Learning to Respect Property[*]

January 2010

Erik O. Kimbrough[‡]

JEL Classifications: B41, C63, C92, D51, D83, F10

Keywords: property rights, conventions, experimental economics, agent-based modeling

Abstract: Agent-based simulations and human-subject experiments explore the emergence of respect for property in a specialization and exchange economy with costless theft. Software agents, driven by reciprocity and hill-climbing heuristics and parameterized to replicate humans when property is exogenously protected, are employed to predict human behavior when property can be freely appropriated. Agents do not predict human behavior in a new set of experiments because subjects innovate, constructing a property convention of "mutual taking" allowing exchange to crowd out theft. When this convention is imposed on agents, they again replicate human behavior. Property emerges as a social convention that exploits the capacity for reciprocity to sustain trade.

---

[*] Data, simulation code, and experiment instructions available upon request. Thanks to the International Foundation for Research in Experimental Economics, the Mercatus Center at George Mason University, and the Institute for Humane Studies for generous financial support. Bart Wilson, Taylor Jaworski, and participants at conferences hosted by the Institute for Humane Studies offered helpful comments and suggestions. Thanks also to Jeffrey Kirchner for his exemplary programming of the experiments and to Jennifer Cunningham for assistance in the laboratory. All errors are my own.

[‡] Interdisciplinary Center for Economic Science, George Mason University; Economic Science Institute, Chapman University. Email: *ekimbrough@gmail.com*

I. Introduction

Property is the glue that holds together an economy based on exchange. For two agents to engage in mutually beneficial trade, it is assumed that each has a right to determine the fate of the goods and services being exchanged. Parties to an exchange rely on the presumption that their counterpart will neither renege on a partially completed transaction nor obviate the need for exchange by forcibly taking the desired goods. Such observations, while perhaps obvious, can be considered trivial only if one does not reflect on the fragility of property and the ease with which it can be violated by sufficiently motivated groups or individuals. The range of violations stretches from the historical predations of Viking marauders to lowly pickpockets in crowded subway stations, and property has always been precarious to a greater or lesser degree. How, then, does the respect for property emerge in order to facilitate specialization and trade? This paper seeks to answer that question with a combination of human-subject experiments and agent-based models. When subjects (or computerized agents) are placed in an environment with imperfect property protection, how do they come to respect one another's possessions in order to reap the gains from trade?

Kimbrough (2009) develops an agent-based model of the two-good production and exchange economy of Crockett, Smith and Wilson (2009, hereafter CSW) to make concrete the dynamics and behavioral rules that guide subjects to discover (or fail to discover) specialization and trade. Hill-climbing and reciprocity heuristics are consistent with the varieties of human behavior in the original CSW environment and predict behavior under environmental variations with exogenous property enforcement. As Kimbrough argues, if agents' decision rules yield model output that is accurate in its depiction of human behavior in one environment, a valid model should predict human behavior in additional environments.[1] Hence this paper asks whether a calibrated agent-based model populated with heuristic-driven agents predicts human-subject behavior in a *new* set of experiments in which property is not exogenously protected.

This interplay between laboratory experiments and agent-based models offers an important method of testing behavioral explanations of economic outcomes because creating agents requires specifying detailed decision rules for which simulated interactions with an

---

[1] Arthur (1991) offers an early agent-based model that replicates human decisions in a simple, two-choice bandit experiment, but the degrees of freedom in programming a model make it relatively trivial (given enough time) to replicate human subject data. If a model is required also to make out-of-sample predictions under environmental variations, then the generality of the underlying decision rules can be better established. See Duffy (2001) and Arifovic and Ledyard (2004) for another implementation of this methodology.

economic environment constitute predictions about human behavior. Many behavioral explanations from psychology, neuroscience and other disciplines can be compared by first formulating them as decision rules for computerized agents and then asking whether agents employing these rules in a given environment are able to predict human behavior in that same environment. Feedback from comparisons to human-subject data facilitates refinement of the decision model and may suggest additional experimental treatments necessary to settle disputes.

My agents operate by applying proven hill-climbing and reciprocity heuristics to discover trade and the benefits of specialization, and these same principles guide their discovery of theft and their decisions about whether to steal or trade. Since these general behavioral motivations are sufficient to characterize human behavior in discovering and implementing trade, they should also be sufficient to predict human subjects' implementation of theft. Surprisingly, agents employing these heuristics alone fail to predict human behavior in the new imperfect property enforcement environment; the simple model paints too bleak a picture and predicts degeneration of specialization and exchange as theft crowds out cooperation. Instead of abandoning trade, human subjects innovate to construct a novel property arrangement that facilitates exchange and specialization. They exploit the absence of property protection to develop a second method of trading (consensual taking, or "steal trading") reflected in a specific property convention that emerges in five of six sessions. However, when the potential adoption of that convention (also guided by reciprocity) is added to agents' behavioral repertoire, they again replicate human behavior.

Subject behavior highlights the importance of shared beliefs and the creation of conventions to support exchange. The social interpretation of the act of taking goods from another individual depends powerfully on the conventions in which the act is embedded. Thus, what is interpreted as theft in one case is interpreted as one-half of an exchange agreement under a different property convention. The results suggest that in open-ended environments, an accurate decision-model will require the introduction of agents that form beliefs about the beliefs of others and seek to coordinate those beliefs to form conventions that provide consistent interpretation of actions taken. In the case of property, agents would develop beliefs about which takings constitute violations of property, and then their behavioral heuristics would guide them to develop respect for property (or not) under the chosen interpretation.

Section II details the notion of property as a convention. Section III describes the economic environment, explains how possibilities for theft and the emergence of property are implemented in the agent-based model of Kimbrough (2009), and compares the new, imperfect property protection model to the original to define hypotheses for the new human subject experiments. Section IV describes the results of the experiments and compares them to the model predictions. Section V details the model as updated by observations from the experiments, and Section VI concludes and summarizes the findings.

II. Property as a Convention

Many societies have developed complex rules and laws to punish violators of property, and in general, students of property have focused on the creation of such rules when attempting to account for property's origin.[2] However, an account of enforcement institutions is not an account of the emergence of property, for any use of enforcement implies that property has already been violated. Punishment may satisfy a need for retribution (Levine 1998) or incentivize future adherence to a norm or convention (Fehr and Gachter 2000 and Bernhard, Fehr and Fischbacher 2006), but both of those hypotheses about the purpose of punishment imply the prior existence of some rule, the violation of which merits punishment. The question remains how to account for the emergence of this original rule and to explain how that emergence can be observed in experimental and agent-based environments.

Young (1993, 1996) details the game-theoretic emergence of conventions (by his definition, convergence to a single equilibrium when many exist) among boundedly-rational agents for a broad class of games. Repeated interactions may (by chance) give an edge to a particular strategy creating a positive feedback loop whereby agents employing that strategy are more successful, and the more agents employ the strategy, the more successful it becomes. In relatively simple agent-based models of games such as the prisoner's dilemma and the stag hunt, it has been demonstrated that naïve agents can converge on simple conventions that come to pervade a population through replicator dynamics (virtual natural selection) or by the epidemiological transmission of behavioral rules (Axelrod 1997, Skyrms 2004). Here I seek to explore the emergence of property conventions in a more complex specialization and trade environment, and it will be important to have criteria by which to judge whether a convention

---

[2] Ranging from Bentham (1802) and Westermarck (1908) to Wyman (2005) and Levine (2005) accounts of property's origin have focused on the *legal* aspects of property and the explicit recognition of rights.

has emerged.

Kimbrough, Smith, and Wilson (2010, hereafter KSW) define property empirically as an agreement or convention (either implicit or explicit) that eliminates theft in their experimental environment. They argue that although property, rooted in convention and observable only in its effects, may not exhibit the traits commonly associated with property rights in the folk wisdom (e.g. enforceable contracts, explicit punishment mechanisms, or arbitration), the fundamental fact of property is that it implies the absence of unwanted appropriations of goods. Property, like money, is a self-referential social *practice* – it can only be explained by reference to itself.[3] Thus, *property emerges empirically as the absence of theft*. This is crucial because this definition of property permits observing the endogenous emergence of an institution in both experiments and agent-based models. With agents, as with human subjects, to understand how respect for property emerges one must understand what behavioral rules and what sorts of interactions lead to a cessation of theft.

Kimbrough's agents employ simple heuristic learning methods and are able to replicate and predict human subjects' propensity to discover and exploit specialization and trade. Agents employing rules of reciprocity in exchange and trial-and-error, hill-climbing in specialization capture the variety of human behaviors in CSW's experimental environment. By extending these underlying behavioral principles to agents' employment of theft, I model how agents develop (or fail to develop) property conventions to support exchange. The model output should then predict human-subject behavior in attempting to solve the same social problem.

III. Simulation Environment, Implementation, and Results

III.A. The Economic Environment

The underlying economic environment is described in detail in CSW, KSW, and Kimbrough (2009), and all the same features are retained here. Agents (and subjects) are assigned one of two types, *odd* and *even*. These types define production functions with increasing returns to one of two goods and Leontief preferences over the goods, with a stronger preference for the good in which they possess increasing returns to specialization. Fully

---

[3] Bloor (1974) explains this as follows: a metal disc is a coin only in the context of money. Without a conventional notion of what it means for a disc to be a coin – embodied in its use to complete transactions – a metal disc would not be money. It would merely be a metal disc. In the same vein, what is property if others freely expropriate it? Howitt and Clower (2000) develop an agent-based model that yields an emergent market economy with universal adoption of commodity money.

specialized *odd* (*even*) agents can earn 90 (80) cents per period if they specialize completely and trade with a suitably specialized partner of the opposite type. Agents working in autarky can maximally earn roughly 1/3 of what they can earn by trading.

Goods appear in subjects' "fields" as they are produced and must be moved into "homes" in order to be consumed. CSW fully enforce property rights on all goods in subjects' homes and fields; that is, no one can take goods from another person at any time. Using agents calibrated to imitate CSW's subjects in their fully-enforced environment, I ask how the absence of property rights enforcement affects specialization and exchange in the agent-based model. Agents may not only reap gains from trade but also may steal from other agents to add to their own coffers.

