Endowment Effects in Evolutionary Game Theory: Enhancing Property Rights

Eric J. Hammer
Department of Economics
George Mason University
Fairfax, VA 22030
ehammer@gmu.edu

Abstract:
There are many instances in nature of animals having a form of property rights, a respect for mine and thine. This paper addresses the question of how a norm of respecting property rights can emerge within a group of individuals without access to third party enforcement. Building upon work in economics and biology, an agent based model is presented that examines the processes by which property rights emerge within a population. Consistent with findings in biology and evolutionary economics, this process dispenses with deliberately planned creation of property rights and demonstrates their spontaneous emergence by virtue of a few simple factors. Key to this emergence is the endowment effect, which causes owners of a good to be more willing to fight against a challenger to retain possession. This asymmetry in action creates a strong tendency towards challengers respecting the possession claims of owners over time, resulting in emergent property rights.

Acknowledgements:
I would like to thank Dr. Robert Axtell and Dr. Bryan Caplan of George Mason University for insightful comments regarding the logic of the model, David Christensen for help with the early draft and design, and Dr. Daniel Klein of George Mason University for his help in organization and framing. I would also like to thank Dr. Georg Vanberg of Duke University and other participants in the Public Choice Society 2016 Vincent & Elinor Ostrom Prize seminar for their insightful comments and suggestions.

1.0 Introduction

Property rights constitute one of the cornerstones of free market economies. Their importance is widely recognized by economists; their operation within markets extensively studied. For all that, their origin remains shrouded in the mists of history – or, as is increasingly argued under a Coasean vision, prehistory (Bowles & Choi, 2003). Building upon game-theoretic work in economics and biology, I examine the emergence of property rights from the bottom up through the use of an agent-based model. By instantiating a population of agents which can choose to play either hawk or dove in pursuit of resources, and which recognize whether or not their opponent is the current owner of said resources, a process by which property rights can emerge is revealed. Consistent with findings in biology and evolutionary economics, this process dispenses with deliberately planned creation of property rights and demonstrates their spontaneous emergence by virtue of a few simple factors.

Key to this emergence is the endowment effect, which causes owners of a good to be more willing to fight against a challenger to retain possession. This asymmetry in action creates a strong tendency towards challengers respecting the possession claims of owners over time, resulting in emergent property rights. Further, the presence of an endowment effect requires no strategy or information on the part of agents beyond which agent is an incumbent.

2.0 Background Literature

Traditional theories of property rights, including theories regarding the genesis of property rights, originated in the fields of economics and law. Early political philosophers including Hobbes and Locke (Buckle, 1991) viewed property rights as dependent upon legal institutions and enforcement. Hume, in contrast, noted that the convention of property "arises gradually … and by our repeated experience of the inconveniences of transgressing it" (Krier, 2008). David Friedman similarly argues that people over time learned to "[contract] out of the Hobbesian jungle" and that the act of contracting itself establishes Schelling points around which people coordinate – property rights themselves being among the most notable of these Schelling points (Friedman, 1994). Among modern economists, no doubt the most well-known theory regarding the origin of property rights is Harold Demsetz's "Toward a Theory of Property Rights" (1967). Demsetz uses the example of the extraction of beavers among certain Native American tribes before and after the arrival of European fur traders. While his article has been subject to various criticisms (Krier, 2008) for our purposes it is sufficient to note that Demsetz offers an explanation of *why* property rights emerged (the benefits of property rights outweighed the costs) but not *how* they emerged, which is the question I wish to illuminate. Just how is it that societies can settle upon norms of property rights before the advent of reliable third party enforcement?

Among legal scholars, the traditional explanation for the emergence of property rights presumes that they were consciously designed or chosen – a presumption that typically assumes the presence of a state or state-like authority to implement the design. This view, which remains dominant in legal thought today, can be seen in classical sources although a prominent exception to this is Blackstone (Krier, 2008). Challenges to this paradigm have emerged in recent years, for instance see (Jones & Brosnan, 2008), (Leeson, et al., 2006), (Leeson, 2009).

In contrast to traditional theories, a growing body of research in economics and law takes advantage of developments in biology making use of game theory and evolutionary approaches. Nearly all of this research is built on hawk-dove game theoretic models, and especially the hawk-dove-bourgeois extension of hawk-dove games. The hawk-dove-bourgeois game takes into account the propensity of individuals to value something they possess higher than its equivalent when not in possession. This propensity is referred to as the endowment effect (Thaler, 1980) (Kahneman, et al., 1990) (Hoffman, 1993) (Huck, et al., 2005). Typically, the author of a game-theoretic paper on property rights will set out to explain the emergence of property rights by showing that the bourgeois strategy (in the presence of the endowment effect, play hawk when the owner of an item, and dove when challenging an owner for an item) is the dominant strategy, or the evolutionarily stable strategy, pioneered by biologist John Maynard Smith (Smith & Parker, 1976).

Biologists, having observed something that looks very much like property rights in species of animals as varied as primates and butterflies, explain the evolution of the endowment effect as the determining factor in the choice of strategy as being rooted in its role as an asymmetrizing factor (DeScioli & Wilson, 2010). In other words, possession of something is a relatively visible state around which players can base their choices (Gintis, 2007) (Burtsev & Turchin, 2006) (Davies, 1978). This line of thought is developed in Sugden (1989) and Yee (2002).

Gordon Tullock (1994) and Peter Leeson (2008) both offer work that bridges the fields of economic and biology. Tullock applies the insights of economics to cooperative animal societies, noting how amazing it is that many animals, even cells, engage in specialized division of labor and other cooperative activities all without command,[1] command that humans are often presumed to require to achieve the same ends. Leeson starts with these insights and turns them back to the realm of humans, presenting evidence for two spontaneously cooperating groups, “human ant nests,” one with great disparity in individuals’ strength, the other with great social heterogeneity. These findings of coordination without a third party to enforce, or ever define,

[1] In the case of ants, without even a brain complex enough to grasp the concept.

coordination follow Leeson's other examples of spontaneous order despite suboptimal conditions (Leeson, 2008), (Leeson, 2009).

Whether in biology, economics, or law, most studies of the emergence (as opposed to design) of property rights are done analytically, using game theory. However, emergent phenomenon lend themselves to algorithmic study. Indeed, if property rights are in fact an emergent phenomenon, an algorithmic approach has the advantage of allowing the theorized process by which property rights develop to be examined as it unfolds.

