B because these agents, who were receiving $\Pi_{i,B}^A > \Pi_P^A$, would now receive Π_P^A. Thus, payoffs below $\Pi_{B,i}^A$ for agent $i \in \hat{B}_{viol}$ can only be supported if there exists a $\hat{B} \in \mathbb{B}$ such that $i \notin \hat{B}$ and, for all $j \neq i \in \hat{B}_{viol}$, $j \in \hat{B}$. In other words, the only possible agent-specific punishment strategy for player i is reversion to some other equilibrium that excludes i and only i from the set of bandits. Because this class of punishment strategy does not, in general, exist, I assume that the lower bound for payoffs is Π_i^A in the remaining analysis.[13]

Proposition 4.3. *For δ sufficiently large, any vector of payoffs $\Pi_t = (\Pi_{1,t}, \Pi_{2,t}, ..., \Pi_{N,t})$ such that $\Pi_{i,t} \in (\Pi_i^A, N - \sum_{j\neq i} \Pi_j^A)$ for all i and $\sum_{i=1}^N \Pi_{i,t} \leq N$ can be supported in a cooperative equilibrium.*

Proof. First note that Π_i^A is agent i's payoff under anarchy, and so, for δ sufficiently high, any payoff greater than this for agent i can be supported in a cooperative equilibrium by the threat of reversion to the anarchic equilibrium. The upper bound on individual payoffs and the total aggregate payoff are pinned down by the maximum feasible output N. As every agent must receive a payoff greater than Π_i^A and N is the maximum feasible output, every agent i must receive less than $N - \sum_{j\neq i} \Pi_j^A$ and all agents collectively cannot receive

[13] Typical results that prove the existence of agent-specific punishment strategies for the case of non-equivalent utilities (e.g., Abreu, Dutta & Smith 1994) do not hold here due to non-convexity of the action space.

a payoff greater than N. $\square$

By Proposition 4.3, we can see that, for δ sufficiently high, any division of aggregate output that guarantees each agent a payoff greater than Π_i^A can be supported in a cooperative equilibrium. This result is weaker than a typical "folk theorem," as bandits must receive a payoff higher than their minimax payoff, Π_P^A, even as $\delta \to 1$. Payoffs below Π_i^A for agent i can not, in general, be supported due to the possible non-existence of agent-specific punishment strategies.

The above result pins down the threat points for a cooperative outcome when agents are sufficiently patient. Aggregate output is N in each period and all agents are guaranteed at least their anarchic equilibrium payoff. The remaining surplus, $N - (N - |\hat{B}_{viol}|)F(\hat{l})$, is somehow divided between the agents. When bandits hold all the bargaining power as is typically assumed in the literature on predatory states (e.g., Grossman & Noh 1994, McGuire & Olson 1996, Moselle & Polak 2001, Robinson 2001, Skaperdas 2014, Konrad & Skaperdas 2012), the producers are no better off than they were under anarchy. Thus, while the model supports Olson's (1993) stationary bandit hypothesis that a bandit would have an incentive to reach a long-run cooperative agreement with the producers, the producers are not necessarily beneficiaries of the arrangement.

The more interesting case is when patience is limited. In this case, each agent must be guaranteed a strictly greater payoff than Π_i^A. The cooperative payoff for agent i, denoted by Π_i^C, is bounded below by the payoff from a one shot deviation to banditry, denoted by Π_i^D, followed by a reversion to the anarchic equilibrium. More formally,

$$\Pi_i^C \geq (1-\delta)\Pi_i^D + \delta\Pi_i^A. \tag{4.10}$$

It is easy to see that the lower bound for Π_i^C approaches Π_i^A only as $\delta \to 1$, otherwise the agents must be paid a premium to ensure long-run cooperation. In addition to δ, this premium depends on Π_i^D which may be influenced by the precise nature of the cooperative pact chosen. This avenue is explored further in Section 4.3.3.

In the following sections, I will distinguish between the *coalition in power* and the rest of society. The coalition in power in my model is analogous to the principal in a principal-agent model, with the caveat that the coalition may consist of multiple individuals working collectively toward the same goal. While the theory is silent as to the relative bargaining power of each agent in the society, it is typical to assume in the literature on predatory states (e.g., Grossman & Noh 1994, McGuire & Olson 1996, Moselle & Polak 2001, Robinson 2001, Skaperdas 2014, Konrad & Skaperdas 2012) that the state is the principal who sets the terms of the cooperative arrangement, holds all the bargaining power, and can extract any surplus beyond the minimum required to ensure cooperation.[14] Thus, in keeping with Olson's (1993) stationary bandit hypothesis, in which the bandit settles down and becomes the *de facto* government, it is natural to assume that the coalition in power is the equilibrium set of bandits $\hat{B}_{viol}$ that is predicted to emerge through long run evolutionary processes. I assume that the coalition in power acts as a single decision maker, and divides the surplus from cooperation according to some sharing rule that satisfies each members individual rationality constraint (4.10).

4.3.3 Deterring Banditry

A key result from Section 4.3.2 is that when patience is limited, each agent i must be paid a premium above Π_i^A to deter one-shot deviations to banditry. When δ is relatively low, this premium can be substantial. For example, when society is at the production possibilities frontier, $\Pi_i^C \geq (1-\delta)(N-1) + \delta\Pi_i^A$ for all i. This is because, at the production possibilities

[14] To quote Oppenheimer (1908, p. 32): "The moment when first the conqueror spared his victim in order permanently to exploit him in productive work was of incomparable historical importance. It gave birth to the nation and State, to right and the higher economics, with all the developments and ramifications which have grown and which will hereafter grow out of them."

frontier, $B = \emptyset$ and $V = N$, and so any agent who deviates to banditry would be uncontested in taking all of society's insecure output. Even if the coalition in power were willing to pay this premium, it may not be feasible to do so. In particular, when the inequality

$$\delta < \frac{\Pi^D - 1}{\Pi^D - \Pi_P^A} \tag{4.11}$$

holds, the required aggregate payoff to ensure cooperation exceeds the maximum feasible output of society. In the case where $\Pi_P^A = 0$ and $B = \varnothing$, this would be problematic for any $\delta < (N-2)/(N-1)$, and in the case where $\Pi_P^A > 0$ this would be problematic for an even larger range of δ.

Even when it is feasible to pay all agents to cooperate at the production possibilities frontier, the coalition in power may not find it optimal to produce N units of output. Instead, the coalition seeks to maximize the remaining surplus after giving each other agent the minimum payoff required to ensure cooperation. This surplus is not necessarily maximized when all agents choose production, because the one-shot deviation payoff Π_i^D for agent i depends on the relative strength of the agents in B. By having a subset of the agents choose a positive level of expropriative effort g_t, the coalition can lower each agent's one-shot deviation payoff, and thus lower the minimum Π_i^C required to ensure cooperation.[15] As all members of the coalition in power are, in a sense, bandits, the notion that some members of the coalition might choose a positive level of expropriative effort as part of a cooperative equilibrium requires additional terminology. I henceforth refer to these individuals as *guards.*

Definition 4.2. *A "guard" is any member of the coalition in power that forgoes production and instead chooses a positive level of expropriative effort as part of a cooperative equilibrium.*

Guards compete in an asymmetric lottery contest with prize V, as does a bandit under

[15] It is important to note that the presence of agents choosing positive levels of expropriative effort does not mean that society reverts back to anarchy. Instead, the effort of these agents can be thought of as a deterrent to banditry. The cooperative payoffs are still acheived through the transfers $\tau_{i,t}$.

anarchy, but are typically required to transfer back a portion of this insecure output as part of the cooperative pact to maintain the incentives for producers to cooperate.

Proposition 4.4. *Let G denote the set of guards in a cooperative equilibrium, and let G^* denote the coalition in power's most preferred set of guards. Then the following claims are true:*

i *G^* consists of the $|G^*|$ agents with the most violence capacity.*

ii *$|G^*|$ is non-increasing in δ.*

Proof. (i) Strong guards provide a greater reduction in the required Π_i^C than weak guards for the same cost of 1 forgone unit of output. Therefore, the surplus accruing to the coalition in power will only be maximized when G consists of the $|G|$ agents with the most violence capacity.

(ii) An increase in δ uniformly decreases the minimum Π_i^C required to ensure cooperation for any given G. This, in turn, decreases the marginal contribution to the surplus of all $i \in G$. Whenever a guard's marginal contribution falls below 1, the coalition in power would prefer for this guard to switch to production, thus decreasing $|G^*|$. □

Part (i) of Proposition 4.4 states that, all else equal, stronger guards are preferred to weaker guards, because a strong guard has a greater discouragement effect on banditry than a weak guard. Part (ii) states that the optimal number of guards depends monotonically on δ. In general, higher δ requires less guards.