III.B. Implementing Theft

The underlying agent-based model in this no property rights environment is exactly the same as that described in Kimbrough (2009).[4] Agents explore their economic environment incrementally, learning specialization and exchange via hill-climbing and reinforcement-learning (reciprocity). The sole difference is the addition of a function that permits agents to take goods from other agents unilaterally and without engaging in trade. Hereafter, such takings are referred to as "theft", and the new model will be referred to as the *T*-model. The following pseudocode gives a general overview of the model, and the italicized lines indicate additions to the Kimbrough model:

Model Pseudocode:

```
Initialize Agents
Set Model Parameters
Begin Loop Over Periods
    Begin Period
        Begin Loop 1
            Reset Indicator Parameters
            Select Prospective Trading Partners With Stochastic Discovery
            Produce Goods According to Choice of $t$ and Production Function
        End Loop 1

        Begin Loop 2
            If Agent is a Thief
                Steal Goods from Target Agents
                Update Probabilities of Theft
            End If
        End Loop 2

        Begin Loop 3
            Consume Goods to Maximize Autarkic Earnings
        End Loop 3
```

---

[4] The model details are included in the appendix below.

```
                        If Trading Pairs Formed Under Stochastic Discovery, then
                                    Begin Loop 4
                                                Trade Goods to Minimize Waste Among Pairs
                                                Update Probability of Future Trades with Interactive
                                                            Reinforcement
                                                *Update Probability of Theft Between Trading Pairs*
                                    End Loop 4
                        End If

                        Begin Loop 5
                                    Update Learning Rules
                                    Update Rates of Specialization with Hill Climbing
                                    Update Willingness to Trade
                        End Loop 5
            End Period
End Loop Over Periods
```

To introduce theft, it must be determined, first, whether each agent will steal; second, from whom agents that choose to engage in theft will steal; third, what impact this has on agents who are stolen from; and fourth, what rules might allow some sets of agents to overcome theft and develop respect for property. In general, the initial probability of theft should be non-zero and in some manner based on empirical data; being stolen from should incite retaliation; and agents should have some chance of crowding out theft via mutually beneficial trade.

Keeping those considerations in mind, theft in the model operates as follows. Agents are initialized either with or without a propensity to steal. Then, with some probability each agent with a propensity to theft will *actually* steal and will acquire all the goods produced by their target agent. Agents who are stolen from, even those with no initial propensity to steal, will become more likely to steal in the future and will direct their future thieving efforts at those who have stolen from them. On the other hand, to offset and potentially crowd out theft, trade relationships will diminish the future probabilities of theft and heal the cracks in inter-agent relationships. Theft occurs after production, but before consumption and exchange.

While KSW have previously performed human-subject experiments with imperfect property enforcement, the environment they employ was designed specifically to increase the social cohesion of the group. Because it is critical to the regular discovery of specialization and trade that each subject find a suitable trading partner, KSW employ a secondary treatment from CSW where subjects begin the experiment in pairs and are slowly merged into groups of four and eventually eight as time passes. In this way, the sociality of pairs is emphasized and the discovery of trade is facilitated. Here, I will ask how agents (and subjects) perform when they interact in groups of eight from the outset, creating an opportunity to test the model against a set of experiments that have not yet been performed. The operational details of the extended model follow.

III.B.1. Deciding Who Steals

As mentioned above, it is important that a randomly-instantiated agent's initial probability of theft be based in some way on empirical data. Some subjects in KSW's experimental treatments begin to steal goods from others almost immediately; others only begin to steal once they've been stolen from; and some never steal at all. As a heuristic with which to construct the agents, I employ the likelihood that an experimental subject engages in theft before using the chat room in an attempt to communicate. This seems a reasonable choice for two reasons.

First, given that their subjects are explicitly made aware of the chat room and must actually experiment with the interface to learn that theft is possible, such behavior suggests something about a subject's approach to the task. Second, subjects from KSW who talk before stealing earn, on average, \$0.11 (roughly 33%) more per period than those who steal before talking, suggesting that the heuristic is useful for categorizing experimental subjects. Thus, because 118 of the 192 subjects in the various treatments of KSW engage in theft before they attempt to communicate, each agent i is instantiated with a variable thief_i $\{0, 0.3\}$ with $P_i(\mathit{thief} = 0.3) = 118/192$, meaning that roughly 60% of agents will attempt to steal in the first period of the simulation. I choose 0.3 because it is roughly equal to one divided by the average period in which KSW's experimental subjects first engage in theft.[5] Thus, of the 60% of agents that attempt to steal, on average 30% will actually steal in the first period of a simulation.

After instantiation, any time an agent is targeted for theft, thief_i is incremented by a value called *stealProbabilityIncrement* to increase the probability of future theft.[6] Furthermore, if an agent has $\mathit{thief}_i = 0$, being stolen from increments this variable and adds another potential thief to the population. On the other hand, any time an agent i engages in trade with another agent j, both thief_i and thief_j are decremented by 2**stealProbabilityIncrement* to reduce the probability of future theft. The idea is that agents will engage in both positive and negative reciprocity, but that agents are more sensitive to the opportunity to forgive past wrongs for the prospect of future

---

[5] It is clear that an increase in this value will lead to an increase in theft and a diminishing of cooperation, so the initial value of this parameter is not systematically varied in the simulations described below.

[6] I fix the value of *stealProbabilityIncrement* at 0.1 for all reported T-model simulations.

benefits.[7] Thus, it is possible that all agents will see their probabilities of theft fall to 0 if enough trade occurs.

III.B.2. Rules to Determine a Thief's Target

As in the case of trade, each agent *i* stores a discrete probability distribution $S_{i,j}$ specifying the likelihood of agent *i* stealing from each other agent *j* in I. In each period, those agents for which *thief* > 0 make a random draw, *z,* from a $\text{uniform}_{[0,1]}$ distribution and compare it to *thief*. If $z_i < thief_i$, agent *i* will choose a target *j* from the set of other agents with probability $S_{i,j}$. Thus, each agent *j* will be chosen as the target of theft by agent *i* with probability = $thief_i * S_{i,j}$. In the first period, for any agent with *thief* > 0, $S_{i,A} = S_{i,B} = \ldots = S_{i,j,}$and $S_{i,i} = 0$ for each agent, and for any agent with *thief* = 0, $S_{i,A} = S_{i,B} = \ldots = S_{i,j,} = S_{i,i} = 0$. These probabilities are altered over the course of the simulation by the following process: 1) if agent B steals from agent A, $S_{A,B}$ is incremented by *stealProbabilityIncrement*, augmenting the probability that A returns the favor and steals from agent B in the future; and 2) if agent A trades with Agent C, $S_{A,C}$ is decremented by 2**stealProbabilityIncrement*.[8] Note the similarity of these effects to that on the probability of theft in general.

III.B.3. What Happens to Stolen Goods?

Recall that all theft decisions occur prior to the consumption and exchange portions of the model. Once an agent elects to steal from another agent, that agent takes all of the target agent's goods and treats them as his own for consumption and trade purposes. In the *experimental* environment, subjects may steal from as many other subjects as they like, but they are limited in their effectiveness by the prospect of real-time retaliation. Because in the agent-based version of the environment presented here theft must happen sequentially and not in real-time, and because theft may be cumulative (i.e. if A steals from B and then C steals from A, C acquires both A's and B's goods), I randomize the order of theft in each period. Furthermore, I allow agents to steal from *only one* other agent. Thus, it is possible (if all agents are stealing, and the ordering of theft is just right) that a single agent will end a period with all of the goods produced by all agents in that period. It is also possible that an agent will attempt to steal from an agent whose goods have

---

[7] The double impact of cooperation is partly a practical attempt to "give peace a chance" because theft is much easier than trade in the model. It only takes one person to steal, but for a trade to occur two agents must each select the other as a trading partner. The notion that people tend to be forgiving for the prospect of gains is based on subjective observation of human behavior in the CSW and KSW experiments, but the rule could be adjusted to examine its impact.

[8] $S_{i,j}$ is bounded below by 0 due to logical constraints. I assume that agents have perfect memories, i.e. that these probabilities do not fade over time.

already been stolen. This will not contribute to the breakdown of cooperation (that is, it will not adjust any of the relevant probabilities) because no actual goods change hands.

III.C. Simulation Results and Experimental Hypotheses

I employ the six parameterizations from Kimbrough (2009) under the new *T*-model, and the next section reports results on 1800 simulations of 35 periods each under each parameterization and compares these to the original *No T*-model. I compare the models in terms of efficiency and specialization, and the data from the *T*-model form my hypotheses for the new experiments. Let π_{it} denote the realized earnings of agent *i* in period *t* and π_i^a and π_i^c denote, respectively, expected earnings in autarky and at the global optimum for agent *i*. Define efficiency in period *p* for the agents of group *N* as $\frac{\sum_{i\in N}\pi_{it}-\sum_{i\in N}\pi_i^a}{\sum_{i\in N}\pi_i^c-\sum_{i\in N}\pi_i^a}\times 100\%$. And if q_{it} denotes the total amount of goods produced by agent *i* in period *t*, and $\overline{q}_i$ is the maximum amount of goods that agent *i* can produce when fully specialized, define the rate of specialization as $\frac{\sum_{i\in N} q_{it}}{\sum_{i\in N}\overline{q}_i}\times 100\%$.

Table 1 below lists each of the six parameterizations of the *T*- and *No T*-models and their average rates of efficiency and specialization as well as average efficiency and specialization in the original eight-subject CSW experiments and the new *Theft* experiments reported here. As expected, the introduction of theft to the model has a strong negative impact on both efficiency and specialization. Average efficiency is -10.66% in the *T*-model and 3.80% in the *No T*-model, and average rates of specialization are 37.40% and 43.05%, respectively. Mann-Whitney tests confirm that the differences are statistically significant.