There has been some work to this effect done, notably by Samuel Bowles and Jung-Kyoo Choi at the Santa Fe Institute (2003). Bowles and Choi use agent-based modelling and invoke collective punishments and between-group conflict to mirror patterns of behavior among nomadic early humans. Their agents can be sharers, grabbers, punishers, or bourgeois. Burtsev and Turchin (2006) created a model that gave agents a neural net in addition to the more standard receptors and actions; their model is noteworthy in that it allows numerous strategies to emerge through a process of evolution. In contrast, the model I propose demonstrates the emergence of respect for property rights under much simpler circumstances, suggesting more complex models are over identified.

3.0 Why an Agent Based Model?

At this point one might ask what the benefit is of using an Agent Based Model (ABM) instead of simply standard game theory. First and most importantly, an ABM allows for agents with limited information, unable to fully strategize. Agents also are able to learn and adapt over time based on changing conditions, although the agents in this model are zero intelligence agents (see Appendix 1) and so are adopted by the environment instead of adapting per se.[2] Further, an agent based model allows us to find path dependent behaviors that emerge over time, rather than the expected outcomes, due to differing agents repeatedly playing the game with different opponents. This allows the model to demonstrate how agents can be adopted based on their tendencies towards coordination with an ever evolving environment.

My model can also examine how likely particular outcomes are and thus what conditions are more likely to lead to a property rights coordination, even when the agents are not fully knowledgeable about the world. This can illuminate why property rights norms, or at least norms that look a lot like property rights, emerge in so many parts of nature, and why we should expect humans to act within these norms as part of their nature.

As a last point, using an agent based model allowed for other variations on the classic hawk-dove game. One was a musical chairs type hawk-dove game where the winners did not get the reward immediately, but rather reaped the benefit of being an owner later. This model represented competition over an investment good such as a field or fruiting tree that only produces value at a certain point in the future. Another variation allowed for the possibility of three-way interactions among agents within a single game, alongside the standard two player interactions. Every interaction had a set probability of involving a third agent, increasing the value of coordination to avoid violence. As it turned out, however, neither variation provided much useful insight, and so have been left out of the current discussion.

---

[2] See (Alchien, 1950) for a discussion on the distinction between adaptation and adoption of agents within an evolutionary system. In this model, the agents themselves are incapable of learning, but the population of agents does evolve over time.

4.0 The Model

After considering a variety of approaches, I decided to examine a relatively simple model that emphasizes the evolutionary emergence of property rights using the NetLogo (Wilensky, 1999) testbed. The model focuses around agents competing for resources in an evolutionary context. Agents search about for a resource stock for which there is zero marginal benefit for more than one stock, such as a stand of berry bushes or a newly killed animal, etc. Agents come across these resources randomly, but not alone; another agent without its own resource can challenge for any currently unowned resource. During this encounter, agents will individually and simultaneously choose to either fight the other agent or back down from a fight, in the manner of a hawk-dove game. The agents' propensities to fight or back down, to play hawk or dove, are a built in trait of the agents, each having a certain probability of playing hawk depending on the situation. If both agents choose to fight they will harm each other, and play another round of the game until one or both back down. If one chooses to fight and the other backs down, the fighter becomes the owner to the resource and gains the benefits thereof, while the other flees. If both choose not to fight they are assumed to share the resource, each getting some portion of the resource and becoming owners in it.

Over time, agents metabolize their resource stores they have built up and agents that run out of resources die, along with agents mortally wounded by fighting. Agents who have built up a sufficient stock of resources may reproduce, spending some of their resources to create a new agent similar to themselves and sharing their remaining resources with their offspring[3]. In this way, as the model progresses the agents with more effective strategies around their mix of hawk/dove plays will produce new agents that cleave to their parent's strategies but with change

[3] This consumption of resources to reproduce also limits the absolute advantage that agents can gain due to resource stockpiling, discussed later.

and the possibility for improvement. Over time the society should settle into a norm where agents display "expectations" regarding other agents' actions and coordinate accordingly; that is, we would expect that a particular set of probabilities that complement each other will emerge. Three basic non-agent parameters affect this process:

D: Damage agents inflict during combat

F: Food payoff received

p: Percentage of the food payoff received by agents after playing dove/dove.

The payoff matrix for the hawk/dove game is shown in Table 1.

Table 1: Payoff Matrix for the Hawk/Dove Game

	Hawk	Dove
Hawk	*-D / -D*	*F / 0*
Dove	*0 / F*	*F * p / F * p*

In all the model runs, damage and the food payoff range from 1 to 5 in units of 1, and the percentage of the food payoff ranges from .1 to 1 in units of .1. Test runs with different parameters outside these ranges did not substantially alter the outcomes.

4.1 Why Random Behavior?

It has been noted that the agents do not make a decision in the classic economic sense; there is no optimization around costs and benefits. The model also does not use the standard game theoretic approach of solving for a mixed-strategy Nash equilibrium (MSNE) result based on full knowledge of the relevant variables for both parties. It is a fair question then to ask, where is the economics? Where is optimization, or at least learning? The short answer is: at the system level, not the individual level.

The agents themselves do not optimize or learn, yet patterns of behavior emerge over time as the system calibrates around the agents, the population of which changes over time. In other words, the environment changes based on the behavior of the agents, behaviors that are selected for based on the environment; optimization occurs without an optimizing agent, through feedback and selection.

To give a more thorough defense of the decision to choose random behavior with reference to hawk and dove, it is necessary to discuss early versions of the model that used a more standard optimization algorithm. See Appendix 1. For our purposes here, it is enough to think of agents as solving for mixed strategies but having individual biases that push the probabilities in one direction or another. The great virtue in this design, other than elegance, is that agents need not have elaborate strategies or comprehensive knowledge of other agents and the world to function. Only two bits of information are necessary to deciding their behavior: the incumbency status of themselves and their opponent. This feature is highly desirable as it means the model can be applied to humans with our wildly different perceptions of reality and strategy as well as to butterflies and other insects that seem incapable of forethought and planning.