Proposition 4.4 suggests that the optimal set of guards G^* can be found using the following iterative process. Starting from the production possibilities frontier, the coalition in power compares its surplus under $G = \varnothing$ to its surplus under $G = \{1\}$. If the latter surplus is higher, then agent 1 is added to G. Then the surplus under $G = \{1\}$ is compared to the surplus under $G = \{1, 2\}$. If the latter surplus is higher, then agent 2 is added to G, etc. This process continues until adding an agent would reduce the surplus. G^* is the former set in this last comparison.

The possibility of the coalition in power choosing $G^* \neq \varnothing$ is particularly notable because, in addition to maintaining internal order and discouraging deviations from the cooperative equilibrium, guards provide society with a natural resistance to the *external* threat of roving bandits. Such an external threat is not explicitly modeled here, but the logic is as follows: suppose some new agent $j \notin \{1, 2, ..., N\}$ with violence capacity θ_j were to enter the society unexpectedly for one period and resort to banditry. This agent's payoff $\Pi_{B,j} = p(g_j, g_{-j})V$ is decreasing in $|G|$, as each agent $i \in G$ decreases $p(g_j, g_{-j})$ and the corresponding share of the insecure output that bandit j could secure. Thus the coalition in power, acting in its own self-interest, ultimately may provide its citizens with protection against roving bandits as Olson (1993) hypothesized. When the threat of roving bandits is explicitly incorporated into the model, the coalition may even increase $|G^*|$ beyond the optimal size discussed here to further discourage entry by roving bandits.

4.3.4 The Limits of Cooperation

When δ is very low, the coalition must continue to expand G to ensure that a cooperative outcome can be reached. In some cases, G^* may even be a superset of $\hat{B}_{viol}$. Recall that even the weakest agent in $\hat{B}_{viol}$ is able to secure payoff above Π_P^A by choosing banditry under anarchy. When the output per producer increases from $\alpha(\hat{x})F(\hat{l})$ to 1, it is possible that some agents $i \in \hat{P}$, who under anarchy can do no better than Π_P^A, will now be able to secure a payoff above Π_P^A by making a one-shot deviation to banditry. Thus, the most extreme case in which cooperation can be maintained is when $|G|$ is increased to the point that no agent remaining in P could profitably deviate to banditry, denoted $G_{\max}$. In the limit, no cooperative outcome is feasible, as eventually δ drops below the threshold where the guards are still willing to make transfers $\tau_{i,t}$ to the producers, but this limiting case is no longer consistent with Olson's (1993) conception of the stationary bandit, and instead approaches the case of a roving bandit. Thus, it seems that in general, under the conditions that Olson (1993) had in mind, some form of cooperation is sustainable.

4.3.5 Extreme Distributions of Violence Capacities

It is clear from the preceding analysis that the distribution of violence capacities plays an important role in determining the division of resources in a cooperative equilibrium. In this section, I consider two polar distributions of violence capacities, uniform and degenerate, and compare the relative payoffs of bandits and producers in the cooperative equilibria. In particular, I only consider the case of limited patience because the cooperative payoffs are invariant to the distribution of violence capacities when $\delta \to 1$.

First, consider the degenerate case in which one agent has all the violence capacity and the others have none (i.e., $\theta_1 = 1$, and $\theta_i = 0$ for all $i > 1$). For N sufficiently large, there is a unique non-cooperative equilibrium where agent 1 chooses banditry ($\hat{B} = \{1\}$) and the remaining agents choose production. Agent 1 receives a payoff of $(N-1)(1-\alpha(\hat{x}))F(\hat{l})$ and all other agents receive a payoff of $\Pi_P^A = \alpha(\hat{x})F(\hat{l})$. To find the cooperative equilibrium, it is useful to note that, when $1 \in G$, the payoff from producer i defecting is 0 because this agent will not be able to expropriate any of the insecure output. Thus, it is easy to see that $G^* = \{1\}$ for all but the highest of discount factors (i.e., the case where $\delta = 1$). Choosing $G \supset \{1\}$ is never optimal because the marginal contribution of any other guard to the bandit's surplus is 0. Therefore, when $G^* = \{1\}$, the minimum payoff required to ensure cooperation is Π_P^A for all $i \neq 1$. In this case, the lone bandit can secure all the benefits from cooperation, the producers are no better off in the cooperative equilibrium, and the emergent order resembles autocracy.

Next, consider the case of a uniform distribution of violence capacities in which all agents are equally adept at violence. The non-cooperative equilibrium is unique up to a permutation of identities and includes $|\hat{B}|$ bandits and $N - |\hat{B}|$ producers. The number of bandits and producers depends on the specification of $\alpha(\cdot)$ and $F(\cdot)$. For example, in the linear case analyzed by Skaperdas (2014), where $\alpha(x) = x$ and $F(l) = l$, $|\hat{B}| = \lfloor N/2 \rfloor$ and $|\hat{P}| = \lceil N/2 \rceil$. In general, when the distribution of violence capacities is uniform, $\Pi_B^A = \Pi_P^A$ whenever possible. If $\Pi_B^A > \Pi_P^A$ it is only because $|\hat{B}|$ and $|\hat{P}|$ are restricted to integer values. The equality will hold as $N \to \infty$, and will hold approximately for any large N. Suppose $\Pi_B^A = \Pi_P^A$. In a cooperative equilibrium, G^* could range from $\varnothing$ to $G_{\max}$, depending on the value of δ. When $G = G_{\max}$, by definition the minimum required Π_i^C to ensure cooperation is Π_P^A because no agent can profitably deviate to banditry. But $G = G_{\max}$ implies that the cooperation payoffs for guards are as close as possible to the cooperation payoff for producers. Thus, whenever possible, $G = G_{\max}$ will result in $\Pi_G = \Pi^C = \Pi_P^A$, effectively driving the surplus accruing to the coalition in power down to 0. Therefore, $G^* \subset G_{\max}$, necessarily implies that the surplus accruing to the coalition in power is greater than 0. This implies that the coalition in power, which faces the same cooperation constraint as the producers, will receive strictly higher payoffs than the producers when the coalition holds the bargaining power. Thus, even in the the case of uniform violence capacities, we would expect the existence of a privileged elite in the emergent order unless the producers were able to collectively organize to secure a share of the rents. However, the emergent order in this case will still be relatively meritocratic compared to the orders that emerge under more extreme distributions of violence capacities. Producers are better off in the cooperative equilibrium than under anarchy for most values of δ, and the rents accruing to the coalition in power are relatively small.

4.4 Implications

Section 4.2 developed a model in which property is insecure and must be protected. In the non-cooperative case, we saw that property is not respected in equilibrium and that instead a subset of society resorts to banditry, taking as much as possible from the producers. We also saw that in this non-cooperative case, producers do not produce the maximum feasible output because they must spend time defending this output from bandits. However, as detailed in Section 4.3, when society is engaged in a long-run interaction (i.e., when bandits are "stationary"), we saw that it could benefit from a cooperative pact. This pact includes a primitive respect for property rights, in which each producer is entitled to keep a fraction of her insecure output. This fraction may depend, in part, on her potential payoff from breaking the pact. While this pact primarily benefits the strongest members of society, even would be producers benefit when patience is limited and a payoff premium is required to ensure cooperation. In addition to providing a primitive respect for property rights, the coalition in power may naturally increase society's resistance to the threat of roving bandits through the use of guards.

Starting from the primitive of a distribution of violence capacities and a simple production economy, we see how a natural state might develop out of anarchy. The orders that emerge from this framework mimic many key features of the primordial government envisioned by Olson (1993) and the limited access order described by North et al. (2009), lending credence to the idea that government may have emerged as a solution to violent conflict. This chapter adds to the literature on predatory states (e.g., Grossman & Noh 1994, McGuire & Olson 1996, Moselle & Polak 2001, Robinson 2001, Skaperdas 2014, Konrad & Skaperdas 2012) in several key ways. Firstly, in this model, the structure of the coalition in power is endogenously determined rather than exogenously assumed. This necessarily yields insights into the constraints faced by the ruler that are difficult to capture when the relationship between ruler and subject is taken as given. Secondly, by allowing the agents to be asymmetric, I am able to study broader class of social orders than when constrained to the symmetric case. One key finding is that differences in the distribution of violence capacities help to determine whether the order that emerges is relatively meritocratic, as typically assumed by proponents of the workmanship model of property (e.g., Locke 1689, Nozick 1977), or relatively oligarchic, as we see in limited access orders (e.g. Olson 1993, North et al. 2009). In extreme cases, even autocracy can emerge. However, given any distribution of violence capacities, there exists the potential for a privileged elite, earning supranormal payoffs, to form. Other forms of government such as pure meritocracy are not predicted to emerge in equilibrium. Thus this model provides support for theories that view various forms of non-democracy as a necessary precursor to democracy (e.g., Acemoglu & Robinson 2009, North et al. 2009, Fukuyama 2011). While this model helps to explain the origin of the natural state, how this state transitions to other forms of government remains an open question for future research.