Finding 1: The T-model simulations are significantly less efficient than the no-T simulations in weeks 3 and 6.

Following CSW and KSW, each 35 trading-period simulation (and experiment) is divided into 6 weeks. Week 3 efficiency measures the performance of each session at the halfway point, and week 6 efficiency measures output by the end. Table 2 reports two-sided Mann Whitney Tests comparing each T-model parameterization to its corresponding no-T model in both week 3 and week 6. These tests reject the null hypothesis of equal mean efficiency for every parameterization; and the difference in means suggests that efficiency is lower in the T-model. Additional, unreported Mann-Whitney tests reject the null hypotheses of equal rates of specialization in week 3 or week 6 for all of the treatments. Goods are also produced at a significantly lower rate with theft. As a robustness check, I also compare the T-model to the CSW experiments, under the hypothesis that the simulations will also be less efficient than the experiments.

Finding 2: The simulations are also significantly less efficient than the original CSW experiments.

Table 3 displays the results of 100,000 one-sided Mann-Whitney tests comparing efficiency in six randomly selected simulations for each week of each T-model parameterization to the CSW data for the same week. After week 1 in which none of the tests leads to a rejection of the null for any parameterization, the model is universally less efficient than the human subject data. Furthermore, property is almost universally disrespected.

Understanding property as an emergent convention revealed by the absence of theft, I define the strength of respect for property in a given period by the percentage of agents that engage in theft. Then, I define three levels of property enforcement over the final week of each simulation run. Simulation runs have perfect property when no agent steals from another agent in the final week, strong property when less than 10% of agents steal on average each day, and weak property when less than 20% of agents steal on average each day.

Finding 3: By week 6, property is respected in only roughly 5% the simulations.

Table 4 reports the number of sessions (out of 1800) that display perfect, strong, and weak respect for property in week 6 of each T-model parameterization. In only two parameterizations does respect for property (the sum of perfect, strong and weak runs) exceed 5%, with an average of 4.45%. The absence of exogenous protection of property leads agents to steal from one another uncontrollably as theft begets retaliation. The failure of reciprocity to create strong property conventions yields wasteful inefficiency and diminishes the returns to specialization.

I argue that the rules by which these computerized agents learn specialization, exchange, and theft constitute behavioral predictions about incremental learning and the positive and negative feedback that result from economic interactions between agents with limited rationality in a two-good production and exchange economy. The decision heuristics were generated via rational reconstruction of observed human behavior and an iterative process aimed at calibrating behavioral parameters so that agents would achieve outcomes commensurate with human behavior in an environment *with* exogenous property protection. By extending this model to a new environment *without* property protection, I have created a set of predictions (under the specified behavioral assumptions) whose validity can be tested by returning to the laboratory. Together, the aforementioned findings form my hypotheses for the new experiments. *The* Theft *experiments will be indistinguishable from the T-model and significantly less efficient than the CSW experiments due to the absence of effective property protection.*

IV. Experimental Treatment and Results

IV.A. The Theft *Treatment*

The new experimental treatment, which I call *Theft*, retains all preferences, production functions, and instructions from the experiments described in CSW and Kimbrough (2009). Subjects must discover the possibility for exchange and specialization in order to reap gains from trade. The *only* variation is that subjects may also, by exploring the computer interface, engage in theft. Based on the results of the simulations, I hypothesize that the absence of exogenous property enforcement will lead to a significant decrease in efficiency. Theft will hinder cooperation, diminishing specialization and exchange and increasing waste.

Eight subjects were recruited at random from the undergraduate student body of a private university in the United States to participate in each of 6 experimental sessions. They sat at visually isolated computer terminals and read instructions from the computer screen. Subjects received $7 for arriving to the experiment on time and received their earnings in cash privately at the end of each 90-minute session. The average experimental earnings were $9.69, ranging from a low of $0.48 to a high of $23.71. No subject participated more than once, and no subject had prior experience with a similar experimental environment. Instructions are included in an appendix below.

IV.B. Experimental Results

Table 1 displays average efficiency and specialization over all six sessions of the *Theft* treatment, and Figure 1 displays average efficiency by week and average specialization by day for each session. Note the high level of efficiency of Session 5 and also that only one session shows negative efficiency in week 6. It appears that the absence of exogenous property enforcement in this experimental environment does not have nearly as powerful a negative effect on efficiency as the results of the simulations suggest. In fact, by the end of session 5, it is more efficient than any CSW sessions!

Finding 4a: I cannot reject the null hypothesis of equal mean week 3 and week 6 efficiency for the Theft and CSW sessions.

Finding 4b: Nor can I reject the null hypothesis of equal mean rates of specialization in weeks 3 and 6 between the CSW sessions and the Theft sessions.

One-sided Mann-Whitney tests fail to reject the null hypothesis of equal mean efficiency in both week 3 ($U_{6,6} = 23$, p-value = 0.24) and week 6 ($U_{6,6} = 16$, p-value = 0.65) in favor of the alternative hypothesis that the *Theft* treatment is less efficient than the *CSW* treatment. Thus, the hypothesis on relative efficiency generated by the *T*-model can be rejected. Human subjects are able to reap gains from trade, even when property rights are not exogenously enforced.

Furthermore, one-sided Mann-Whitney tests fail to reject the null hypothesis of equal mean specialization in both week 3 ($U_{6,6} = 13$, p-value = 0.80) and week 6 ($U_{6,6} = 11$, p-value = 0.88) in favor of the alternative hypothesis that the *Theft* means are lower. Thus, contrary to the simulated hypotheses, human subjects also maintain relatively high rates of specialization in the face of insecure property rights.

By using the *T*-model to formulate hypotheses about the relationship between the *Theft* and *CSW* treatments, I also create the subsidiary hypothesis that the *T*-model will be indistinguishable from the *Theft* experiments. However, given that the new experiments show no difference from the original, it is unsurprising that this hypothesis also fails.

Finding 5: The T-model simulations are less efficient than the Theft *experiments.*

Table 5 reports bootstrapped 95% confidence intervals comparing average efficiency in each week over all six parameterizations. The intervals are computed by taking 100,000 samples of 6 T-model simulations each, from each treatment, and then dropping the lowest and highest 2.5% of sample averages to find 95% confidence limits for each week of each treatment. The average of the confidence limits for a given week defines the simulation interval. The experimental mean is contained in the confidence interval only in week 2, and it is actually lower in week 1. However, from week 3 to week 6, the experiments are more efficient than the simulations.

These findings beg the question of why the simulated data of the T-treatment compares so poorly to that of the human subjects in the *Theft* sessions. To answer this question, I return to the experiments and examine subject behavior in more detail. Because I have access to complete data on the flow of goods to and from experimental subjects and their expressed thoughts and intentions (in the form of chat room transcripts), I have a clear view of their behavior as it develops in real-time.

As the experimental sessions unfolded, it became clear that human subjects displayed ingenuity of which the simulated agents were incapable. Subjects learned quickly that while the ability to take goods from other subjects allowed them to *steal*, it also provided them an additional means of *exchange*. If I take from you with your consent and you take from me with my consent, this is economically equivalent to my actively giving you something in return for something that you give me, and this is precisely the arrangement that a number of the experimental sessions agree upon. In five of the six sessions I observe the emergence a property convention that permits mutual taking *with the same specific content*. The convention is embodied in the following chat transcript segments, each from a different session:

Person 7>: heres what to do
Person 7>: only take from other people's fields, not houses
Person 7>: whatever is in teh house at the end of the round is what you make money off
Person 7>: so don't jack other people's house stuff, just fields, and at the end of the round
Person 7>: feel me?
Person 1>: yeah, but you can jack from the fields too
Person 7>: yeah thats what im saying
Person 7>: jack from the fields
Person 7>: then everyone can still make profit from the houses

---

Person 6>: #1, can i steal 9 reds?
Person 5>: #2 can i take 5 blue?
Person 5>: #1?
Person 1>: take from my domino if you need red
(…)
Person 8>: so now that we've agreed to not steal from each other
Person 7>: can we take from dominos?
Person 6>: ok can we take form dominos now?
Person 1>: yeah, take just from dominos, not houses
Person 5>: take from other people's dominoes at THE END
Person 6>: maybe we should have a "grace" period
Person 6>: like the first 40 seconds don't take from anybody but yourself
Person 6>: then the last 20 seconds, use the left overs

Person 5>: the fields are fairgame, lets decide on that from now on, nobody takes ANYTHING from a house, if you have stuff to share you can put it in your field
Person 5>: your field is your yard sale lol
Person 4>: lets say after 30 seconds go by though
(…)
Person 4>: WE SAID THE FIELDS ARE FAIR GAME AFTER LIKE... 20 SECONDS

Person 7>: maybe people should keep their extra on the field and not in their house, and people can take from that?
Person 3>: that works better for you
Person 2>: that's a good idea
Person 4>: good idea
Person 2>: then we won't be stealing from each other
Person 7>: so if you have a whole bunch of red or blue you don't need, move it to the field so people can take it
Person 4>: i need red. put in field if you don't need it

Person 5>: i think ppl need to stop pulling from the houses
Person 3>: k
Person 1>: lets try not moving stuff
Person 2>: only add what u need
Person 3>: add from just the fields?
Person 5>: ya

In each of these transcript selections the content of the property convention is clear and specific: subjects agree first to consume the goods they have produced to the best of their abilities by placing them into their homes and *only then* to allow the others to take whatever is leftover in their fields in order to meet their own needs.