4.2 Pseudo-code and Process

4.2.1 Pseudo-code

The basic model is organized as follows:

- Stage 1: Agent Setup

- Stage 2: Two agents are selected at random to interact. At least one agent in the pair must be a non-owner
    - Agents select hawk or dove based on the owner profile
    - Interaction is resolved by strategy choices
        - If hawk/hawk, agents damage each other, reducing their food value
            - If food is reduced to 0 or below, that agent dies, and the opponent becomes the owner and receives a food payoff
        - If hawk/dove or dove/hawk, the agent playing hawk becomes the owner and receives food payoff, while the dove receives nothing
        - If dove/dove, both agents become owners and receive a portion of the food payoff
    - Repeat interaction if agents both chose to play hawk
- Stage 3: Surviving agents metabolize and reproduce if they have sufficient food, and die if food falls to 0 or below
    - Agents consume 1 unit of food. If this reduces their food to 0 or below, that agent dies
    - If an agent has >=30 food, create a new agent
    - Owners have a 75% chance to stop being owners
- Repeat Stages 2 and 3 until there are zero agents, 2000 agents or 750 iterations have passed

There is a variation on the base model, “Endowment” which alters the agents’ starting probability of playing hawk, modeling the endowment effect by increasing the average probability that an agent will play hawk when it is the owner at Stage 1. The effects are detailed below.

4.2.2 Process

Stage 1**:** Agent Setup

The agents used in this model are very simple having only five unique attributes:

1. Ownership Status**:** Whether the agent is currently in possession of a resource
2. Food**:** A running total of resources gathered. Gained when in possession of a resource, lost due to metabolism, damage incurred in combat, or reproduction
3. Probability of Playing Hawk**:** When the agent is an owner facing a challenger
4. Probability of Playing Hawk**:** When the agent is a challenger facing an owner
5. Probability of Playing Hawk**:** When the agent is a newcomer facing a newcomer

At model initiation, every agent is assigned 10 Food and assigned as an owner with a probability of 1%. In the base model, each agent is also assigned a random probability of playing hawk for each of the three situations, evenly distributed between 0 and 1. In the endowment variant, the probability of playing hawk when the agent is the owner is evenly distributed between .5 and 1.[4]

Stage 2**:** Agents Interact

In this stage, two agents are selected at random from the pool of agents. Any two agents may be selected, provided that the second agent is not an owner if the first agent is to represent the zero marginal benefit of an additional food resource.

---

[4] See the Results section for a brief discussion of how this number was chosen.

After the two agents have been selected to interact, each agent chooses to play hawk or dove randomly based on their ownership status and their individual probability of playing hawk in the configuration of ownership they are in. Once the agents have chosen to play hawk or dove the interaction is resolved.

As shown in Table 2, if both play hawk, each agent loses food equal to the parameter value of damage, *D*. Each agent then checks to see if this reduces its food value to zero or below, and if so the agent dies. It is possible for both agents to simultaneously kill each other off, but if only one dies the surviving agent becomes an owner and receives food equal to the food payoff parameter *F*. If neither agent dies the agents will repeat the interaction, selecting hawk or dove and inflicting damage on the other until either one or both agents die or one or both select dove.

If one agent plays hawk while the other plays dove, the agent playing hawk becomes an owner and receives *F* food, while the agent playing dove neither takes damage nor receives food.

If both agents play dove both become owners and each receives some percentage of the food payoff *F*, as defined by the parameter *p*. This changes the nature of the game as at values of $p < .5$ the game is negative sum for any plays other than hawk/dove or dove/hawk, and positive sum for values of $p > .5$ when dove/dove is played.

The process of selecting agents is repeated a number of times equal to the number of agents when Stage 2 was started. For instance, if the population is 300 agents there will be 300 random selections for interactions, even if 50 agents die off during the process. This and the random selection of agents allows for some agents to be selected multiple times in a single iteration while others are not selected at all.

Table 2: Payoff Matrix for the Hawk/Dove Game (Baseline Model)

Baseline Model		
	Hawk	Dove
Hawk	*-D / -D*	*F / 0*
Dove	*0 / F*	*F * p / F * p*

Stage 3**:** Agents Metabolize**,** Die and Reproduce

After the completion of Stage 2, all surviving agents metabolize one unit of food. If this takes their food value equal to or below zero, the agent dies.

After agents metabolize and check for death all agents with at least 30 units of food reproduce. To do so, agents consume 2/3 of their current food value and create a copy of themselves with the same food level. The new agent is never an owner and generates new probabilities of playing hawk by using a normal distribution with a mean of the parent's probabilities and a standard deviation of .1, bound at 0 and 1 inclusive.

At the final step of Stage 3, all agents that are owners become non-owners with a probability of .75, representing the exhaustion of their resource.

After Stage 3, if there are zero agents, 2000 or more agents, or 750 iterations have been completed, the simulation ends. The constraint of 2000 agents is due largely to computational pressures, but tests have suggested that the end state of 2000 agents is very close to the end state of 750 iterations. In either case, any equilibrium that will be reached has been reached for the most part.

5.0 Results

20,000 model runs were executed, sweeping the possibility space by running each combination of the three global variables of damage (D, range 1-5), food pay-off (F, range 1-5) and the percentage of the food pay-off received by agents both playing dove (p, range .1-.9), as well as a Boolean variable for whethe the endowment effect treatment was used. Each combination of variables was run 80 times, for a total of 20,000 runs. I analyze the results two ways, regression analysis and frequency analysis, to understand the general behaviors of the system and how the variable combinations affect the outcome.

5.1 Results: Regression Analysis

To analyze the results of the simulation runs I utilize ordinary least squares (OSL) regression. I define an emergence of property rights as a norm of behavior where interactions between an owner and a challenger (newcomer) results in the challenger choosing to play dove while the owner plays either hawk or dove. Arguably, the owner's decision to play hawk or dove may further define the strength of the rights, but for our purposes, it seems sufficient to focus on the challenger's behavior instead of whether the owner decides to chase the challenger off or share. The question of whether the behavior change in challengers is driven by owner actions or the underlying nature of the world is relevant, however, as well as the difference between the probabilities of the owners and the probabilities of the challengers.

The three regression models then use as regressand average probability of the owner playing hawk (Owner), average probability of challenger playing hawk (Challenger), and the difference between the averages (Diff.I.C.). The regressors are the three key parameter values, *D* (damage from combat) with a range of 1 to 5 inclusive, *F* (food payoff) with a range of 1 to 5 inclusive, and *p* (percentage of the food payoff from dove/dove plays) with a range of .1 to .9 inclusive. In addition, Boolean dummy variables are added to compare the effects of using the variant treatment Endowment, with values of 1 if used and 0 if not.