Chapter 5

Conclusion

The studies described in Chapters 2, 3, and 4 collectively and individually contribute to our understanding of human behavior in large group interactions, but these studies are only the starting point for a larger research agenda.

Recent innovations in the penny auction industry have yielded new auction mechanisms that attempt to better exploit behavioral biases. Studying these mechanisms would not only further our understanding of penny auctions, but also our understanding of human behavior in general.

There remain many open questions about the role that limited observation plays in network formation. Further research could study whether the findings from Chapter 3 generalize to other network formation settings, such as unilateral connections, one-way flow networks, and limited information about the gains from connection.

The model of the origin of the state described in Chapter 4 generates empirically testable predictions. In particular, virtual world experiments provide a promising tool to better understand the ecological process by which social order could be achieved in a stateless environment. The model also helps to inform our understanding of the dynamics of state formation and the potential problems that can arise when a state-building intervention creates a power vacuum.

Bibliography

Abreu, D., Dutta, P. K. & Smith, L. (1994), 'The folk theorem for repeated games: A neu condition', *Econometrica* **62**(4), 939–948.

Acemoglu, D., Dahleh, M. & Lobel, I. (2011), 'Bayesian learning in social networks', *Review of Economic Studies* **78**(4), 1201–1236.

Acemoglu, D. & Robinson, J. A. (2009), *Economic Origins of Dictatorship and Democracy*, Cambridge University Press.

Alos-Ferrer, C. & Netzer, N. (2010), 'The logit response dynamics', *Games and Economic Behavior* **68**, 413–427.

Amaldoss, W. & Rapoport, A. (2009), Excessive expenditures in two-stage contests: Theory and experimental evidence, *in* I. N. Haugen & A. S. Nilsen, eds, 'Game Theory: Strategies, Equilibria, and Theorems', Nova Science Publishers, chapter 9, pp. 241–266.

Arkes, H. R. & Blumer, C. (1985), 'The psychology of sunk cost', *Organizational Behavior and Human Decision Processes* **35**, 124–140.

Augenblick, N. (2012), Consumer and producer behavior in the market for penny auctions: A theoretical and empirical analysis. Working Paper.

Bala, V. & Goyal, S. (2000), 'A noncooperative model of network formation', *Econometrica* **68**, 1181–1229.

Baye, M. R., Kovenock, D. & de Vries, C. G. (1996), 'The all-pay auction with complete information', *Economic Theory* **8**(2), 291–305.

Berninghaus, S., Ehrhart, K.-M. & Ott, M. (2006), 'A network experiment in continuous time: The influence of link costs', *Experimental Economics* **9**, 237–251.

Bondonio, D. (1998), 'Predictors of accuracy in perceiving informal social networks', *Social Networks* **20**, 201–220.

Byers, J. W., Mitzenmacher, M. & Zervas, G. (2010), Information asymmetries in pay-per-bid auctions, *in* 'Proceedings of the 11th ACM Conference on Electronic Commerce', pp. 1–11.

Callander, S. & Plott, C. (2005), 'Principles of network development and evolution: An experimental study', *Journal of Public Economics* **89**, 14691495.

Calv-Armengol, A. & Ilkili, R. (2009), 'Pairwise-stability and nash equilibria in network formation', *International Journal of Game Theory* **38**, 51–79.

Calv-Armengol, A. & Jackson, M. (2004), 'The effects of social networks on employment and inequality', *American Economic Review* **94**, 426–454.

Camerer, C. F., Ho, T.-H. & Chong, J.-K. (2004), 'A cognitive hierarchy model of games', *The Quarterly Journal of Economics* **119**(3), 861–898.

Cameron, C., Gelbach, J. & Miller, D. (2011), 'Robust inference with multiway clustering', *Journal of Business & Economic Statistics* **29**, 238–249.

Carpenter, J., Holmes, J. & Matthews, P. H. (2011), An introduction to 'bucket auctions' for charity. Working Paper.

Carpenter, J., Holmes, J. & Matthews, P. H. (2012), Bucket auctions for charity: A field trial. Working Paper.

Carrillo, J. & Gaduh, A. (2012), The strategic formation of networks: Experimental evidence. Working Paper.

Casciaro, T. (1998), 'Seeing things clearly: Social structure, personality, and accuracy in social network perception', *Social Networks* **20**, 331–351.

Conley, T. & Udry, C. (2001), 'Social learning through networks: The adoption of new agricultural technologies in ghana', *American Journal of Agricultural Economics* **83**, 668–673.

Davis, D. D. & Reilly, R. J. (1998), 'Do too many cooks always spoil the stew? an experimental analysis of rent-seeking and the role of a strategic buyer', *Public Choice* **95**(1-2), 89–115.

Dow, G. K. & Reed, C. G. (2013), 'The origins of inequality: Insiders, outsiders, elites, and commoners', *Journal of Political Economy* **121**, 609–641.

Easterly, W. (2001), *The Elusive Quest for Growth: Economists' Adventures and Misadventures in the Tropics*, MIT Press.

Eckel, C. C. & Grossman, P. J. (2008), 'Forecasting risk attitudes: An experimental study using actual and forecast gamble choices', *Journal of Economic Behavior and Organization* **68**, 1–17.

Falk, A. & Kosfeld, M. (2012), 'It's all about connections. evidence on network formation', *Review of Network Economics* **11**, 1–36.

Field, E. (2007), 'Entitled to work: Urban property rights and labor supply in peru', *The Quarterly Journal of Economics* **122**(4), 1561–1602.

Fischbacher, U. (2007), ‘z-tree: Zurich toolbox for read-made economic experiments’, *Experimental Economics* **10**, 171–178.

Foster, D. & Young, H. P. (1990), ‘Stochastic evolutionary game dynamics’, *Theoretical Population Biology* **38**, 219–232.

Francetich, A. & Troyan, P. (2010), A network formation game with private information. Working Paper.

Friedkin, N. (1983), ‘Horizons of observability and limits of informal control in organizations’, *Social Forces* **62**, 54–77.

Friedman, D., Pommerenke, K., Lukose, R., Milam, G. & Huberman, B. A. (2007), ‘Searching for the sunk cost fallacy’, *Experimental Economics* **10**, 79–104.

Fukuyama, F. (2011), *The Origins of Political Order: From Prehuman Times to the French Revolution*, Farrar, Straus and Giroux.

Gneezy, U. & Smorodinsky, R. (2006), ‘All-pay auctions - an experimental study’, *Journal of Economic Behavior and Organization* **61**(2), 255–275.

Goeree, J., Riedl, A. & Ule, A. (2009), ‘In search of stars: Network formation among heterogeneous agents’, *Games and Economic Behavior* **67**, 445–466.

Gonzalez, F. (2007), ‘Effective property rights, conflict and growth’, *Journal of Economic Theory* **137**, 127–139.

Goodman, J. (2012), Reputations in bidding fee auctions. Working Paper.

Gradstein, M. (1995), ‘Intensity of competition, entry and entry deterrence in rent seeking contests’, *Economics and Politics* **7**(1), 79–91.

Grandjean, G., Mauleon, A. & Vince (2011), ‘Connections among farsighted agents’, *Journal of Public Economic Theory* **13**, 935–955.

Grossman, H. & Noh, S. J. (1994), ‘Proprietary public finance and economic welfare’, *Journal of Public Economics* **53**, 187–204.

Grotius, H. (1625), *De Jure Belli ac Pacis Libri Tres*, Vol. 2 of *The Classics of International Law*, Carendon Press. Reprint, 1925.

Heath, C. (1995), ‘Escalation and de-escalation of commitment in response to sunk costs: The role of budgeting in mental accounting’, *Organizational Behavior and Human Decision Processes* **62**, 38–54.

Herings, J.-J., Mauleon, A. & Vannetelbosch, V. (2009), ‘Farsightedly stable networks’, *Games and Economic Behavior* **67**, 526–541.

Herrmann, B. & Orzen, H. (2008), The appearance of homo rivalis: Social preferences and the nature of rent seeking. Working Paper.

