Testing and Refining Strategic Decision Theory

by

Alex Roomets

A Dissertation Submitted to the Faculty of the

Department of Economics

In Partial Fulfillment of the Requirements
For the Degree of

Doctor of Philosophy

In the Graduate College

The University of Arizona

2011

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation

prepared by Alex Roomets

entitled Testing and Refining Strategic Decision Theory

and recommend that it be accepted as fulfilling the dissertation requirement for the

Degree of Doctor of Philosophy

____ Date: 5/4/2011
Martin Dufwenberg

____ Date: 5/4/2011
Price Fishback

____ Date: 5/4/2011
Mark Stegeman

____ Date: 5/4/2011
John Wooders

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

____ Date: 5/4/2011
Dissertation Director: Martin Dufwenberg

Statement by Author

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at The University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgment of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the copyright holder.

Signed: Alex Roomets

Acknowledgements

I would like to thank all the people who helped bring this dissertation into being. In particular I'd like to thank Martin Dufwenberg and the members of my committee for feedback and guidance, and of course Carl Kitchens for his contributions to the third chapter.

Table of Contents

[1] Joint work with Carl Kitchens

List of Figures

List of Tables

Abstract

In many important economic situations, decision makers influence each other. The subject of game theory offers a mathematical framework to describe such strategic interaction. This dissertation focuses largely on helping to answering the question, "What will someone do in a particular strategic situation?" In order to do this, it is useful to interweave theory with experimentation. After all, observation of what people really do is a necessity when attempting to create models of what people really do. At the same time, theory can help significantly when formulating interesting hypotheses to test. The chapters in my dissertation illustrate this interweaving of theory and experiments.

Introduction

In order to fully appreciate this dissertation it is important to understand the distinction between strategic and individual decision theory. While there is a similarity between strategic decisions and decisions under uncertainty, the two are distinct. When analyzing decisions under uncertainty it is common to exogenously fix the structure of the uncertainty being faced. On the other hand, the structure of the uncertainty in strategic situations is endogenous. This means that a large part of studying strategic decisions is trying to understand the nature of the inherent uncertainty of the situation.

Another difference between strategic and individual decision theory can be seen when looking closely at the underlying question of a line of research, and the types of assumptions being made. For example, the assumption of rationality can be either appropriate or inappropriate based on the question. Decision theory often asks the question, “What should someone do in a particular situation?” Being rational is a reasonable answer to such a question and so the assumption of rationality makes sense. Because of this, economic theory has often done a good job of giving advice to independent economic agents in non-strategic settings. On the other hand, when asking the question, “What will someone do in a particular situation?” an assumption of rationality only makes sense if the "someone" is rational, which is rarely the case. Of course, when looking at topics such as mechanism design, and strategic decision theory in general, the second question is more important. This is because to figure out what an economic agent should do in a strategic situation, one must first understand the behavior of the agents he is interacting with. It follows that strategic decision theory, whether proscriptive or descriptive, must be based on realistic assumptions about behavior.

The first chapter, "On Limited Foresight in Games," is an introduction to limited foresight models that are designed to relax the rationality assumption. In particular, the models assume that a player's ability to foresee possible future contingencies is limited. Perhaps the most obvious example of this limitation is in games like Chess, where even the best players in the world can only see a relatively small number of moves ahead. In such cases models must try to assign value not just to final outcomes but to intermediate positions in hopes of accurately describing decision making.

The limited foresight models developed in the chapter are applied to a number of economic and game theoretic applications including the game of 21, a centipede game, and bargaining. Commonly observed behavior in these scenarios is consistent with limited foresight even when not consistent with backward induction. Specifically, in 21 early play is likely to be random despite a backward induction solution resulting in a win for one player, as few people can see far enough to distinguish one action from another in any meaningful way. In the centipede game agents with limited foresight are likely to continue the game for multiple rounds despite a backward induction solution resulting in immediate termination of the game. In a stylized bargaining game players may propose lopsided splits before eventually agreeing to a depreciated fair split if they have limited foresight.