A total of 20,000 model runs were generated, however only the 13,690 that did not end in universal agent death were considered in the regressions due to the lack of data on the explanatory variables. (The features of runs ending in universal agent death are examined later in the discussion section.) The results of the OLS regression are shown in Table 3.

Table 3: Regression Results by Treatment

	Owner coefficient	Challenger coefficient	Diff.I.C coefficient
D	-0.027	-0.051	0.025
F	0.022	0.049	-0.027
p	-0.297	0.072	-0.369
Endowment	0.303	-0.145	0.448
constant	0.641	0.256	0.385
R-sqr	0.683	0.498	0.552
dfres	13685	13685	13685
BIC	-17892.8	-21648.9	-2348.2

Due to the very large data set, all coefficients are highly statistically significant with minor standard errors, thus I omit measures of significance here.

Examining the coefficients for various regressors, the first thing to notice is that the constant for the average probability of the owner playing hawk is over two times that of the challenger. This suggests a common norm of owners more frequently playing hawk overall compared to challengers, a suggestion supported by the frequency analysis below.

The effects on behavior of the three of the world parameters (damage from combat, food payoff, and percentage of the food payoff received from sharing) have roughly the direction one might expect with higher cost combat tending to reduce the frequency of risking combat and higher food payoff tending to increase the frequency. A surprising result is that while a better payoff for sharing decreases the likelihood of choosing to fight for owners, it actually slightly increases the likelihood for challengers. This somewhat counterintuitive result, perhaps, reflects the challengers taking advantage of less aggressively defensive owners. In any case the non-symmetrical outcome is interesting in itself.

The Endowment treatment offers the greatest magnitude changes of all with a full 30% increase in owner choice of hawk and a 14.5% reduction in challengers playing hawk. Bearing in mind that the change in the average probability of the owner playing hawk was 25%, from .5 to .75, this reinforcement effect is substantial.

The effects of the regressors on differences between the average plays by owners and challengers are consistent with the model's aims: the coefficients are all the Owner coefficients less the Challenger coefficients and move in the expected directions. This suggests that the agents are interacting and the system is adapting behaviors, as opposed to agents behaving entirely randomly or in some strange strange fashion contra to the model. The very large increase in the difference with the Endowment treatment shows clearly how the endowment effect can create a Schelling point for expectations, moving away from the tendency of the model to rest around mixed Nash equillibriums of approximately 50% probability to play hawk. The frequency analysis demonstrates this quite clearly.

5.2 Results: Frequency Analysis

The treatment effects are most visible in the frequency analysis. Figure 1 presents a histogram of the difference in the average probability of playing hawk between owners and challengers per run sorted by the treatment variants. The base model alone often results in owners and challengers playing hawk with nearly the same probability, as shown in Table 4, although the standard deviation is quite large, with a small modal spike near .8 and a tiny one near -.8. This demonstrates one of the interesting tendencies of the model: random events can produce fairly strong norm formation, and those norms nearly always favor the owner. Even when the difference is relatively small there exists fairly heavy clustering in favor of the owner playing hawk more than the challenger, as the mean of .134 demonstrates.

Instantiating the endowment variant onto the base model, we see a very sharp shift towards owners playing hawk relative to challengers, as might be expected given the initial conditions, to a mean difference of .582 and a standard deviation of .140. The magnitude of this shift in mean difference is quite important as it marks a very strong shift in challenger behavior. The shift is relatively consistently born out, but not perfectly born out, as Table 5 demonstrates. The minimum difference is still negative, with owners' probabilities shifting from a mean of .75 to .304 in the minimum case. Still, the Endowment treatment's effects are to create a very strong tendency towards challengers respecting owners' positions.