Hinnosaar, T. (2010), Penny auctions are unpredictable. Working Paper.

Hrisch, H. & Kirchkamp, O. (2010), 'Less fighting than expected: Experiments with wars of attrition and all-pay auctions', *Public Choice* **144**(1-2), 347–367.

Hume, D. (1739), *A Treatise of Human Nature*, Oxford University Press.

Jackson, M. (2010), *Social and Economic Networks*, Princeton University Press.

Jackson, M. & Watts, A. (2002), 'The evolution of social and economic networks', *Journal of Economic Theory* **106**, 265–295.

Jackson, M. & Wolinsky, A. (1996), 'A strategic model of social and economic networks', *Journal of Economic Theory* **71**, 44–74.

Kahneman, D. & Tversky, A. (1979), 'Prospect theory: An analysis of decision under risk', *Econometrica* **47**(2), 263–292.

Kahneman, D. & Tversky, A. (1991), 'Loss aversion in riskless choice: A reference- dependent model', *The Quarterly Journal of Economics* **106**(4), 1039–1061.

Kahneman, D. & Tversky, A. (1992), 'Advances in prospect theory: Cumulative representation of uncertainty', *Journal of Risk and Uncertainty* **5**(4), 297–323.

Kimbrough, E., Smith, V. & Wilson, B. (2010), 'Exchange, theft, and the social formation of property', *Journal of Economic Behavior and Organization* **74**(3), 206–229.

Konrad, K. & Skaperdas, S. (2012), 'The market for protection and the origin of the state', *Economic Theory* **50**(2), 417–443.

Kranton, R. & Minehart, D. (2001), 'A theory of buyer-seller networks', *American Economic Review* **91**, 485–508.

Krishna, V. & Morgan, J. (1997), 'An analysis of the war of attrition and the all-pay auction', *Journal of Economic Theory* **72**(2), 343–362.

Ku, G. (2008*a*), 'Before escalation: Behavioral and affective forecasting in escalation of commitment', *Personality and Social Psychology Bulletin* **34**(11), 1477–1491.

Ku, G. (2008*b*), 'Learning to de-escalate: The effects of regret in escalation of commitment', *Organizational Behavior and Human Decision Processes* **105**(2), 221–232.

Kumbasar, E., Romney, K. & Batchelder, W. (1994), 'Systematic biases in social perception', *American Sociological Review* **100**, 477–505.

Lam, T. (2011), Pay-per-bid auctions with an exit option: An experimental and empirical investigation, Master's thesis, The University of New South Wales.

Laumann, E. (1969), 'Friends of urban men: An assessment of accuracy in reporting their socioeconomic attributes, moral choice, and attitude agreement', *Sociometry* **32**, 54–69.

Locke, J. (1689), *Two Treatises of Government*, Edes and Gill. Reprint, 1773.

Mantovani, M., Kirchsteiger, G., Mauleon, A. & Vannetelbosch, V. (2013), Limited farsightedness in network formation. Working Paper.

Maskin, E. & Tirole, J. (2001), 'Markov perfect equilibrium: I, observable actions', *Journal of Economic Theory* **100**(2), 191–219.

Mayshar, J., Moav, O. & Neeman, Z. (2012), Transparency, appropriability and the early state. Working Paper.

McBride, M. (2006*a*), 'Imperfect monitoring in communication networks', *Journal of Economic Theory* **126**, 97–119.

McBride, M. (2006*b*), 'Limited observation in mutual consent networks', *Advances in Theoretical Economics* **6**, 1–29.

McBride, M. (2008), 'Position-specific information in social networks: Are you connected?', *Mathematical Social Sciences* **56**, 283–295.

McGuire, M. C. & Olson, M. (1996), 'The economics of autocracy and majority rule: The invisible hand and the use of force', *Journal of Economic Literature* **34**, 72–96.

Moselle, B. & Polak, B. (2001), 'A model of a predatory state', *The Journal of Law, Economics, & Organization* **17**(1), 1–33.

Myerson, R. (2008), 'The autocrat's credibility problem and foundations of the constitutional state', *American Political Science Review* **102**, 125–139.

North, D. C., Wallis, J. J. & Weingast, B. R. (2009), *Violence and Social Orders: A Conceptual Framework for Interpreting Recorded Human Hist*, Cambridge University Press.

Nozick, R. (1977), *Anarchy, State, and Utopia*, Basic Books, Inc.

Olson, M. (1993), 'Dictatorship, democracy, and development', *American Political Science Review* **87**(3), 567–576.

Olson, M. (2000), *Power and Prosperity: Outgrowing Communist and Capitalist Dictatorships*, Oxford University Press.

Oppenheimer (1908), *The State*, Fox and Wilkes. Reprint, 1997.

Pantz, K. (2006), The Strategic Formation of Social Networks: Experimental Evidence, PhD thesis, Humbolt Universitt.

Peterson, M. (2009), 'Estimating standard errors in finance panel data sets: Comparing approaches', *Review of Financial Studies* **22**, 435–480.

Platt, B. C., Price, J. & Tappen, H. (2013), 'The role of risk preferences in pay-to-bid auctions', *Management Science* .

Price, C. R. & Sheremeta, R. (2011), 'Endowment effects in contests', *Economics Letters* **111**(3), 217–219.

Pufendorf, S. (1672), *De Jure Naturae et Gentium Libri Octo*, Vol. 2 of *Classics of International Law*, The Clarendon Press. Reprint, 1934.

Reinikka, R. & Svensson, J. (2004), 'Local capture: Evidence from a central government transfer program in uganda', *The Quarterly Journal of Economics* **119**(2), 679–705.

Robinson, J. (2001), When is state predatory. Working Paper.

Rong, R. & Houser, D. (2013), Emergent star networks with ex ante homogeneous agents. Working Paper.

Seabright, P. (2008), Warfare and the multiple adoption of agriculture after the last ice age. Working Paper.

Sheremeta, R. (2010), 'Experimental comparison of multi-stage and one-stage contests', *Games and Economic Behavior* **68**(2), 731–747.

Sheremeta, R. (2011), 'Contest design: An experimental investigation', *Economic Inquiry* **49**(2), 573–590.

Shubik, M. (1971), 'The dollar auction game: a paradox in noncooperative behavior and escalation', *Journal of Conflict Resolution* **15**, 109–111.

Simon, H. (1976), *Administrative Behavior*, Cambridge University Press.

Skaperdas, S. (2014), *Coercion and Social Welfare in Public Finance: Economic and Political Dimensions*, Cambridge University Press, chapter Proprietary Public Finance: On its Emergence and Evolution Out of Anarchy. forthcoming.

Smith, J. M. (1974), 'Theory of games and evolution of animal conflicts', *Journal of Theoretical Biology* **47**, 209–221.

Song, Y. & van der Schaar, M. (2013), Dynamic network formation with incomplete information, Technical report, University of California Los Angeles.

Staw, B. M. (1981), 'The escalation of commitment to a course of action', *Academy of Management Review* **6**(4), 577–587.

Stein, W. (2002), 'Asymmetric rent-seeking with more than two contestants', *Public Choice* **113**, 325–336.

Sugden, R. (2004), *The Economics of Rights, Cooperation and Welfare*, 2nd edn, Palgrave Macmillan.

Thaler, R. H. (1980), 'Toward a positive theory of consumer choice', *Journal of Economic Behavior & Organization* **1**, 39–60.

Thompson, S. (2011), 'Simple formulas for standard errors that cluster by both firm and time', *Journal of Financial Economics* **99**, 1–10.

Tsai, M.-H. & Young, M. (2010), 'Anger, fear, and escalation of commitment', *Cognition & Emotion* **24**(6), 962–973.

van Leeuwen, B., Offerman, T. & Schram, A. (2013), Superstars need social benefits: An experiment on network formation. Working Paper.

Wang, Z. & Xu, M. (2013), Selling a dollar for more than a dollar? evidence from online penny auctions. Working Paper.

Watts, A. (2001), 'A dynamic model of network formation', *Games and Economic Behavior* **34**, 331–341.

Wilson, B., Jaworski, T., Schurter, K. & Smyth, A. (2012), 'The ecological and civil mainsprings of property: An experimental economic history of whalers' rules of capture', *Journal of Law, Economics, & Organization* **28**(4), 617–656.

Wong, K.-F. E. & Kwong, J. (2007), 'The role of anticipated regret in escalation of commitment', *Journal of Applied Psychology* **92**(2), 545–554.

Wong, K.-F. E., Kwong, J. & Ng, C. (2008), 'When thinking rationally increases biases: The role of rational thinking style in escalation of commitment', *Applied Psychology* **57**(2), 246–271.