The second chapter, "A Mechanism to Elicit Continuation Values in Experiments," presents a mechanism to elicit the value experimental subjects' place on continuing a game (continuation value) at any point in games with discrete time and bounded payoffs. Under certain assumptions, the mechanism has two important properties: truthful revelation maximizes subjects' expected payoffs, and behavior in the underlying game is unchanged when implementing the mechanism.

Being able to identify these values should greatly aid in testing and refining the models presented in the first chapter. The models in question currently use values that are based on psychological intuition. Basing values on continuation value data would potentially improve the predictive power of these types of models, and might help refine intuition in the process.

Beyond the connection to the first chapter, the mechanism presented may prove useful in understanding behavior in a more precise way that could help determine the relevance of many theories. Indeed many theories either implicitly or explicitly define continuation values at various points in games. By comparing reported continuation values with those derived by theory, researchers can gain more insights than by comparing actions alone. For example, mixed Nash Equilibria are often thought of as equilibria in beliefs. These beliefs imply a certain continuation value that can be compared empirically to results using the mechanism presented in chapter two.

The final chapter, "Dealing with Eminent Domain," is an example of how economic models and experimentation can be brought to bear on important sociopolitical issues. It is intended to replicate a real-world land-assembly scenario with two regimes (with and without eminent domain). In light of the recent U.S. Supreme Court decision, Kelo vs. New London, there has been a renewed interest in the anti commons, specifically in the context of land assembly. Using a sequential Nash Bargaining model the chapter examines a scenario where a buyer can purchase N identical perfectly complementary goods from N queued sellers. The chapter examines the scenario with respect to two bargaining protocols, one where each contingent price must be agreed upon by buyer and seller, and one where the buyer has an additional option to execute a transaction at a predetermined price for a fee. In the first protocol, theory predicts, given equal bargaining weights, that sellers who are later in the queue will receive lower prices than those earlier in the queue. Using the second protocol, theory predicts that prices should be equal when sellers have equal bargaining weight. An experimental test of these predictions finds evidence that matches predictions in the second protocol, but not the first.

All together the dissertation provides insights that can be used to better understand strategic decision making through experimentation. Perhaps more importantly it provides others with tools that can potentially lead to more insights, better experiments, better models, and ultimately better policy.

Chapter 1
On Limited Foresight in Games

1.1 Introduction

There is overwhelming experimental and observational evidence that humans' ability to use backwards induction is limited. Experiments involving finite games of perfect information such as the game of 21 (Dufwenberg, et. al., 2010) and centipede games (Rosenthal, 1981) almost always fail to support predictions made using backward induction. Furthermore, observational evidence of common finite zero-sum games of perfect information such as Tic-Tac-Toe, Go, finite versions of Chess, etc., shows that wins for either player, as well as ties if possible, occur in non-trivial proportions. Backward induction would predict only one player should ever win these games, or the game should always end in a tie. Backward induction is also often violated in some finite negotiations.[1] For example, if a deal is eventually reached after a drawn out negotiation process, backward induction would suggest that the same deal should have been agreed on at the start, saving both sides bargaining costs. So if humans do not backward induct in these situations, what do they do? How are decisions made in finite games of perfect information when the game trees become large. This paper will introduce a class of models that bound decision makers' ability to analyze the full strategic situation they are in. We call models within this class Limited Foresight Models (LFMs).

The class of limited foresight models is one where agents do not fully consider actions within a strategy for all possible future histories. Instead, agents consider only the elements within "visual range" of the current history, a subset of the complete set of histories of the game. This results in play that is not consistent with backward

[1] In cases where the negotiations could be modeled as a finite game of perfect information.

induction. Generally, the visual range will determine a "visual game." A visual game will be a new game that starts at the current history of the original game in question and only includes histories of the original game that are within visual range. A player will then act as if playing that visual game while at that history. When a new history is reached, a new visual game is created, and the player will act as if he is playing the new visual game. By looking at the initial action of the visual game at each history, play of the original game can be predicted.