Human subjects innovate on their ability to unilaterally take goods from other individuals' homes and fields by adapting the social meaning of "taking" to their circumstances. Rather than treating all takings as malicious violations of property worthy of rebuke and retaliation, subjects come to agree that some takings (particularly those that occur after autarkic consumption has been optimized and explicit trade agreements have been completed) are not actually violations of property at all.[9] By altering the social meaning of property to permit some

---

[9] That this interpretation constitutes a conscious and radical alteration of subjects' initial views on property in this environment is evidenced by the vocal reaction to early unilateral takings in the chat transcripts. For example, by period 4 of Session 5, one subject has already instructed his counterparts to "stop stealing" and "JUST LEAVE EVERYONES STUFF ALONE". In Session 6 a subject laments "everyone just steals it all from everyone else how

types of unilateral taking, subjects in the *Theft* treatment are able to achieve unexpectedly high efficiency. On the other hand, my T-model agents are concerned only with the *act* of taking goods or having goods taken from them when making their decisions. The fact that agents in my simulations do not admit the possibility of steal-trading is likely the reason that they do not achieve equivalently high levels of efficiency. With this fact in mind, I develop a third simulation model that permits agents to engage in steal-trading.

V. Steal-Trading Model

I extend the T-model with a variation called the ST-model (for steal-trading). In addition to stealing from and trading with one another, ST agents may also consensually take unused goods from other agents' homes after the initial consumption and exchange period. Steal-trading agents respect other agents' rights to goods in their "houses", but they consider goods that are leftover in the fields to be "fair game". This additional behavior reflects the innovative convention that emerged in the experiments reported above; by choosing to engage in steal-trading, agents may eliminate wasteful theft and achieve efficiency on par with human subjects.

V.A. The Mechanics of Steal-Trading

After agents produce, steal, consume and trade, but before they update their learning rules and record their data for the next period, I introduce the possibility of steal-trading. Only those agents that have a positive probability of theft and are also willing to trade will attempt to steal trade. Agents that possess unconsumed goods search the set of other agents to determine which other agent has the highest amount of whichever good they are presently wasting the least. The intuition is that an agent steal trades with those agents that can best help it satisfy its preferences, *given* the set of goods it presently possess. If an agent is wasting a larger amount of red than blue, then the agent needs more blue, so it takes from the agent who has the largest amount of leftover blue.

A steal trading agent that has selected a target takes all of the target agent's unconsumed goods and consumes them according to its preferences. If unconsumed goods remain, other steal-trading agents may later take them as part of a second cooperative taking. Furthermore, since steal-trading is simply a second means of exchange, any agent presently unwilling to trade that is party to a steal trade, becomes willing to trade in the next period. Furthermore, if neither thief

---

pointless". But later in both of these sessions, subjects agree that taking is only a problem when goods are moved from houses. Some takings aren't pointless at all.

nor target engaged in trade in the present period, the probability that they trade in the future increases and their probabilities of theft decrease. Thus, both steal-trading *and* trading can ward off theft and crowd out its detrimental effects.

V.B. ST Treatments – Results

I perform 1800 simulations under the ST-model with each of the six parameterizations used for the T and No-T models. Table 6 displays average efficiency and specialization for each ST parameterization. I group the 1800 simulations for the *1515ST* parameterization into six sets of 300 sessions sorted by final period efficiency, and figure 2 plots average efficiency by week, and average specialization by day for each of the six groups. Note that in all but the two least efficient groups, efficiency is positive by the end of the session and generally increasing in time.

Furthermore, respect for property is extremely strong under the ST-model. Table 7 shows the number of sessions in each parameterization of the ST-model with perfect, strong, and weak property rights in week 6. More than 70% of these simulations yield perfect respect for property, and on average nearly 95% of sessions respect property at least weakly. The failure of agents to in the T-model to overcome theft in order to reap the gains from trade appears to have been solved by the imposition of the potential for steal-trading. I now compare the ST-model data to the human subject data of the *Theft* treatment. If my observations from the experiments about the importance of steal-trading are accurate, then the new simulation should be indistinguishable from the *Theft* experiments.

Finding 6: The null hypothesis of equal mean efficiency of the Theft *experiments and all* ST-model *parameterizations cannot be rejected in weeks 3 and 6.*

Table 8 displays bootstrapped 95% confidence intervals for average efficiency in 100,000 random samples of six simulation runs each, for all parameterizations, and compares those intervals to the experimental means from each week of the *Theft* treatment. After week 2, mean efficiency from the *Theft* experiments falls within the confidence interval for *all* parameterizations. Thus it appears that the introduction of steal-trading has solved the problems of the T-model. Furthermore, unreported confidence intervals suggest that, like the *Theft* experiments, the ST-model is indistinguishable from the CSW experiments with exogenous property rights enforcement after week 1.

By creating a subset of takings that are not interpreted as theft, agents are able to overcome the temptation of stealing to develop specialization and trade. Thus, it can be argued on the basis of these data that a lack of property rights enforcement is insufficient to diminish the productive power of a two-good economy populated with boundedly-rational agents. Because insecure property rights enforcement permits additional means of exchanging (via steal-trading), subjects and agents are able to eliminate the threat of rampant, costless theft despite their inability to directly protect their goods. If trade and steal trading have some probability of crowding out theft, then high levels of efficiency can attain.

VI. Summary and Discussion

In Kimbrough (2009) agents were created that accurately reproduce the patterns of human behavior in CSWs experimental environment. Here, I extend the agent-based model to a new environment in which the computer program no longer enforces property rights over goods. This model, called the T-model, maintains all features and rules of the original model, but it adds the ability of agents to steal from one another. The T-model yields a sharp decrease in efficiency relative to the original model and also to the original set of experiments. Agents steal from one another and this behavior escalates uncontrollably – in fact only 5% of simulation runs display even weak respect for property. Because the validity of the model depends on its ability to make predictions about behavior in a novel environment, the results of the T-model constitute predictions for a new set of human experiments employing the same institutional variation.

I perform the second set of experiments (the *Theft* treatment), and another battery of statistical tests rejects the hypotheses produced by the T-model. Not only are the *Theft* sessions more efficient than the T-model simulations, they are also indistinguishable from the human experiments with fully-protected property. Human subjects innovate to exploit a feature of the environment in a way that is impossible for the computerized agents; eliminating exogenous property enforcement permits theft, but it also permits cooperative taking, or steal-trading. The subjects develop a convention that alters the social interpretation of some instances of theft and thereby define a subset of unilateral takings that are permissible. In fact, 5 of 6 human sessions with theft develop a steal-trading convention with the *exact same content*: goods in houses (i.e. goods to be consumed) are inviolable, while goods in fields (i.e. wasted production) are “fair game”. Thus imperfect property enforcement allows subjects a second method of welfare-improving exchange.

I then create a third version of the agent-based model (the ST-model) that captures the spirit and effects of this convention. When steal-trading of wasted (i.e. unconsumed and untraded) goods is permitted and such steal-trading may offset the detrimental effects of theft, the ST-model agents once again replicate human experimental behavior. Thus, a third iteration of the model has taken a hypothesis about human social conventions from the observed chat room behavior (which has the potential to be mere cheap talk) and applied it effectively to create agents that mimic human subjects. Statistical tests demonstrate that the model is now indistinguishable from both the human *Theft* treatment and the original *CSW* experiments where theft is impossible. Boundedly-rational agents engage in welfare improving specialization and trade despite the absence of exogenous property enforcement – so long as they employ the appropriate convention. Importantly, when theft and steal-trading are impossible, the steal-trading version of the agent-based model reduces to the original model with perfect property rights, and it still accurately predicts human behavior in the *CSW* experiments.

Human subject experiments and agent-based models both offer unique views into the processes of economic behavior. Whereas field data on economic systems must be captured in a series of snapshots at various instants throughout a process, both computerized human-subject experiments and simulations allow one to observe the evolution of an economy as it happens. This paper combines the power of the two methods by extracting (or hypothesizing) rules of behavior for simulated agents from observed behavior in the lab.

Whereas Kimbrough (2009) derived behavioral rules for individual agents from the actions of lone experimental subjects, this paper extracts an economy-wide property convention to supplement individuals' rules. Respect for property emerges as a convention when subjects make clear the gains from trade, express to others that unilateral takings will only harm their ability to engage in mutually beneficial exchange, and then act in accordance with their stated views. Human groups develop property as part of a process that takes advantage of what Grotius calls our inherent "sociableness" (Buckle 1991). Subjects in the *Theft* treatment employ their sociableness to develop a steal-trading convention that allows them to overcome the incentives for theft and to reap the gains from trade, and it is clear that computerized agents without this social inclination cannot discover such conventions independently. However, when the convention is made available to boundedly-rational agents who learn incrementally to specialize, trade, and respect the possessions of others, those agents become indistinguishable from human actors in the same environment.

In that sense, the model is successful because it abstracts from the social aspect of exchange in order to capture the incremental process by which property conventions spread across a population as trade (in the form of both barter and steal-trading) crowds out theft. On the other hand, the model highlights the potential limitations of an agent-based approach to relatively open-ended social problems. Key to the subjects' discovery of steal-trading was their ability to build consensus on the interpretation, or social meaning, of the act of taking goods in various contexts. One way to move closer to an accurate decision model would be to create agents that bargain over and converge to an interpretation of various kinds of takings, but even in such a complex model, the possibility remains that agents would be unable to predict human behavior because the set of possible interpretations must be specified ahead of time.

References:

Arifovic, Jasmina, and John O. Ledyard. (2004) "Scaling Up Learning Models in Public Good Games," *Journal of Public Economic Theory*, 6: 203-238.

Arthur, W. Brian. (1991) "Designing Economic Agents that Act Like Human Agents: A Behavioral Approach to Bounded Rationality," *The American Economic Review, Papers and Proceedings*, 81(2): 353-359.

Axelrod, Robert. (1997) *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton University Press: Princeton, NJ.

Bengtsson, H. (2003) "The R.oo Package – Object-Oriented Programming with References Using Standard R Code" *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*. ed. Hornik, Kurt, Freidrich Leisch, and Achim Zeilis. Vienna, Austria.

Bernhard, Helen, Ernst Fehr and Urs Fischbacher. (2006) "Group Affiliation and Altruistic Norm Enforcement," *American Economic Review*, 96(2): 217-221.