Wong, K.-F. E., Yik, M. & Kwong, J. (2006), 'Understanding the emotional aspects of escalation of commitment: The role of negative affect', *Journal of Applied Psychology* **91**(2), 282–297.

Young, H. P. (1998), *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*, Princeton University Press.

Appendix A

Supplement to Chapter 2

A.1 Screenshots

Figure A.1: Screenshots

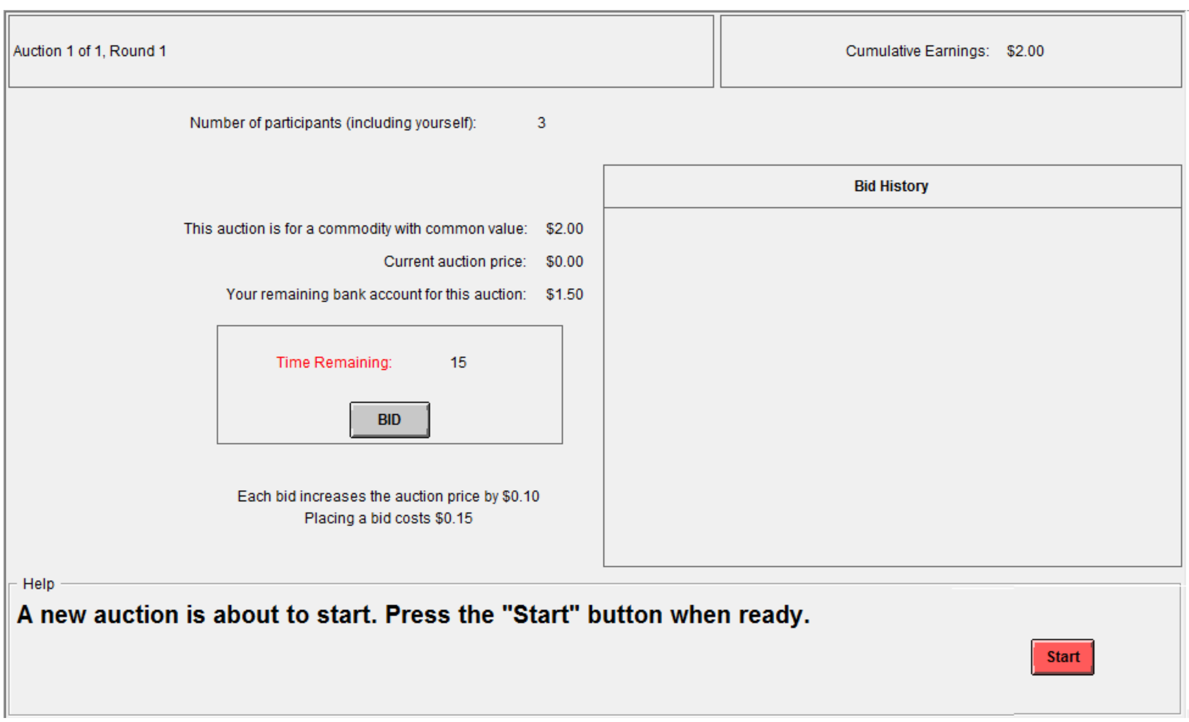

(a) Start Screen

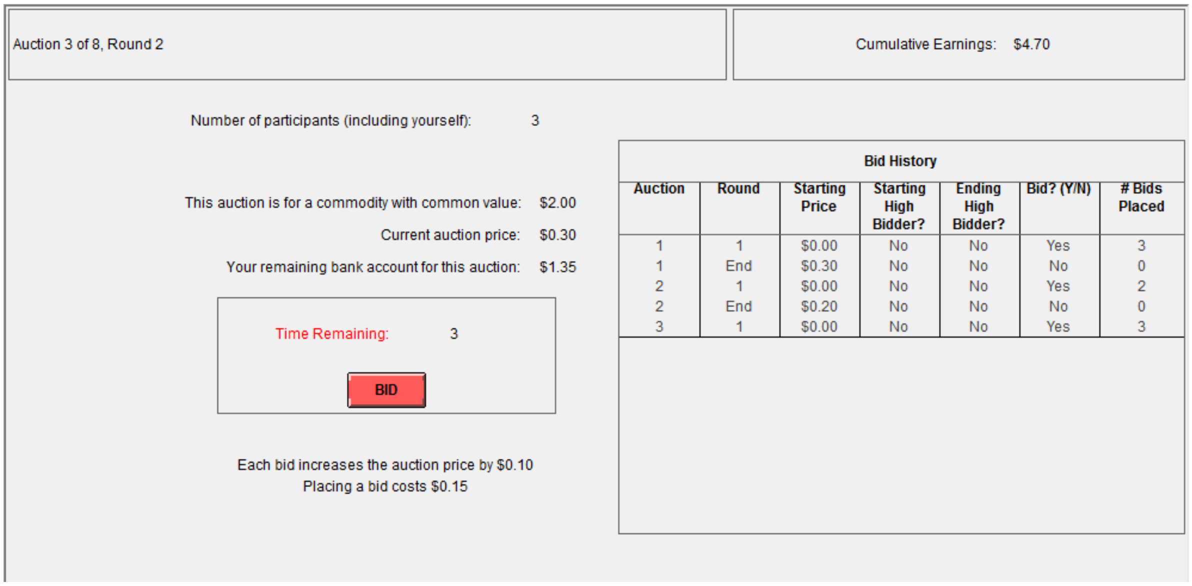

(b) Decision Screen

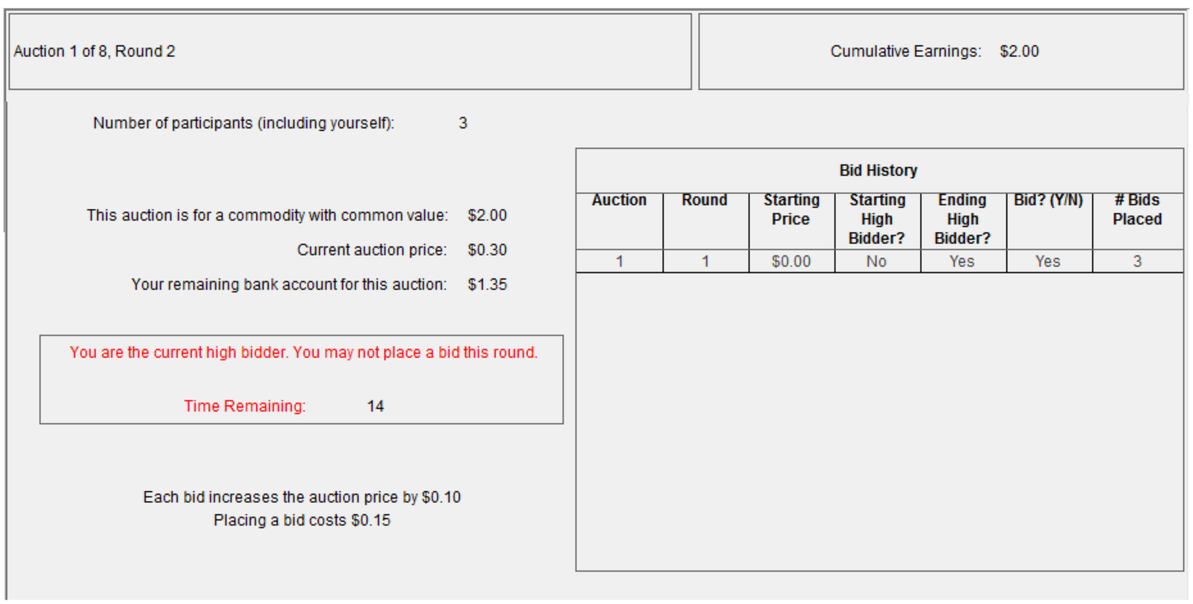

(c) High Bidder

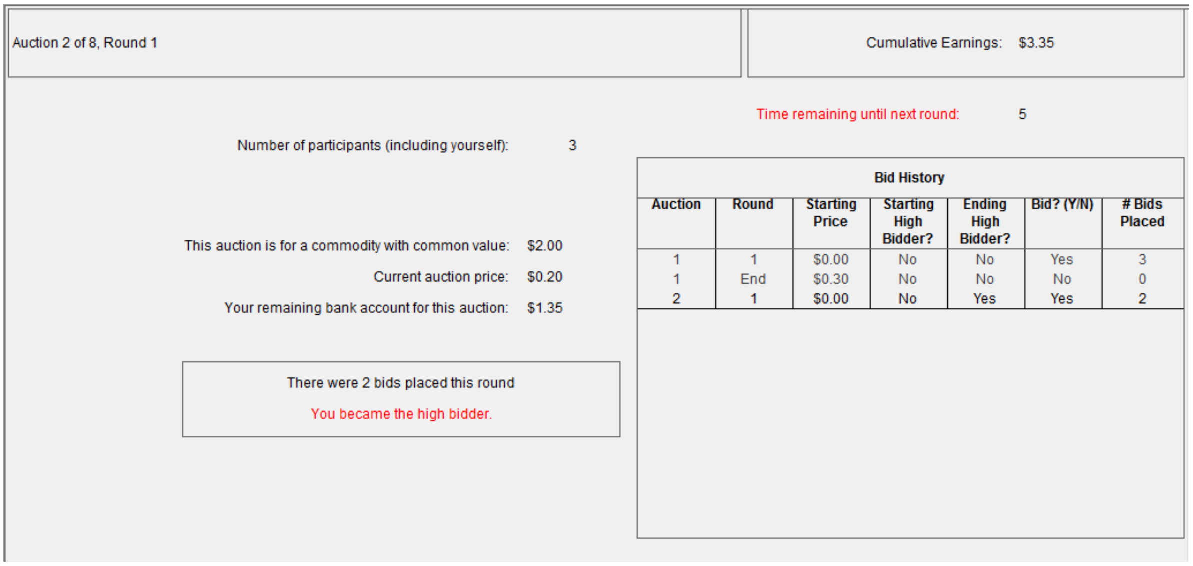

(d) Results Screen

A.2 Instructions

SCREEN 1:

Welcome to this experiment at UC-Irvine. Thank you for participating.

You are about to participate in a study of decision-making, and you will be paid for your participation in cash, privately at the end of this session. What you earn depends partly on your decisions and partly on chance.

Please turn off your cell phone.

The entire session consists of multiple practice rounds and real rounds. You will be paid for the real rounds only. The practice rounds are only for you to familiarize yourself with the decision-making environment.

All rounds will take place through the computer terminals. It is important that you do not communicate with any other participants during the session.

When you are ready, please click continue to go to the instructions.

SCREEN 2:

Over the course of this session, you will be participating in a series of auctions along with several other participants. Participants are randomly matched at the start of each auction, so the other participants may not be the same across auctions. There will be 8 auctions and 1 practice auction in total.

In each auction, you have the opportunity to bid on a commodity with known common value. That is, you and all other participants know the value, and this value is the same for all participants. The common value and number of participants will be displayed on your screen throughout the auction.

At the end of an auction, the high bidder immediately receives the dollar value of the commodity and then pays the end auction price. The resulting profit for the high bidder is given by the following formula:

Profit = Common Value - End Auction Price

Any profits you earn will be added to your bank account and any costs you incur will be taken from it. The starting balance in your bank account for each auction is $1.50. The balance in your bank account at the end of an auction does NOT carry over to future auctions. Instead, at the end of an auction you will earn $1.00 for each $1.00 in your bank account.

SCREEN 3:

A sample decision screen is shown below. The countdown timer near the "BID" button would normally be counting down to zero, and the screen would end when it reaches zero. In this example the common value of the commodity is $2.00 and there are 5 participants (including yourself).

Please press the "BID" button to continue this example.

SCREEN 4:

A sample results screen is shown below. The countdown timer to the top right of the screen would normally be counting down to zero, and the screen would end when it reaches zero. In this example you placed a bid and became the current high bidder. This information is displayed in the center box on the screen. Note how the current auction price has increased by $0.10 and your bank account has decreased by $0.15 as described previously.

Please press the "Continue" button to continue this example.

SCREEN 5:

It is the decision screen once again, but now you are the high bidder. Since the high bidder does not participate in the decision screen, the "BID" button is not available to you. If bids are placed in this round, the auction will continue and you will participate in next round's decision screen. If no bids are placed, the auction will end and you will immediately receive the dollar value of the commodity and then pay the current auction price.