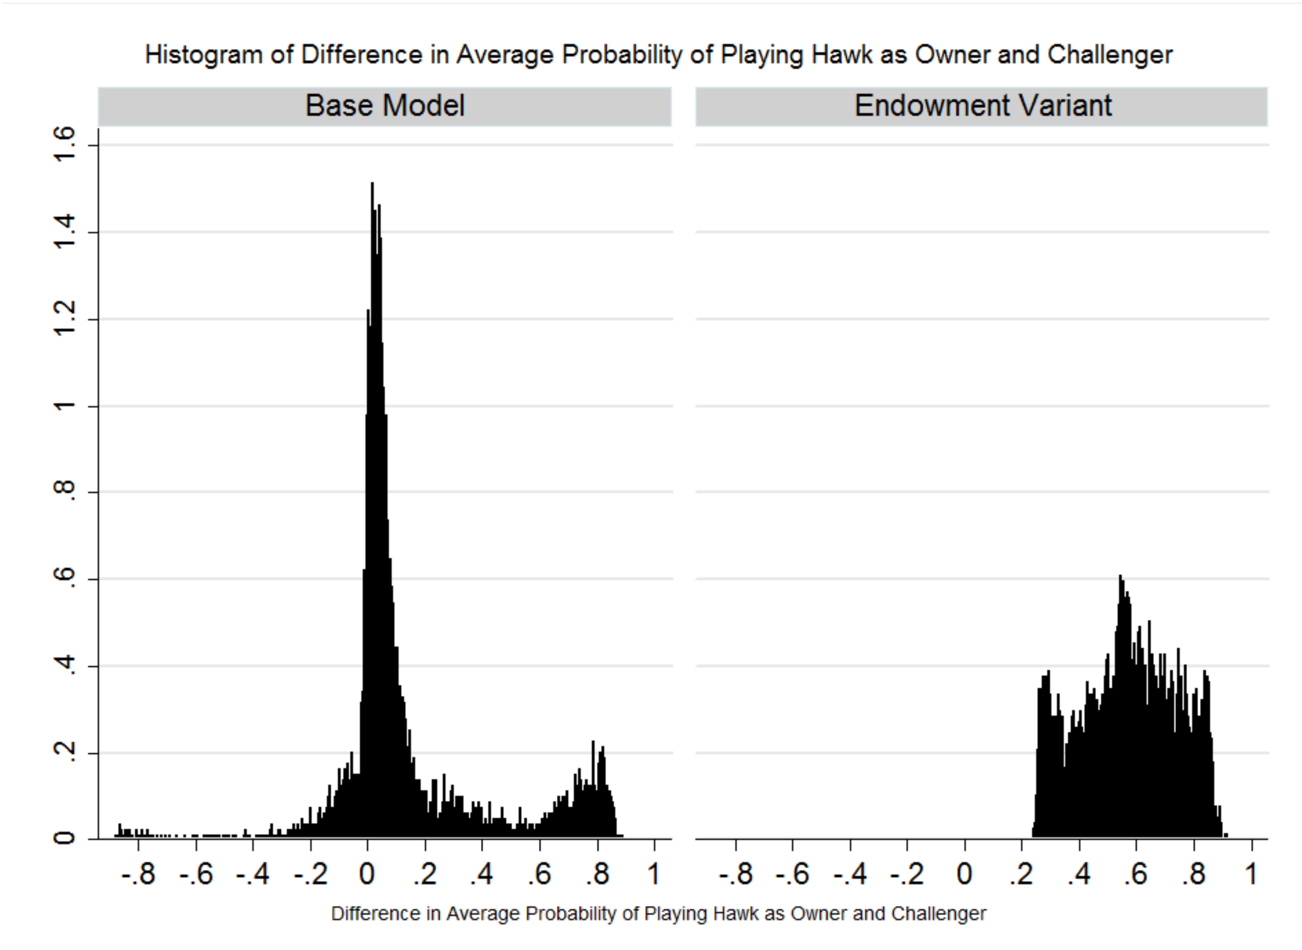

Figure 1: Difference in Average Probability of Playing Hawk as Owner vs. Challenger

Table 4: Standard Deviation and Mean

Base Model		Endowment Variant	
Mean	Std. Dev.	Mean	Std. Dev.
0.134	0.311	0.582	0.140

Table 5: Summary Statistics of Simulation Outcomes

	Base Model		
	Owner	Challenger	Difference
mean	0.460	0.330	0.134
sd	0.210	0.1632	0.3106
min	0.038	0.036	-0.896
max	0.934	0.937	0.861

	Endowment Variant		
	Owner	Challenger	Difference
mean	0.772	0.189	0.582
sd	0.077	0.106	0.140
min	0.305	0.014	-0.165
max	0.947	0.515	0.880

As previously noted, a number of model runs resulted in universal agent death. I excluded these from the regression analysis because they do not produce usable data. I can, however, examine what causes these population crashes. Table 6 shows a probit regression for the agent population against the world parameters and treatment variants. As expected, increasing the damage from combat increases the likelihood of a population crash, but not by a great deal, and is dwarfed by the effects of food payoff and percentage of the food payoff achieved by sharing. Interestingly, the Endowment variant increases the probability of a crash by a small amount, increasing the number of runs that do not reach max agents by approximately 13.5%, and the number of population crashes by ~16.5%. Table 7 shows the effect of the Endowment variant on crashing population.

Table 6: Probit Regression for Population Crash

D	0.357
	(22.81)**
F	-2.597
	(52.92)**
p	-6.487
	(48.60)**
Endow	0.153
	(3.82)**
cons	6.242
	(46.67)**
N	20,000

* $p<0.05$; ** $p<0.01$

Table 7: Effect of the Endowment Variant on Crashing Population

	Endowment Variant	
	Off	On
0 Live agents	2914	3396
>0, < 2000	102	28

5.3 Results: Alterations to the Magnitude of the Endowment Effect

The change in the probability of the owner playing hawk due to the endowment effect is somewhat nebulous to model, being inherently arbitrary due to the random nature of the agents' choices to begin with. Kahneman et al. find willingness to accept an offer is approximately two times willingness to pay (1990). How does that translate into increased willingness to fight for what is our own? That depends on how you model the endowment effect. I assume a reference point model, where once something is possessed its continued possession is valued at zero, but its loss is a negative value. Assuming then that the costs of fighting do not change (which itself is debatable), the hawk/dove framework changes, with playing dove always associated with a loss. Our hawk-dove game then has the features of the game as shown in Table 8 below if the row player is the owner:

Table 8: Baseline Model (Row Player Owner)

	Hawk	Dove
Hawk	*-D / -D*	*0 / 0*
Dove	*-F*e / F*	*(- (1-p) * F * e) / F * p*

Here *e* is the strength of the endowment effect. Although the expected value of playing hawk for the row player is now negative (the owner has literally nothing to gain from playing hawk) the expected value of playing dove is also negative, and possibly more so depending on the parameter values. As Gintis (2007) points out, there will be ranges of values for *D*, *F* and *e* where hawk is the dominant strategy for the owner. In this simple case, so long as *F*e* is greater than *D* the owner's dominant strategy is to play hawk. Assuming e is roughly equal to 2, for all cases where the expected cost *D* is less than twice the value of the reward *F*, hawk dominates.

My model, however, does not ask agents to play strategically, but rather finds the strategies that survive. In a sense the model could be calibrated to follow the emergence of the endowment effect, as opposed to the endowment effect's part in the rise of property rights norms. However, with the model as built, implementing a change in agent predisposed behaviors based on an increased willingness to fight requires some amount of judgement in choosing just how much of a change there is. Asserting that hawk was the dominant play for owners and therefore should be chosen 100% of the time seems far too heavy handed, but a fairly large shift seemed reasonable given previous work. After testing various values for the modifier associated with the endowment effect, I settled upon moving the average probability to .75 from .5, at the same time limiting the range to .5 to 1 from 0 to 1, using a formula of

$$Probability\ of\ Playing\ Hawk = random(1 - X) + X$$

where X ranges from 0 to 1, 0 being no effect, and 1 being always play hawk. Setting X to .5 seemed like a reasonably focal decision, and did not alter the overall shape of the distribution too drastically. Figure 2 below shows outcomes of runs with various values of X. As demonstrated in Figure 3, even small changes in the base propensity to play hawk can have fairly significant changes in the results. Changing the probability of playing hawk when the agent is an owner from a range of 0-1 with a mean of .5 to a range of .1-1 with a mean of .55 causes a significant shift in the probability of a challenger playing hawk over time, for instance (see Figure 3). The histogram for challenger probabilities shifts to the left much more quickly than the probability of playing hawk for owners increases. This demonstrates that a small change in owner behavior can very effectively drive large changes in the emergence of a property right norm.