Bentham, Jeremy. (1802) *The Theory of Legislation*, C.K. Ogden (ed.). Kegan, Paul, Trench, Trubner & Co.: London, UK. (1931)

Berkelaar, Michael and others. (2008) "lpSolve: Interface to Lp_solve v. 5.5 to Solve Linear/Integer Programs." R package version 5.6.4.

Bloor, David. (1997) *Wittgenstein, Rules and Institutions*. Routledge: New York, NY.

Buckle, Stephen. (1991) *Natural Law and the Theory of Property: Grotius to Hume*. Clarendon:

Oxford, UK.

Crockett, Sean, Vernon L. Smith, and Bart J. Wilson. (2009) “Exchange and Specialisation as a Discovery Process,” *Economic Journal*, 119(539): 1162-1188.

Duffy, John. (2001) “Learning to Speculate: Experiments with Artificial and Real Agents,” *Journal of Economic Dynamics and Control*, 25: 295-319.

Fehr, Ernst and Simon Gachter. (2000) “Fairness and Retaliation: The Economics of Reciprocity,” *Journal of Economic Perspectives*, 14(3): 159-181.

Henningsen, Arne. (2008) “micEcon: Microeconomics.” R package version 0.5-6.

Howitt, Peter, and Robert Clower. (2000) “The Emergence of Economic Organization,” *Journal of Economic Behavior and Organization*, 41: 55-84.

Kaufmann, Stuart and Simon Levin. (1987) “Toward a General Theory of Adaptive Walks on Rugged Landscapes,” *Journal of Theoretical Biology*, 128: 11-45.

Kimbrough, Erik O. (2009) “Heuristic Learning and the Discovery of Specialization and Exchange,” working paper, Chapman University.

Kimbrough, Erik O., Vernon L. Smith, and Bart J. Wilson. (forthcoming) “Exchange, Theft, and the Social Formation of Property,” *Journal of Economic Behavior & Organization*.

Levine, David K. (1998) “Modeling Altruism and Spitefulness in Experiments,” *Review of Economic Dynamics*, 1(3): 593-622.

Levine, Ross. (2005) “Law, Endowments, and Property Rights,” *Journal of Economic Perspectives*, 19(3): 61-88.

Neuwirth, Erich. (2007). “RColorBrewer: ColorBrewer palettes.” R package version 1.0-2.

R Development Core Team (2008). *R: A language and environment for statistical computing.* R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL http://www.R-project.org.

Skyrms, Brian. (2004) *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press: Cambridge, UK.

Starr, Ross M. (1976) “Decentralized Nonmonetary Trade,” *Econometrica*, 44 (5): 1087-1089.

Tesfatsion (1997) “A Trade Network Game With Endogenous Partner Selection,” in *Computational Approaches to Economic Problems: Advances in Computational Economics*. H. Amman, B. Rustem and A.B. Whinston, eds. Springer: New York, NY.

Warnes, Gregory R., Ben Bolker and Thomas Lumley. (2008) "gtools: Various R Programming Tools." R package version 2.5.0.

Westermarck, Edward. (1908) *The Origin and Development of the Moral Ideas*. Kessinger Publishing: Whitefish, MT. (2007)

Wyman, Katrina Maryam. (2005) "From Fur to Fish: Reconsidering the Evolution of Private Property," *New York University Law Review*, 80: 117-240.

Young, H. Peyton. (1993) "The Evolution of Conventions," *Econometrica*, 61(1): 57-84.

------ (1996) "The Economics of Convention," *Journal of Economic Perspectives*, 10(2): 105-122.

Figure 1: Average Efficiency by Week and Rate of Specialization by Day -- Theft Treatment

1(a) Efficiency (Blocks -- By Week)

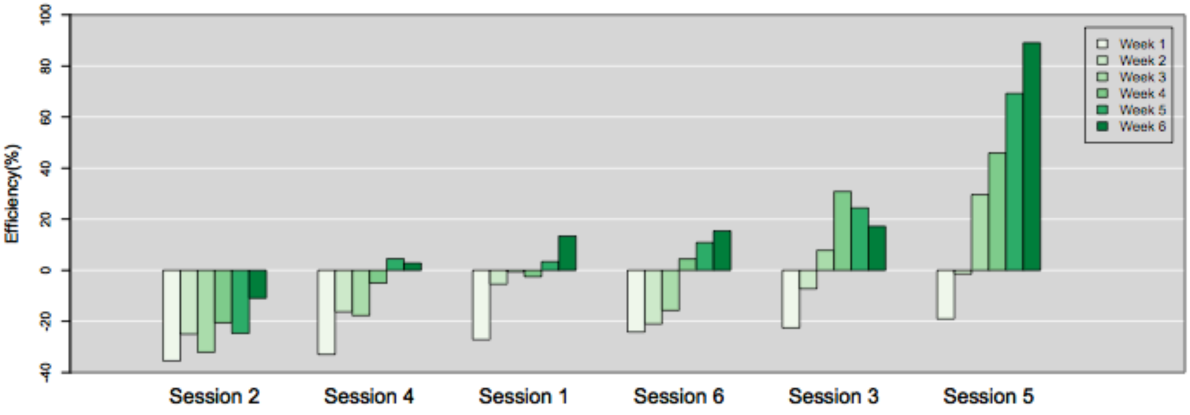

1(b) Specialization

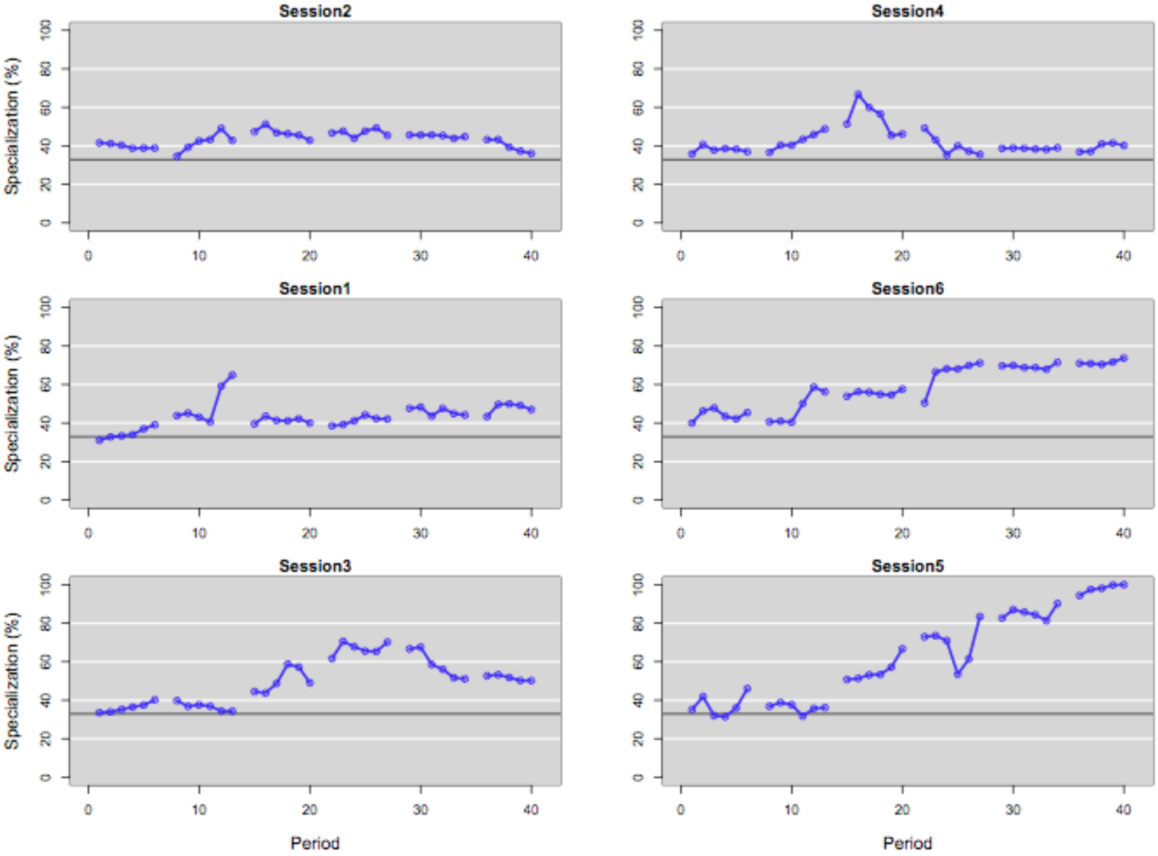

*I have ordered the experimental sessions by average efficiency in the last week, but I name them by the order in which they were performed.

Figure 2: Average Efficiency by Week and Rate of Specialization by Day -- 1515ST Treatment

2(a) Efficiency (Blocks -- By Week)