Please press the "Continue" button to proceed.

SCREEN 6:

Click continue to participate in 1 practice auction. After completing the practice auction, you will participate in 8 real auctions.

You will begin the session with $2.00 in earnings. During the session, any gains you make will be added to your earnings and any losses you incur will be deducted from your earnings.

A.3 Theory

A.3.1 Equilibrium Concept

To simplify notation, define $c = C/\varepsilon$, $p_t = (P_t - P_0)/\varepsilon$, and $\pi = (\Pi - P_0)/\varepsilon$. Let m denote the number of non-leading participants, $q(p_t)$ denote the probability of a non-leader bidding at price p_t, $V(p_t)$ denote the continuation value of the auction for a non-leader at price p_t, and $V^*(p_t)$ denote the continuation value of the auction for the high bidder at price p_t. Then the risk neutral symmetric Markov perfect equilibrium is recursively characterized by (A.1) and (A.2).

$$V(p_t) = \max \left\{ \begin{array}{c} \sum_{k=0}^{m-1} \left[\binom{m-1}{k} q(p_t)^k (1-q(p_t))^{m-1-k} \left(\frac{V^*(p_t+k+1)+kV(p_t+k+1)}{k+1} \right) \right] - c, \\ \sum_{k=1}^{m-1} \left[\binom{m-1}{k} q(p_t)^k (1-q(p_t))^{m-1-k} V(p_t+k) \right] \end{array} \right\}, \quad \text{(A.1)}$$

$$V^*(p_t) = (1-q(p_t))^m (\pi - p_t) + \sum_{k=1}^{m} \left[\binom{m}{k} q(p_t)^k (1-q(p_t))^{m-k} V(p_t+k) \right] \quad \text{(A.2)}$$

Denote the first and second arguments of the max function by A and B respectively. Argument A is the expected continuation value of bidding with certainty given the bid probabilities of others. Argument B is the expected continuation value of not bidding given the bid probabilities of others. For $q(p_t)$ to be a part of the equilibrium bidding strategy it must satisfy one of 3 conditions:

1. $q(p_t) = 1$ and $A \geq B$

2. $q(p_t) = 0$ and $A \leq B$

3. $q(p_t) \in (0, 1)$ and $A = B$

This equilibrium is not in general unique and so there may be multiple $q(p_t)$ which satisfy one of these conditions for a given p_t. By utilizing the fact that $V(p_t) = 0$ for $\pi - p_t < 0$ this system can be solved numerically using backwards induction.

A.3.2 Proof of Theorem 2.1

Proof of Theorem 2.1 (i):

Start with the following theorem from Hinnosaar (2010):

Theorem A.1. *(Hinnosaar 2010) Let $E(R)$ denote the expected auction revenue conditional on sale in the symmetric risk neutral equilibrium characterized by (A.1) and (A.2). Then $E(R)$ has the following properties:*

i $E(R) \leq \Pi$

ii *if* $q(p) < 1, \forall p$, *then* $E(R) = \Pi$

iii *for some values of the auction parameters* $E(R) < \Pi$

The proof of Theorem A.1 does not require the symmetry assumption and instead makes use of the fact that $V(p) = 0$ for $q(p) \in [0, 1)$. This leads to the following generalization of Hinnosaar (2010):

Corollary A.1. *In any asymmetric risk neutral equilibrium characterized by (A.1) and (A.2), the expected auction revenue conditional on sale $E(R)$ has the following properties:*

i $E(R) \leq \Pi$

ii if $q_i(p) < 1, \forall i, p$, then $E(R) = \Pi$

iii for some values of the auction parameters $E(R) < \Pi$

Proof. Any strategy combination $q_i(p) < 1, \forall p$ assigns a positive probability to participant i never bidding. This strategy will yield 0 with certainty. Because a participant will only use a mixed strategy when she indifferent between her strategy choices, when $q_i(p) \in (0, 1)$ this strategy must yield 0 in expectation. Thus, $V_i(0) = 0$. When $q_i(p) < 1, \forall i, p$, $V_i(0) = 0, \forall i$. This implies that $E(R) = \Pi$ when $q_i(p) < 1, \forall i, p$. Because $q_i(p) = 1$ only when $V_i(p) > 0$, it follows that $E(R) \leq \Pi$ and that $E(R) < \Pi$ in equilibria which involve a participant bidding with certainty. □

Theorem 2.1*(i)* follows from Corollary A.1 *(i)*.