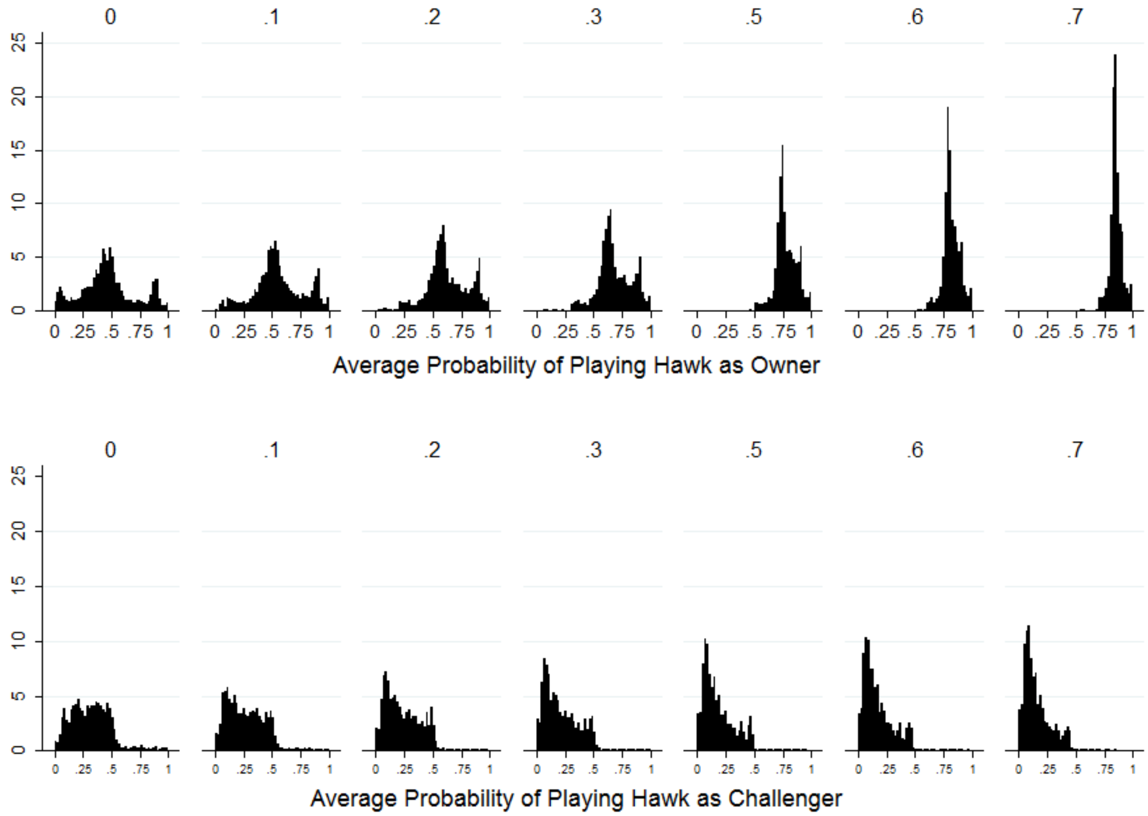

Figure 2: Histograms of Run Outcomes by Strength of Endowment Effect *X*

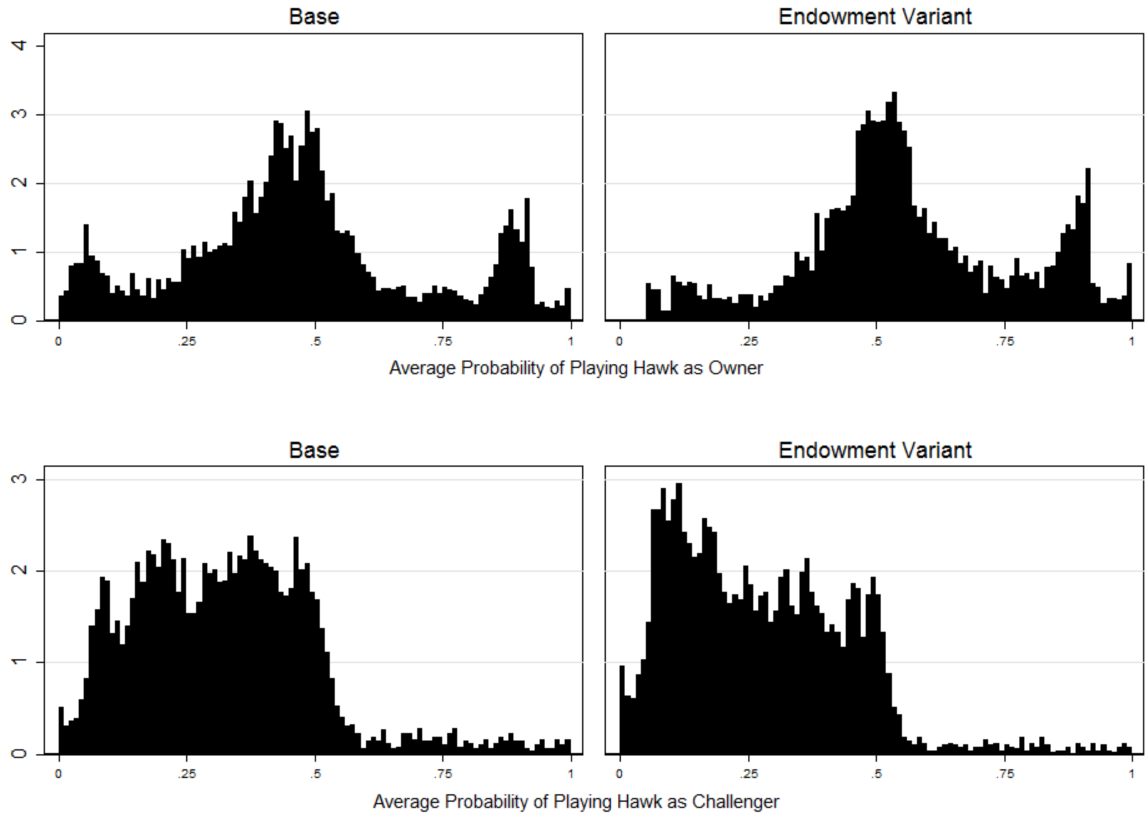

Figure 3: Histograms of Average Probability of Playing Hawk for Owners and Challengers with Endowment Effect *X* = .1

This result is important to the emergent story here. A very small change in the likelihood of owners to play hawk drives a much sharper change in the likelihood of challengers to play hawk over time. This suggests that a relatively small endowment effect is sufficient to achieve social norms where violent conflict is comparatively rare. As Figure 2 shows, the effect is not symmetrical, but far more pronounced on the challenger side, suggesting that observed behaviors such as altruistic punishment[5] might actually benefit the actor more than initially thought as well as having a very large effect on the overall societal norms.