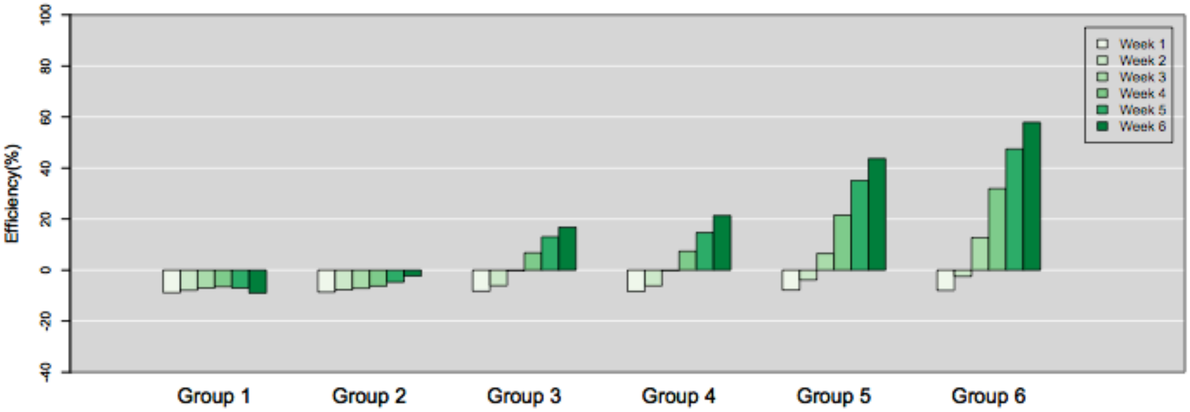

2(b) Specialization

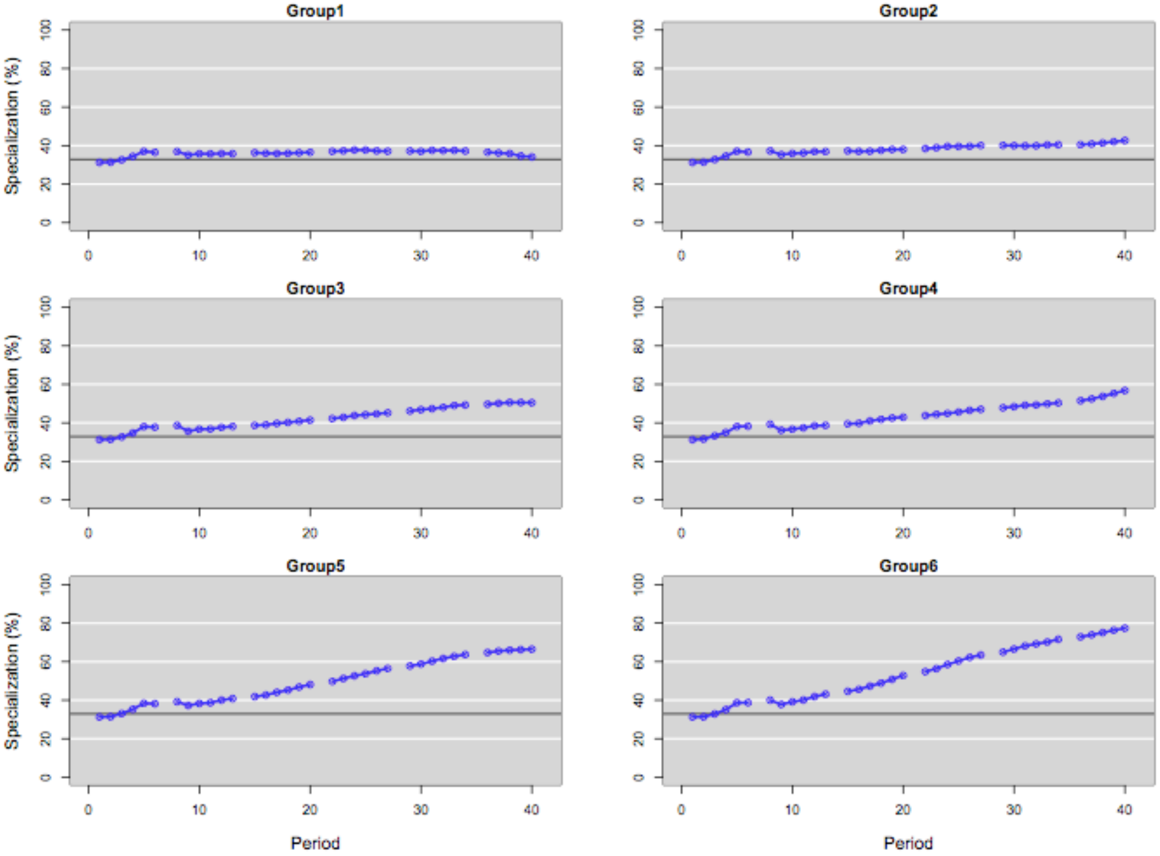

Table 1: Theft (T) Treatments – Parameters, Efficiency and Specialization

Treatment	*MinEarnAlone*	*MinEarnTrade*	*Efficiency*	*Specialization*
1015T	10	15	-11.44%	37.22%
1515T	15	15	-10.80%	37.48%
2015T	20	15	-10.36%	37.47%
1020T	10	20	-10.82%	37.22%
1520T	15	20	-10.34%	37.50%
2020T	20	20	-10.19%	37.53%
1015	*10*	*15*	*2.83%*	*44.23%*
1515	*15*	*15*	*4.38%*	*42.93%*
2015	*20*	*15*	*4.28%*	*41.66%*
1020	*10*	*20*	*3.21%*	*43.72%*
1520	*15*	*20*	*3.53%*	*43.76%*
2020	*20*	*20*	*4.58%*	*41.99%*
CSW	**NA**	**NA**	**4.24%**	**44.49%**
Theft	**NA**	**NA**	**0.05%**	**50.24%**

*Italicized entries are from Kimbrough (2009)

**Bolded entries are human subject experiments

Table 2: Two-Sided Mann-Whitney Results Comparing Efficiency of Theft and No-Theft Simulations by Treatment (1800 each)

Treatment	*Week 3 (efficiency)*	*Week 6 (efficiency)*
1015T	U = 2689740 ***	U = 2906682 ***
1515T	U = 2880796 ***	U = 3034669 ***
2015T	U = 2976736 ***	U = 3069754 ***
1020T	U = 2920274 ***	U = 2990344 ***
1520T	U = 2897438 ***	U = 2979198 ***
2020T	U = 3017562 ***	U = 3073506 ***

*** Significant at = 0.001

Table 3: # of Rejections (p=.05) -- 100000 Samples Each – T-Model vs. CSW Experiments

Treatment (Week)	*1015T*	*1020T*	*1515T*	*1520T*	*2015T*	*2020T*
CSW (1)	0	0	0	0	0	0
CSW (2)	98769	97924	97264	96880	96532	96358
CSW (3)	99994	99996	99979	99995	99993	99980
CSW (4)	99959	99976	99959	99951	99974	99965
CSW (5)	99953	99939	99973	99949	99962	99941
CSW (6)	99941	99960	99964	99950	99979	99974

Table 4: Strength of Respect for Property T-Model (Week 6) – Number of Sessions with Perfect, Strong, and Weak Property

Treatment	*1015T*	*1020T*	*1515T*	*1520T*	*2015T*	*2020T*
Perfect (0%)	9 (0.5%)	6 (0.33%)	7 (0.39%)	6 (0.33%)	9 (0.5%)	8 (0.44%)
Strong (<10%)	50 (2.78%)	44 (2.44%)	38 (2.11%)	37 (2.06%)	25 (1.39%)	27 (1.5%)
Weak (<20%)	42 (2.33%)	48 (2.67%)	30 (1.67%)	45 (2.5%)	23 (1.28%)	26 (1.44%)
Total	**101** **(5.61%)**	**98** **(5.44%)**	**75** **(4.17%)**	**88** **(4.89%)**	**57** **(3.17%)**	**61** **(3.39%)**

Table 5: 95% Bootstrapped Efficiency Confidence Intervals – T-Model vs. *Theft* Experiments (Bolded Entries Do NOT Contain the *Theft* Experimental Mean)

	T-Model	*Theft*
Week 1	**[-12.71, -9.37]**	-26.87
Week 2	[-13.54, -7.91]	-12.72
Week 3	**[-13.72, -6.07]**	-4.77
Week 4	**[-14.00, -4.69]**	8.87
Week 5	**[-14.08, -3.70]**	14.61
Week 6	**[-13.96, -2.34]**	21.17

Table 6: Steal Trade (ST) Model -- Efficiency and Specialization

Treatment	*Efficiency*	*Specialization*
1015ST	4.18%	44.49%
1515ST	5.63%	44.02%
2015ST	4.18%	42.14%
1020ST	1.86%	42.40%
1520ST	4.71%	43.79%
2020ST	5.17%	42.86%

Table 7: Strength of Respect for Property ST-Model (Week 6) – Number of Sessions with Perfect, Strong, and Weak Property

Treatment	*1015ST*	*1020ST*	*1515ST*	*1520ST*	*2015ST*	*2020ST*
Perfect (0%)	1338 (74.33%)	1345 (74.72%)	1305 (72.5%)	1373 (76.28%)	1292 (71.78%)	1303 (72.39%)
Strong (<10%)	306 (17%)	304 (16.89%)	318 (17.67%)	273 (15.17%)	328 (18.22%)	323 (17.94%)
Weak (<20%)	69 (3.83%)	64 (3.56%)	81 (4.5%)	71 (3.94%)	74 (4.11%)	68 (3.78%)
Total	**1713 (95.17%)**	**1713 (95.17%)**	**1704 (94.67%)**	**1717 (95.39%)**	**1694 (94.11%)**	**1694 (94.11%)**

Table 8: 95% Bootstrapped Efficiency Confidence Intervals – ST Model vs. *Theft* Experiments (Bolded Entries Do NOT Contain the *Theft* Experimental Mean)

Treatment	*1015ST*	*1020ST*	*1515ST*	*1520ST*	*2015ST*	*2020ST*	*Theft*
Week 1	**[-10.96, -7.04]**	**[-9.76, -6.13]**	**[-9.85, -6.41]**	**[-8.98, -5.77]**	**[-9.04, -5.90]**	**[-8.98, -5.83]**	-26.87
Week 2	**[-11.06, -0.01]**	**[-9.92, -0.90]**	**[-10.12, 1.68]**	**[-8.40, 0.72]**	**[-8.92, 0.06]**	**[-8.66, 0.02]**	-12.72
Week 3	[-9.73, 11.59]	[-8.55, 7.63]	[-8.43, 13.11]	[-7.00, 10.15]	[-8.29, 9.24]	[-7.53, 9.80]	-4.77
Week 4	[-6.24, 23.40]	[-7.37, 18.22]	[-4.67, 25.37]	[-4.53, 22.04]	[-5.87, 21.54]	[-4.34, 23.07]	8.87
Week 5	[-2.41, 34.19]	[-4.36, 26.68]	[-0.48, 35.43]	[-1.25, 31.55]	[-1.99, 32.25]	[-0.07, 33.99]	14.61
Week 6	[0.69, 40.93]	[-2.18, 33.73]	[2.39, 41.84]	[1.87, 39.47]	[1.15, 39.59]	[3.07, 41.14]	21.17