Proof of Theorem 2.1 (ii):

When restricting attention to Markov perfect equilibria, the condition $E(R) = \Pi$ cannot in general be satisfied. Consider the case where $m = 2$, the symmetric bid function $q(p_t) = 1$ for some t, and suppose that $q(p_t) = 1$ is the only symmetric strategy that satisfies conditions 1-3 above. Suppose the symmetry restriction is dropped, and that participant i deviates to $q_i(p_t) \neq 1$. Participant $-i$'s best response is still $q_{-i}(p_t) = 1$ as $V_{-i}(p_t) > 0$ for all strategies $q_i(p_t)$. But by similar logic, when $q_{-i}(p_t) = 1$, participant i's best response is $q_i(p_t) = 1$. Thus no symmetric or asymmetric strategy exists in this case such that $q_i(p) < 1, \forall i, p$.

When the Markov perfect and subgame perfect restrictions are removed, it is easy to see that equilibria exist such that $E(R) = \Pi$ for an arbitrary (n, π, c, ε). Pick any strategy set s such that the expected continuation value for all players at time zero equals zero and such that $E(R) = \Pi$. Now define the strategy s_i^* which is: play strategy s_i until some player

$j \neq i$ deviates from s_j and then bid with probability 1 until $\pi - p_t \leq 0$. The strategy set s^* constitutes a Nash equilibrium where $E(R) = \Pi$.

A.4 Supplemental Figures

Table A.1: Primary Hypothesis Tests

	Statistic	Number of Participants: 3	Number of Participants: 5
(1) H_1: Revenue ≤ 2			
H_a: Revenue > 2			
Entire Session			
Discrete	Student t (p-value)	**0.140**	**0.000**
	[df]	7	15
Continuous	Student t (p-value)	**0.046**	**0.000**
	[df]	7	15
Predicted Revenue: Last Period			
Discrete	Student t (p-value)	**1.0000**	**0.011** | **0.093**
	[df]	7	7 | 7
Continuous	Student t (p-value)	**0.1557**	**0.087** | **0.057**
	[df]	7	7 | 19
(2) H_2: Revenue Discrete = Revenue Continuous			
H_a: Revenue Discrete ≠ Revenue Continuous			
	Student t (p-value)	**0.680**	**0.507**
	[df]	15	43
(3) H_0: Bids/Auction Risk Averse = Bids/Auction Risk Neutral = Bids/Auction Risk Seeking			
H_3: Unequal Populations			
Joint Test	Kruskal-Wallis (p-value)	**0.022**	**0.428**
	[df]	2	2
H_{0a}: Risk Averse ≥ Risk Neutral	Mann-Whitney U (p-value)	**0.004**	**n/a**
H_{3a}: Risk Averse < Risk Neutral	[df]	1	
H_{0b}: Risk Neutral ≥ Risk Seeking	Mann-Whitney U (p-value)	**0.936**	**n/a**
H_{3b}: Risk Neutral < Risk Seeking	[df]	1	

Notes: Bold denotes hyptheses that are predicted to be true. All t-tests are computed using clustered standard errors. Session 1 and Session 2 results are seperated by "|" when applicable.

Table A.2: Additional Hypothesis Tests

H_0: Revenue Session 1 = Revenue Session 2			
H_a: Revenue Session 1 ≠ Revenue Session 2			
Discrete	Student t (p-value)	**0.838**	
	[df]	15	
Continuous	Student t (p-value)	**0.782**	
	[df]	15	
H_0: Revenue 3 Participants = Revenue 5 Participants			
H_a: Revenue 3 Participants ≠ Revenue 5 Participants			
Discrete	Student t (p-value)	**0.007**	
	[df]	23	
Continuous	Student t (p-value)	**0.632**	
	[df]	35	
H_0: Revenue Trend = 0			
H_a: Revenue Trend ≠ 0			
Discrete3	Student t (p-value)	**0.014**	
	[df]	7	
Discrete5	Student t (p-value)	**0.006**	**0.588**
	[df]	7	7
Continuous3	Student t (p-value)	**0.071**	
	[df]	7	
Continuous5	Student t (p-value)	**0.039**	**0.078**
	[df]	7	19

Notes: **Bold** denotes hyptheses that are predicted to be true. All t-tests are computed using clustered standard errors. Session 1 and Session 2 results are seperated by "|" when applicable.

Table A.3: Average Unique Visitors Per Month (in Thousands)

Year	Pay-to-Bid Auction Website										Industry	
	Quibids	Beezid	Swoopo	Dealdash	Pennygrab	Orangebidz	Madbid	Bidcactus	Happybidday	Other	Total	Ebay
2010	0.0	13.8	30.9	1.4	0.0	0.0	11.6	0.0	0.0	6.8	64.4	71953.3
2011	19.0	71.0	22.4	10.8	0.4	2.5	1.2	3.5	0.0	15.2	146.0	51663.5
2012	688.7	324.2	0.0	18.2	5.5	14.1	0.8	15.4	6.9	19.9	1093.6	89552.0

Note: Data collected from SEMrush.com in April 2013. Unique visitors = google organic traffic/(% organic traffic from google * % organic traffic). Percentages were interpolated from the industry average when missing. "Other" category consists of Bidstick, Bidray, Bidhere, Bidfun, Bigdeal, Haggle, Skoreit, Oohilove, For10cents, and Zerobid.

Appendix B

Supplement to Chapter 3

Table B.1: Trend in Group Outcomes (Match)

Treatment	Dependent Variable		
	Convergence (%)	Efficiency (%)	Cycles (%)
Neighbor Observation - Low Cost			
Average Marginal Effect	5.5	-0.6	-5.6
Standard Error	(1.9)	(0.8)	(1.8)
p-value	0.005	0.437	0.002
Neighbor Observation - High Cost			
Average Marginal Effect	3.2	-0.7	-8.2
Standard Error	(1.8)	(0.7)	(1.1)
p-value	0.075	0.308	< 0.001
Full Observation - Low Cost			
Average Marginal Effect	2.8	0.9	-8.6
Standard Error	(2.5)	(0.9)	(1.1)
p-value	0.271	0.311	< 0.001
Full Observation - High Cost			
Average Marginal Effect	5.0	3.7	-4.2
Standard Error	(2.2)	(0.8)	(2.2)
p-value	0.026	< 0.001	0.055
Hypothesis Test	Wald statistic (p-value)		
H_L: Trends Jointly Equal H_{La}: Different Trends	0.756	< 0.001	0.189

Note: Bold denotes hypothesis predicted to be true.

Table B.2: Paired Comparisons: Trend in Group Outcomes (Match)

Comparison	Bonferroni Adjusted p-value
$\text{Efficiency}_{\text{NO-LC}} = \text{Efficiency}_{\text{NO-HC}}$	> 0.999
$\text{Efficiency}_{\text{NO-LC}} = \text{Efficiency}_{\text{FO-LC}}$	> 0.999
$\text{Efficiency}_{\text{NO-LC}} = \text{Efficiency}_{\text{FO-HC}}$	< 0.001
$\text{Efficiency}_{\text{NO-HC}} = \text{Efficiency}_{\text{FO-LC}}$	0.925
$\text{Efficiency}_{\text{NO-HC}} = \text{Efficiency}_{\text{FO-HC}}$	< 0.001
$\text{Efficiency}_{\text{FO-LC}} = \text{Efficiency}_{\text{FO-HC}}$	0.141

Table B.3: Primary Hypothesis Tests

Hypothesis Test	Statistic (p-value)		
H_1: $\text{Convergence}_{\text{LC}} \leq \text{Convergence}_{\text{HC}}$		Observation Level	
H_{1a}: $\text{Convergence}_{\text{LC}} > \text{Convergence}_{\text{HC}}$		Neighbor	Full
	Barnard's Exact	0.002	0.938
Logistic Regression: Group Convergence	t-test	0.001	0.939
H_2: $\text{Convergence}_{\text{NO}} \leq \text{Convergence}_{\text{FO}}$		Tie Cost	
H_{2a}: $\text{Convergence}_{\text{NO}} > \text{Convergence}_{\text{FO}}$		Low	High
	Barnard's Exact	0.279	> 0.999
Logistic Regression: Group Convergence	t-test	0.264	> 0.999
H_3: $\text{Efficiency}_{\text{LC}} \leq \text{Efficiency}_{\text{HC}}$		Observation Level	
H_{3a}: $\text{Efficiency}_{\text{LC}} > \text{Efficiency}_{\text{HC}}$		Neighbor	Full*
	Mann-Whitney U	< 0.001	0.008
Regression: Individual Share of Group Efficiency	t-test	< 0.001	< 0.001
H_4: $\text{Cycles}_{\text{NO}} \leq \text{Cycles}_{\text{FO}}$		Tie Cost	
H_{4a}: $\text{Cycles}_{\text{NO}} > \text{Cycles}_{\text{FO}}$		Low	High
	Barnard's Exact	0.001	< 0.001
Logistic Regression: Group Cycles	t-test	< 0.001	< 0.001
H_5: $\text{Drop Stem}_{\text{NO}} \geq \text{Drop Stem}_{\text{FO}}$		Tie Cost	
H_{5a}: $\text{Drop Stem}_{\text{NO}} < \text{Drop Stem}_{\text{FO}}$		Low	High
Regression: Dropped Stems (Match Proportion)	t-test	0.666	0.914
H_6: $\text{Drop Stem}_{\text{LC}} \geq \text{Drop Stem}_{\text{HC}}$		Observation Level	
H_{6a}: $\text{Drop Stem}_{\text{LC}} < \text{Drop Stem}_{\text{HC}}$		Neighbor	Full
Regression: Dropped Stems (Match Proportion)	t-test	0.001	0.014
H_7: Time as Stem Equal by Risk Category		Pooled Data	
H_{7a}: Different Proportions			
Regression: Time as Stem (Match Proportion)	F-test	0.456	

Note: Bold denotes hypotheses predicted to be true. All regressions utilize subject level data with standard errors two-way clustered at the group and subject level. Regressions also include controls for the match number.