[5] See (Fowler, 2005).

6.0 Conclusion

This paper considers property rights as a norm of challengers refraining from challenging owners for resources. Using this definition I test several different regimes to study the circumstances under which such a norm is likely to arise. By comparing a baseline model against models which asymmetrize decision making (the endowment variant), I am able to uncover circumstances under which property rights are likely to arise. While this norm can occasionally arise under any of the model variants, an endowment effect makes such norms much more likely to obtain, as well as other conditions that raise the cost of fighting.

Most importantly, the property rights outcome requires no strategic consideration on the part of the agents themselves. Individuals do not need farseeing wisdom, complete information, or a complex strategy to settle upon a norm of "mine and thine." Instead, only a slight tendency to fight a little more for incumbents rapidly creates a social norm of ownership rights.

Making violence extremely costly relative to the gains from violence increases sharing of resources and respect for ownership, however it also increases the probability of a population collapse, though this result is likely quite sensitive to the combinations of F and p. This carries potential implications for policy across a variety of fields, such as the choice of rules regarding self-defense in the home for instance, implications worth future exploration.

The most crucial finding is the large change in the average willingness to challenge that a small increase in the endowment effect causes over time. My model and simulation runs suggest that self-help remedies to enforce property rights even in relatively small amounts drive large changes in challenger behavior; a small change in owners' willingness to stand up for themselves makes violent taking a much less successful behavior evolutionarily. Owners need not be willing to fight to the end for their goods, but rather merely need to be slightly more willing on average to fight to make theft not worth the trouble. Violent taking lead to trouble, and is actively selected against by the environment.

Significantly, the endowment effect, at least in my model, does have a dark side: it tends to result in higher likelihoods of everyone dying. Whether this feature is an artifact of the model or is in fact inherent in the endowment effect is an important question worthy of further research. If in fact the endowment effect, while increasing the likelihood of property rights also increases the likelihood of a population crash, important practical questions, as well as moral and ethical questions, arise. These questions are beyond the scope of this paper, but further models could be developed to study whether it is worthwhile for all societies to develop property rights, when, and for what sorts of resources.

Appendix 1: Discussion of Agent Structure

The decision to have agents make their choice of playing hawk or dove based on a random, unchanging probability fixed at birth, deserves some discussion. After all, agents themselves are not strictly optimizing in every instance, and in some sense can hardly be said to make a decision at all. However, once perfect information of the situation and outcomes is off the table, the difference is actually very small.

Initially, the model used agents with a wide array of attributes and a very different distribution of resources. Agents could be stronger or weaker, metabolize food faster or slower, see greater or lesser distances, and move more or less quickly. Resources were distributed over an area, and agents had to actively search for them, sometimes encountering other agents that might challenge them. Upon initiation of a confrontation, the participants would engage in a hawk-dove game, with their strengths affecting the cost of fighting relative to their store of reserves, and some random value of the resource. If the interaction is more than a one off game, repeating if both choose hawk, there is quickly a complication, one that spawns many others.

It is not sufficient to find the mixed-strategy Nash equilibrium (MSNE) of just one game when the decisions of the players can affect the length, but an agent has to consider their ability to endure a drawn out conflict as compared to their opponent. For instance, if two agents are equally matched in terms of strength, but one has reserves enough to last through three rounds of combat while the other only has enough for one, the agent with more reserves is guaranteed victory so long as he decides to fight. When the situation is closer, however, the decision begins to depend not just on one agent's perception of the situation, but also the other's. When this perception is known and shared (i.e. all knowledge is open and available), this is not a problem, as agents quickly have a dominant strategy set of fighting to the end if the rewards are greater than the costs, or capitulating. No MSNE needs to exist unless the opponents are equally matched.

So far so good, except that the evolutionary process has little of interest to act upon, and the model doesn't really address the limited knowledge issue that agent based models are so effective at engaging, two problems that are in fact linked. In terms of evolution the pressure works on the margins of differences in the model: strength, metabolic rates, ability to find food, etc. The evolved ability of agents to extract resources from the environment improves, and in a somewhat interesting way on occasion, as some specialize in finding unclaimed food more effectively and others specialize in taking claimed resources. Generally, though, the end state is agents with very low metabolism that can find food effectively and all have strength such that fighting each other for the food resources is largely pointless. The reason is that we have accidentally fixed the most interesting aspect of culture, institutions and decision making: the ability of agents to deal with uncertainty. The agents have perfect knowledge of the situation and perfect ability to compute the optimal results, and so the only possibility for error lies in their random selection in the MSNE when equally matched. In effect they are still not making decisions, but calculating the proper strategy and executing, and doing so perfectly. They cannot learn or form better strategies, because they are effectively already perfect.

In order then to create interesting evolutionary results, I have to introduce limited information so that the agents can make meaningful mistakes around their perceptions. The perception question, and just how much knowledge there is open or otherwise, creates a problem. One example of this problem was already built into the model: agents did not know how likely they were to find food again if they back down from a conflict. Let us assume an agent was within one or two activations of starving due to metabolizing if it did not acquire some food.

The decision to fight for any given resource it comes across must be colored by this fact; if food is plentiful it doesn't make sense to fight for fear of violent death, but if food is scarce and the agent is likely to starve before finding more the cost of fighting is effectively now zero. The agent might die fighting, but he will surely die if he doesn't, so he might as well give combat a go.

Implementing an agent rule where agents keep track of the moves they make between locating food, and deciding to always fight if they are within a certain probability of starving is not terribly problematic. That change, however, affects how other agents view the encounter, as now if they know everything about the other agent and its relevant variables they know exactly what it will do and the outcome. At that point, of course, we are assuming agents know perfectly the interior physical and mental states of other agents during a conflict, except that they do not know what choice the other agent will play in the hawk-dove game in some situations. This is unsatisfactory, not only because it is divorced from any real interaction we could imagine, but because it creates alternate possibilities that the hawk-dove game does not allow. Imagine a starving agent coming upon another agent with some food. Both are perfectly aware of each other's knowledge, and the agent with the food knows that there is another, smaller source of food nearby, closer than the starving agent's expectations.. The sensible thing would be for the starving agent to just go over to that food once the other agent makes him aware of it. However, that is impossible because the hawk-dove game does not allow that sort of communication. We must now either add a step to the game, that of sharing mental states of the world, or limit the amount of information about each other that agents have available.