Appendix I. Model Details

A. Preferences, Production Functions, and Optimality:

Both types of agent (*Odd* and *Even*) have Leontief preferences (U_{type} over r units of *red* and b units of *blue*), and: $U_{odd} = \min\{r, 3b\}$ and $U_{even} = \min\{2r, b\}$. Each type produces with increasing returns to one of the two available goods, *red* and *blue*. Specifically, *odd* agents produce according to: $R_{\text{odd}} = \frac{13}{10\sqrt{10}} t^{\frac{5}{2}} \approx 0.41t^{\frac{5}{2}}$ and $B_{\text{odd,}} = \frac{10}{10-\left(\frac{300\sqrt{10}}{13}\right)^{\frac{2}{5}}}(10-t) \approx 2.25(10-t)$ and *even* agents produce according to: $R_{\text{even}} = \frac{13}{10-\sqrt{\frac{260}{11}}} t \approx 2.53t$ and $B_{\text{even}} = \frac{11}{10}(10-t)^2$ where t represents time devoted to production and is chosen by each agent between 0 and T = 10. Production is constrained to integer values

In autarky, *odd* (*even*) agents optimally spend 56% (51%) of their time producing red to create 30 (13) reds and 10 (26) blues. This yields utility of 30 per period for *odd* agents and 26 per period for *even* agents. At the optimum, on the other hand, each type specializes in the production of one good and trades to acquire the other good. *Odd* agents specialize completely in red ($t = 10$) to produce 130 red, and even agents specialize completely in blue ($t = 0$) to produce 110 blue. By trading at a price of 4 red to 3 blue, odds consume 90 red and 30 blue for a utility of 90, and evens consume 40 red and 80 blue for a utility of 80 per period.

B. Model Details

B.1. Agent Initialization:

At the beginning of each simulation a number of parameters must be assigned to each agent. Some of these parameters help agents make their decisions and may be altered via interactions with other agents and with the environment. Table 1.A lists these parameters, what they mean, and how they are initialized. I will refer back to these variables in explaining the model below and will italicize them to ease the reader's recognition.

B.2. Loop 1 –Parameter Reset, Trading Partner Selection, and Production

In the first loop agents set the relevant parameters for the interactive stages of the model. They reset indicator parameters used in the previous period to their defaults (currently stealing = FALSE, this period's earnings = 0, etc.), choose prospective trading partners, and produce according to their choice of t.

When choosing a trading partner, an agent first checks its value of *Willing* to determine whether it is willing to engage in trade with another agent. Each willing agent then selects a prospective trading partner. Specifically, each agent $i \in I$ stores a *TradeVector*, Δ_i, containing probabilities $\delta_{i,j}$ of attempting to initiate an exchange relationship with each other agent *j* in the simulation. In the first period, $\delta_{i,A} = \delta_{i,B} = \ldots = \delta_{i,j}$ and $\delta_{i,i} = 0$. That is, the initial probability of choosing any given agent as an exchange partner is equal for each other agent. This suggests that agents have no prior reason to choose any one agent over another, and the probability of an agent attempting to trade with itself is zero for obvious reasons.

Using this probability distribution, each willing agent randomly selects one agent ϕ as a prospective trading partner. For example, agent A draws a random index from the set of other agents with probability $\Delta_{A,j}$. Agent B also draws a random index, given its own distribution over the set of other agents. If two agents each choose one another in the same period, i.e. $\phi_A = B$ *and* $\phi_B = A$, then the agents will exchange later in the period. This method of developing exchange relations can be described as a two-sided stochastic discovery process by which agents on either side of a prospective exchange sample the other available agents as potential trading partners until a pair share similar goals and are able to initiate a trade.[10]

After selecting trading partners, agents move to the production phase of the period. Based on its *Type* and *SpecializationRate* each agent produces according to the relevant production function specified in the Appendix. At the beginning of the simulation, each agent *i* sets $t = 5$ in its production functions, meaning that the agents devote 50% of their time to the production of each good and are extremely close to the optimal autarkic level of production.[11] They store their red and blue production for later phases of consumption and exchange.

B.3. Loop 2 – Consumption

In this loop, each agent maximizes autarkic earnings by mapping production to its utility function. Given an agent's *Type*, consumption is maximized over the relevant utility function. Given an *odd* agent's stock of red and blue, it will consume the minimum of two values: 1) its quantity of red and 2) triple its quantity of blue. An *even* agent will consume the minimum of either 1) twice its quantity of red or 2) its quantity of blue. The remaining goods that cannot be

---

[10] Tesfatsion (1997) develops an agent-based "trade network game" with endogenous partner selection using a variation on the Gale-Shapley mechanism (1962) to select trading pairs.

[11] CSW allow their subjects to set their levels of specialization after exploring their production functions for a few minutes before the experiment. If subjects do not choose their own rate, a 50-50 split is the default.

consumed in appropriate proportion are set aside for trading (if possible) and may be wasted if no trade occurs (since unconsumed goods do not carry from one period to the next).

B.4. Loop 3 – Trade

To carry out trades, I employ a bilateral process similar to that which generates the decentralized non-monetary general equilibrium described in Starr (1976). The model identifies the set of possible trades, (i.e. it identifies all pairs of agents (i, j) for which $\phi_i = j$ & $\phi_j = i$) and applies an optimization to distribute the goods each agent has available for trade. Remember that agents have already consumed their available goods to their best autarkic ability. The goods they exchange are only those they could not consume themselves.[12] When two agents exchange, their leftover goods are pooled, and an integer-programming problem is solved to minimize total waste and distribute the goods across the subjects' utility functions accordingly.[13] All remaining goods that are not consumed either privately or as a result of exchange are wasted and destroyed

---

[12] I chose this method because it maps nicely into behavior observed in CSW, where subjects tend to first consume all they can of what they have produced and *only then* begin to look for opportunities to gain from exchange. I discovered Starr's paper late in this project while trying to understand the implications of subjects' trading intuition and was struck by how well the model matched up with subjects' apparent logic. In Starr's model, each agent trades with each other agent *repeatedly* until stocks have been exhausted and equilibrium is achieved. However, an economy operating under this principle is not guaranteed to converge in finite time. Here I use only a single-exchange approximation of his process.

[13] CSW often observe subjects making statements like, "I need x blue so I don't waste any red." Furthermore, subjects often come to the idea of trade because they have leftover goods that would otherwise be wasted. Thus, waste-minimization seems a reasonable assumption about how agents approach the opportunity to exchange. Of interest to both price theorists and game theorists, employing this algorithm means testing first for each agent's *Type* and then applying a different optimization depending on the types of the agents engaging in the exchange. The program has, by default, perfect knowledge of agent preferences and employs a computational method that far exceeds in complexity anything that human subjects can perform mentally. The most efficient allocation of goods is computed and the goods are distributed among the traders accordingly, without regard to who possessed what before the exchange took place. Thus, distributional, strategic, and computational issues that may impact real-world exchanges are simplified away. I want to be clear that my choice of price determination method avoids the subtle implication that subjects in an experiment solve the price determination problem in this way. In general, economics assumes that people are intelligent actors who seek to optimize consumption relative to their preferences. While people may not use these algorithms to solve the problem in practice, their behavior is supposed to approximate the algorithmic solution just the same. However, if the social nature of exchange matters or the relative information sets of economic agents differ, exchange prices and quantities may differ, and hypotheses on these factors are interesting points that merit study in their own right. For example, if someone hypothesizes that concerns for fairness play a role in people's exchange decisions, I could impose a rule requiring agents to refuse exchanges that disproportionately benefit the other party. Or if one believes that strategic factors impact bargaining behavior, rules could be imposed requiring agents to bargain over their exchange rate in proportion to their production or the number of trade relationships available to them. On the other hand, if the concern were that the intelligence of the agents too far exceeds that of real-world actors (in that humans can't be expected to perform rapid optimizations), one way to address this would be to add noise to the exchange price derived from the optimization. I do not explore these variations here, but they suggest avenues for future research. The integer constraint is drawn from CSW.

at the end of the period. Note that this rules out an agent having multiple trading partners in a single period.[14]

Using a rule similar to "interactive reinforcement" as described in Skyrms (2004), agents adjust the probability distribution (their *TradeVector*) to reflect the outcome of their exchange. An initial version of the model employed Skyrms' reinforcement model directly, augmenting the probability of future exchanges by the payoffs resulting from the present exchange. However, such a simple rule could not replicate human behavior. For the final version of the model, I employ an updated version of the rule, introducing a notion of acceptable risk and a willingness of agents to simply give up on trade altogether.

Specifically, each agent *i* in an exchange compares its earnings, e_i, to a global value called *minimumEarningsTrade*, meaning the minimum earnings an agent is willing to accept as the outcome of a period in which it has exchanged with another agent. This value is a proxy for both a discount rate and a level of risk-aversion for agents engaged in exchange, and is homogeneous in the population of each simulation. If $e_i \geq$ *minimumEarningsTrade* agent *i* will increment the probability $\delta_{i,j}$ of trading with agent *j* in the future and offset that increment by decrementing equally their probability of trading with each other agent in Δ_i. On the other hand, if $e_i <$ *minimumEarningsTrade*, agent *i* will decrement $\delta_{i,j}$ and increment equally all other agent probabilities in Δ_i.[15] If any $\delta_{i,j} < 0$, the $|\delta_{i,j}|$ is redistributed equally over the remaining positive, non-zero probability agents that were not involved in the exchange. If all probabilities fall to zero, the agent becomes an autarkist. The agent refuses to trade (*Willing* = FALSE) and becomes short-sighted in exploring its production function (i.e. paying attention only to the last period's earnings when deciding which direction to move its *SpecializationRate* rather than the global values of *minimumEarningsTrade* or *minimumEarningsAlone*).