*Excludes 9 networks containing a subject who never connected. Including these networks changes p-values to 0.017 and 0.003. The treatment of this subject does not qualitatively affect the result of any hypothesis test.

B.1 Instructions

Text enclosed in square brackets, e.g., [string1/string2], depends on the treatment.

SCREEN 1:

Welcome to this experiment at UC Irvine. Thank you for participating.

You are about to participate in a study of decision-making, and you will be paid for your participation in cash, privately at the end of this session. What you earn depends partly on your decisions and partly on chance.

Please turn off your cell phone.

The first part of this experiment consists of 10 matches with 15 rounds per match. You will be paid according to the outcome of *each* round and you will receive the sum of all your individual round earnings, in addition to the show up fee and your earnings in the second part, at the end of the experiment. You will receive further instruction when the second part of the experiment begins.

All rounds will take place through the computer terminals. It is important that you do not talk with any other participants during the session.

When you are ready, please click "Continue" to go to the instructions.

SCREEN 2:

In each match, you will be randomly grouped with 11 other participants. The participants that you are grouped with will stay the same for the 15 rounds in this match, but you will be randomly grouped with a new set of participants at the start of the next match. Each participant will be assigned a color and will retain this color throughout the match.

During each match, you and the other participants will form a network. In each round, you will be presented with a [partial view of the network/complete view of the network]. At the start of the first round, the network has no connections. In all following rounds, you will be shown [part/all] of the network formed at the end of the prior period.

The network consists of 12 colored circles, each representing a participant in your group, which may or may not be connected via lines (we refer to these lines as "links"). Participants are either directly linked, denoted by a line connecting two participants, or not directly linked, denoted by the absence of a line.

Please press "Continue" to proceed.

SCREEN 3:

After observing [part/all] of the network, each participant is then randomly paired with another participant in the network and given the opportunity to change his or her link status with that participant. If you already have a link with that participant, you can maintain that link or remove that link. If you do not have a link with that participant, you can initiate a link to that participant or not initiate a link. Your link status with all other participants stays the same. Links that have already been initiated are maintained.

NOTE: Links are initiated and maintained bilaterally, meaning that a link is initiated or maintained only if BOTH participants choose to do so. If one of the participants does not want the link then it will not form.

You will have 15 seconds per round to make your decision, and you must press the "Continue" button before time runs out for your decision to be recorded. If you do not press "Continue," then the computer will assume you do not want to make any changes to the network that formed last period. Please press "Continue" to proceed.

SCREEN 4:

We say that two participants are “indirectly linked” if these participants are not directly linked but are connected through a series of links with other participants.

In each round you will [only see your own links and the links of participants that you are directly linked to/be shown the links of every participant].

Your payoff for each round will be 1 point for every participant you are either directly or indirectly linked to minus [0.7/1.5] points for every participant that you have attempted to initiate or maintain a link to. Note that you will pay this cost *even if your link attempt is not successful.*

EXAMPLE 1: If A and B are both directly linked to each other and no-one else, each receives 1 point for being linked to one other participant and pays a cost of [0.7/1.5] for having one direct link. The resulting payoff is [0.3/-0.5].

EXAMPLE 2: If A and B are both directly linked to C and no-one else, then A and B both receive 2 points for being linked to two other participants and pay a cost of [0.7/1.5] for having one direct link. The resulting payoff is [1.3/0.5].

NOTE: You will receive a payoff and pay the cost of all your direct links (including those links to participants with whom you were not paired) *in every round.* [You receive 1 point for every participant you are linked to even if you cannot see how you are linked to that participant./] Please press “Continue” to view an example.

SCREEN 5:

Below is a sample network. You are red and have exactly one link. This link is to yellow, and you can see that yellow is linked to lime. [There may be other links in the network that you cannot see./] At the bottom of the screen, there is a box around yellow denoting that in this round you must choose whether to change your link status with yellow. If you were paired with blue, then the box would be around blue.

In this example, you can maintain the link with yellow or remove it. Please select "No" to remove the link and then press "Continue."

SCREEN 6:

You are no longer linked to yellow because you chose not to maintain this link. [Assuming no other participant has changed their link status and that there were no other links you couldn't see in the network, the new network is shown to the left below. Your view of the new network is shown to the right below. Notice that you no longer see the link between yellow and lime because you are no longer linked to yellow./Assuming no other participant has changed their link status, the new network is shown below.]

Please press "Continue" to proceed.

SCREEN 7:

This is an example of what your decision screen will look like at the start of the first round. In this example, you are red and you are paired with yellow. You will ordinarily have 15 seconds to make a decision on this screen.

SCREEN 8:

The first of the 10 matches will now begin. Each match has 15 rounds and you will be paid for *each* round. **You will be paid $0.02 for each point you earn during the session.** You will also receive [$2.00/$4.50] at the beginning of the first match.

REMINDER: You will only have 15 seconds each round on the decision screen. You must press "Continue" before time runs out for your link decision to be recorded.

You will receive 1 point for each participant you are directly or indirectly linked to and you will pay a cost of [0.7/1.5] points for each link attempt you initiate or maintain. You pay the cost to maintain a link even if you were not able to change the status of that link in this round. You do not pay the cost for link attempts that failed in a prior round.

You will start with no links in the first round of each match. In future rounds, you will start with the links you had in the prior round.

In each round you will [only see your own links and your neighbors' links, but you will still receive points if you are indirectly linked to a participant that you cannot see./be shown the links of every participant, regardless of whether or not you are linked to that participant.]

Please press "Continue" to start the experiment.

Appendix C

Supplement to Chapter 4

C.1 Implication of Asymmetric Violence Capacities for Natural Law

In Section 4.1 I claimed that natural law theory was consistent with the idea that concessions might need to be made to powerful individuals when entering into the social contract. Here I provide several quotations (emphasis added) from prominent natural law theorists to support this claim.

> As soon as community ownership was abandoned, and as yet no division had been made, it is to be supposed that *all agreed, that whatever each had taken possession of should be his property.* – Grotius

> There was need of an external act or seizure, and for this to produce a moral effect, that is, an obligation on the part of others to refrain from a thing already seized by someone else, *an antecedent pact was required* and an express pact, indeed,

> *when several men divided among themselves things open to all.* – Pufendorf

> The possession of all external goods is changeable and uncertain; which is one of the most considerable impediments to the establishment of society, and is the reason why, *by universal agreement, express or tacit, men restrain themselves* by what we now call the rules of justice and equity. – Hume

> Since Gold and Silver, being little useful to the Life of Man in proportion to Food, Rayment, and Carriage, has its value only from the consent of Men, whereof Labour yet makes, in great part, the measure,*it is plain, that Men have agreed to disproportionate and unequal Possession of the Earth.* – Locke

As the above quotes illustrate, universal consent to respect property is pre-supposed by natural law theorists, but, as is demonstrated in Sections 4.2 and 4.3, it may not be individually rational for a relatively strong agent in an asymmetric society to agree to these terms unconditionally. Namely, if the strong agent can receive a higher payoff under anarchy than she would in a cooperative pact to respect property, we would not expect her to consent to this pact. Concessions might need to be made in the interest of establishing order that lead to *ex post* asymmetries between payoffs. Thus, Locke's workmanship model of property, whereby each individual keeps the fruits of her own labor, should only apply *after* individual rationality constraints have been met. To presuppose that the workmanship model alone is sufficient to establish order, as does Nozick, is to implicitly assume that relative strengths are approximately equal *ex ante.*