Assuming that perfect telepathic communication is not a useful feature of a model (and in fact presents some other strange problems), we must have agents with imperfect perception.

Clearly, two agents will have some idea of the features of other agents they meet, but it must be imperfect, and randomly so. Perhaps they can only interpret strength in a general way, ordinally instead of cardinally. Likewise, they may only have a rough estimate of how long their opponent is likely to be willing and able to fight, and certainly a rough estimate of how hungry he is. Then again, it might be useful to also have the ability to bluff, to signal different attributes than one actually has. Then it becomes useful to be able to see through these bluffs, and to perceive how your opponent perceives you. There becomes a spiral of various capacities for accurate perception of reality that can affect the outcome, the selection of a mixed strategy weighted on expectations. Knowing that there will be incorrect perception, agents should weigh their decisions based on past experiences with agents similar to the current contender, requiring either a number of rounds with little evolutionary effect (analogous to human primary and secondary school perhaps) before the "real" conflicts begin, or the creation of a history for each agent. We might even attach a concept of bias one way or the other. Perhaps a particular agent always underestimates an opponent's willingness to fight, or over estimates his strength. That would change things, and considering that the root idea is to examine the effects of a valuation bias (the endowment effect) on cultural norms, seems necessary at some level.

At this point, however, the MSNE calculation starts to depend very heavily on the internal characteristics of a given agent and very little on the external world, and particularly on the different weights each agent puts on the various aspects of their decision making under uncertainty. Agents will change their behaviors slightly based on changes in the real world, but only a small amount. Their internal states drive their decisions, not an uncertain world.

In the end, skipping all the calculations of various estimates and biases involved in a situation and instead using similar agents while simply stating "This agent plays hawk against a challenger 64.6% of the time" produces observationally equivalent results, without the vast array of variables needing calibration and testing. Whether one agent or another is better at calculating mixed strategies is largely irrelevant. The behavior of the system of agents is what is interesting here, not the behavior of individuals per se, as the culture is not embodied in the decisions of any one agent. To put it another way, culture need not be a result of any individual's strategy, but selection of what behaviors work given a changing social environment. The success of an individual agent is a subset of the success of a behavior set in evolutionary terms.

Works Cited

Alchien, A., 1950. Uncertainty, evolution, and economic theory. *The Journal of Political Economy,* pp. 211-221.

Bowles, S. & Choi, J. K., 2003. The first property rights revolution. In: *Workshop on the Co-evolution of Behaviors and Institutions.* s.l.:Santa Fe Institute.

Buckle, S., 1991. *Natural Law and the Theory of Property: Grotius to Hume.* 2nd ed. New York: Oxford University Press.

Burtsev, M. & Turchin, P., 2006. Evolution of cooperative strategies from first principles. *Nature,* 404(7087), pp. 1041-1044.

Davies, N. B., 1978. Territorial Defense in the Speckled Wood Butterfly (Pararge Aegeria): the Resident Always Wins. *Animal Behavior,* Volume 26, pp. 138-147.

Demsetz, H., 1967. Toward a theory of property rights. *The American Economic Review,* pp. 347-359.

DeScioli, P. & Wilson, B. J., 2010. *Mine and Thine: The Territorial Foundations of Human Property.* [Online]
Available at: https://chapman.edu/ESI/wp/Wilson_FoundationsProperty.pdf
[Accessed 23 December 2013].

Fowler, J. H., 2005. Altruistic punishment and the origin of cooperation. *Proceedings of the National Academy of Sciences of the United States of America,* 102(19), pp. 7047-7049.

Friedman, D., 1994. A Positive Account of Property Rights. *Social Philosophy and Policy,* Volume 11, pp. 1-16.

Gintis, H., 2007. The evolution of private property. *Journal of Economic Behavior & Organization,* 64(1), pp. 1-16.

Hoffman, E. &. S. M. L., 1993. Willingness to pay vs. willingness to accept: legal and economic implications. *U.L.Q.,* 71(59).

Huck, S., Kirchsteiger, G. & Oechssler, J., 2005. Learning to like What You Have: Explaining the Endowment Effect. *The Economic Journal,* 115(505), pp. 689-702.

Jones, O. D. & Brosnan, S. F., 2008. Law, Biology, and Property: A New Theory of the Endowment Effect. *William and Mary Law Review,* 49(6), pp. 1935-1990.

Kahneman, D., Knetsch, J. L. & Thaler, R. H., 1990. Experimental Tests of the Endowment Effect and the Coase Theorem. *Journal of Political Economy,* 98(6), pp. 1325-1348.

Krier, J. E., 2008. The evolution of property rights: a synthetic overview. *U of Michigan Law & Economics, Olin Working Paper No. (08-021).*

Leeson, P. T., 2008. Coordination without Command: Stretching the Scope of Spontaneous Order. *Public Choice,* 135(1/2), pp. 67-78.

Leeson, P. T., 2008. Social Distance and Self-Enforcing Exchange. *The Journal of Legal Studies,* 37(1), pp. 161-188.

Leeson, P. T., 2009. The Laws of Lawlesness. *The Journal of Legal Studies,* 38(2), pp. 471-503.

Leeson, P. T., Coyne, C. J. & Boettke, P. J., 2006. Converting Social Conflict: Focal Points and the Evolution of Coordination. *Review of Austrian Economics,* 19(2-3), pp. 137-147.

Smith, M. J. & Parker, G., 1976. The logic of asymmetric contests. *Animal Behaviour,* 24(1), pp. 159-175.

Sugden, R., 1989. Spontaneous Order. *The Journal of Economic Perspectives,* 3(4), pp. 85-97.

Thaler, R., 1980. Toward a positive theory of consumer choice. *Journal of Economic Behavior and Organization*, 1(1), pp. 39-60.

Tullock, G., 1994. *The Economics of Non-Human Societies.* Tucson: Pallas Press.

Wilensky, U., 1999. *NetLogo.* Evanston(IL): Center for Connected Learning and Computer-Based Modeling, Northwestern University.

Yee, K. K., 2003. Ownership and trade from evolutionary games. *International Review of Law and Economics,* 23(2), pp. 183-197.