B.5. Loop 4 – Parameter Updates

Next, on the basis of earnings in the present period, each agent makes decisions about its method of learning, rate and direction of specialization, and willingness to trade. Throughout each session, agents track their best autarky earnings and record the associated level of

---

[14] CSW observe that exchange is largely (though not completely) bilateral and becomes increasingly so over time. It would certainly be possible to extend this model to second- and third-order exchange relationships, but that route has not been pursued here, even though the model would then be closer to Starr (1976). Furthermore, leftover goods from the first optimization will tend to be small and secondary trades would often be of little marginal value.

[15] The default increment is 0.3 so reinforcement is fairly strong. The decrementing process differentiates my reinforcement from that of Skyrms because it allows some probabilities to fall to zero.

specialization. If an agent is not involved in exchange in the present period, it compares its earnings e_i to its previous best in autarky and updates that memory if the present period exceeds its current value.

B.5.a. Learning Rule

Once eligible agents have decided whether to commit their present circumstances to memory, each agent evaluates the effectiveness of the learning rule it applied in the previous period. If an agent has traded in the present period, it compares its present earnings to *minimumEarningsTrade* and changes its learning rule to short-sighted if it has earned less than this value. If a trading agent has already decided to become short-sighted, it also compares its earnings in the *previous* period to *minimiumEarningsTrade*. If $e_{i,p-1} <$ *minimumEarningsTrade*, it becomes unwilling to trade in future periods and reverts its specialization rate to its autarkic best. On the other hand, a short-sighted agent that has benefited from exchange ($e_i >$ *minimumEarningsTrade*) will adjust its learning rule to become far-sighted.

If, instead, the agent has acted as an autarkist, it adjusts its learning rule on the basis of comparison to *minimumEarningsAlone*, becoming short-sighted if its earnings are too low. When agents make the decision to become short-sighted, they also change their *t*-value to match their previous best autarkic performance.

B.5.b. Specialization

Next, each agent applies its learning rule to adjust its *SpecializationRate* via a modified hill-climbing process. [16] Initial experimentation with the basic hill-climbing model failed to replicate human behavior, so I developed the more complex version described here, which allows agents to adjust the parameters of their hill-climbing process in response to events. There are two possible values of *learningRule*, each of which can be thought of as describing an agent's discount rate and/or level of risk-aversion in hill climbing. If agents employ a far-sighted *learningRule* they compare their earnings in the present period to a fixed, global value called *minimumEarningsAlone* that represents the minimum an agent would be willing to earn (without adjusting its strategy) when working as an autarkist. The agent makes this comparison to decide whether to continue moving in the same direction on its production function or whether to go the other direction (to increase or decrease *t* in the next period's production phase). If $e_{i,p} <$

---

[16] Kaufmann and Levin (1987) explore hill-climbing processes on "fitness landscapes" in the context of biology.

minimumEarningsAlone, the agent alters its trajectory; if $e_{i,p} \geq$ *minimumEarningsAlone* the agent makes no change. On the other hand, if agents employ a short-sighted *learningRule* they compare $e_{i,p}$ to $e_{i,p-1}$. If $e_{i,p} \geq e_{i,p-1}$ agent *i* continues to move in the same direction on its production function in period $p+1$. If $e_{i,p} < e_{i,p-1}$ the agent begins to move its rate of specialization in the other direction. Once this decision has been made, the agent updates its *t*-value for the next period.

A specific example of how the short-sighted rule may detrimentally impact the development of specialization and trades is as follows: An *odd* agent that is incrementally increasing its rate of specialization while its trading partner is already fully specialized (or nearly) may actually experience a decrease in earnings at one stage in the process relative to the previous period. Here the algorithm takes goods from an under-specialized *odd* agent and transfers them to the more specialized *even* agent with which it is trading. Although joint profits are maximized and steadily increasing as odd specialization increases, for some levels of specialization, the *odd* agent experiences a decrease in welfare relative to a case in which it is less specialized. In such cases, an odd agent employing the short-sighted learning rule would reverse the direction of its increase in specialization, and might never achieve the optimal outcome. (See Kimbrough 2009 for graphics that support this argument)

Path-dependence resulting from the initial rates of specialization of new trading partners (initial location on the profit landscape) may prevent a given pair from converging to the global optimum, instead stranding them on a local peak or breaking their trading relationship apart before it can become mutually beneficial. An agent employing the far-sighted learning rule, on the other hand, will more likely ignore this temporary decrease and continue specializing.

B.5.c. Willingness to Trade

Finally, if other agents in the simulation are trading in this period, unwilling or short-sighted agents may decide to learn from their examples. Agents have a randomly assigned level of *Conservatism* between 0 and 1 that increases each time they change from willing to unwilling. They draw a random number from a $\text{uniform}_{[0,1]}$ distribution and compare this to their level of *Conservatism*. If the random draw exceeds their conservatism, they revert from unwilling back to willing.[17] Furthermore, agents that have become willing to trade have some random probability

---

[17] Additionally, agents that became autarkists due to their trade probability distribution (their *tradeVector*) falling to 0 for all other agents have that vector reset to the default.

that they will mimic the production decision of another agent in the model. A second random draw is compared to their probability of *Mimicry*, and if the draw is greater, the agent randomly copies the production decision of one agent whose earnings were higher in the present period. After this loop is completed, relevant data is recorded and the simulation begins a new period.

C. Model Parameterizations

I consider 6 different parameterizations, where each varies two values that guide agents' learning rules: 1) the level of earnings below which each agent chooses to adjust its learning rule when acting as an autarkist, or *minimumEarningsAlone*, and 2) the level of earnings below which each agent chooses to adjust its learning rule when trading with another agent, or *minimumEarningsTrade*. These two parameters are homogeneous in each session's agent population, exogenously imposed and, when combined with random draws that guide behavior, determine the evolution of the economy.

MinimumEarningsAlone and *minimumEarningsTrade* act as proxies for an agent's risk aversion and/or discount rate; think of them as representing an agent's maximum acceptable risk at any given level of production and conditional on whether that agent acts in autarky or trades. A higher value of either means that agents will more quickly abandon incremental production-space searches that produce relative decreases in short-run earnings. Agents with lower values, on the other hand, will assume additional risk in the short run when exploring the behavior space in hopes of finding long-term benefits.[18]

Table 1A: Agent Behavioral Variables

Variable Name	*Purpose*	*Initialization*
ID	Agent's index	Index in agentVector
Type	Odd or even	1/2 odd and 1/2 even
SpecRate	Value of t in production	Set to 5
LearningRule	Far-sighted or short-sighted	Random
Increase	Increase or Decrease t when updating	Random
Willing	Willing or unwilling to trade	2/3 willing
TradeVector	Vector of trade partner probabilities	1/(N-1) for each other agent
ProbSwitch	Probability of updating Willing	Random

---

[18] Note that agents in this model never completely abandon their hill-climbing specialization procedure. Some human subjects, foiled in early attempts, simply stop looking for welfare improvements altogether and settle into autarky. While agents may eventually become unwilling to trade, they can always discard this conservatism and reengage, and even when they employ a short-sighted learning rule, they never completely stop their search of the production space. Only when they are fully specialized and trading at the competitive equilibrium with no theft, do they actually settle into an unchangeable pattern of behavior.

Mimic	Probability of mimicking another agent	Random

Appendix II. Experiment Instructions

<page 1>
This is an experiment in the economics of decision making. The instructions are simple, and if you follow them carefully and make good decisions you may earn a considerable amount of money which will be paid to you in CASH at the end of the experiment.

In this experiment you are **Person 2**. You and the other **7** people in this experiment each have the ability to produce two fictitious items: red and blue. For the first **10** seconds of each period, you will produce items in the upper left portion of your screen. Using the scroll bar, you can change the proportion of each second allocated to producing red and blue. Each person's production is displayed in the domino-shapes at the bottom of your screen.

When a domino-shape or house is selected, its contents are displayed in the top left portion of your screen. To select a domino-shape or house, left click on it and it will be become highlighted in yellow.

<page 2>
After the production phase ends, the period continues for another **90** seconds. When the clock expires, you earn cash based upon the number of red and blue items that have been moved to your house. To select items to be moved, *left* click on an item or click on the red or blue buttons at the top of the screen. The yellow highlighted items can be moved by dragging with the *right* mouse button. The maximum number of red or blue items a house or field can hold is 170. (You cannot move items until the experiment has started or during the production phase.)

The specific information on how the red and blue items in your house generate earnings is given in the upper right corner of your screen. You personally earn (in cents) the minimum of the following two numbers:

2 times number of red items,
number of blue items.

Or, think of it this way. You earn by consuming what's in your house in the proportion of 1 red to 2 blue items. For every 1 unit of red you need 2 units of blue to earn 2 cents. Your potential profit updates as items, unit by unit, are moved into your house.

<page 3>

Everyone in this experiment can send text messages. Everyone can read all posted messages. In the center of the screen, you can type a message in the text box next to the send button. To send a text message press the Send button. There are two chat rooms. Messages sent to Chat Room A will only appear in chat room A. Message sent to Chat Room B will only appear in chat room B

Under your house you can also post a one-line message that will be visible at all times to the other players.

You are free to discuss all aspects of the experiment, with the following exceptions: you may not reveal your name, discuss side payments, make threats, or engage in inappropriate language (including such shorthand as 'WTF'). If you do, you will be excused and you will forfeit your earnings.

<page 4>
During the experiment, every 7 periods will be a "break period" in which nobody produces anything but that the chat rooms are still open.

You can open a table of your production possibilities by clicking on the Show %s button. This table will fill in every time you change the proportion of time allocated using the scroll bar.

This is the end of the instructions. If you wish to explore how you produce red and blue items, click the Practice button. You may change the proportion of time allocated to producing red and blue items using the scroll bar, and you may Practice as many times as you wish. (You will not be able to move items until the experiment has begun.)

If you wish to review the instructions, you may go back at this time. If you feel you are prepared to proceed with the actual experiment, click on the Start button. The experiment will begin once everyone has clicked on the Start button. If you have a question that you feel was not adequately answered by the instructions, please raise your hand and ask the monitor before proceeding.