These models are attractive because they capture a certain psychological intuition. In particular, they model the intuition that cognitive resources are not infinite. In order to consider the future result of an action taken in the present, cognitive resources must be spent. At some point, as more hypothetical actions are analyzed, cognitive resources run out. Decisions then must often be made without a complete understanding of their strategic importance. By modeling decision making in this way, we hope to better predict the actions of real economic agents that are subject to cognitive limitations.

Our models assume that cognitive resources will be spent in a fashion that begins at the current history, and moves continuously through possible future histories in a manner consistent with the idea of foresight. It is not unreasonable however to assume that a player might start at some beneficial outcome, and try to transform that outcome backward to the current history. Generally, this type of thought process would be difficult to implement as games where limited cognitive resources come into play often have a prohibitively large set of end games. That said, this backward induction concept is not wholly incompatible with our models, as the reader will see when we discuss intermediate valuation.

This paper focuses on the previously stated intuition. To evaluate the outcome of an action or a series of actions, a player must mentally transform the current state into the resulting future state. This transformation draws from a limited pool of cognitive resources. As a result, strategic decisions must often be made under conditions where not all possible future states are known or fully understood.

In the coming sections we define two subclasses of LFMs for finite games of perfect information with discrete action sets, as well as discuss ways of extending such models to infinite games of imperfect information and continuous action sets. Models in the first subclass are called Horizon-Based LFMs (HBLFMs). This class assumes that you can see all possible action sequences out to some "horizon", and nothing beyond that. Models in the second subclass are called Directed LFMs (DLFMs). This class assumes that you must "investigate" possible action sequences before you can see them. You must mentally focus on an action in order to see the consequences the action would have. What is then limited is your ability to investigate.

With each subclass of limited foresight models, there are a number of further assumptions that must be made for a specific model, in order to make a prediction. So, while all HBLFMs will use some horizon to set the visual range at any history, a specific model will have to define how intermediate histories at the edge of that visual range are evaluated, and how actions at histories inside the visual range will be taken. DLFMs are more complicated in that they need not only to evaluate edge histories and predict intermediate actions. They must also predict how available foresight will be directed.

Besides the models in this paper, limited foresight has been employed occasionally in economics literature. Jehiel (2001) looks at repeated games where strategies are formed over some limited forecast range. The model is similar in many respects to our horizon based models. Unlike our models, it is focussed on repeated stage games and flow payoffs. Jehiel does present the concept of a "state of mind," which takes the form of a random variable drawn from some density function. This state of mind is used to value post-horizon play. The paper also uses equilibrium analysis, while our models do not impose equilibrium conditions. Overall, Jehiel's model could roughly fit into our horizon-based class. Kochov (2009) looks at individual decision makers with limited foresight. While the paper is focused on individual decision makers as opposed to games, it includes foresight that is both horizon based and somewhat directed in that subjects ignore low probability outcomes. Wichardt (forthcoming) also looks at a horizon-based limited foresight range, and uses it to put constraints on strategy profiles.

We now present a simple HBLFM to give the reader an idea of the formal structure of LFMs. We will apply this model to the game of 21. We then take a detailed look at HBLFMs, and their applicability to a centipede game and a stylized bargaining game, before moving on to DLFMs, some simple results, and various extensions.

1.2 A Horizon-Based Limited Foresight Model

As this is the first LFM the reader is likely to have seen, our first model will trade off psychological accuracy for simplicity. Furthermore as it is the first HBLFM presented, we will sometimes refer to it as HBLFM#1. It is meant primarily to solidify the concept of an LFM formally and to expose the reader to the necessary notation. Now let's consider a game.

Let $G = (I, H, L, ((u_i); i \in I))$ be a finite extensive form game of perfect information, where I is the set of players, H is the set of histories, $L : H^P \to I$ is the mapping of who moves at each partial history where $H^P \subset H$ is the set of intermediate or partial histories, $u_i : H^T \to \mathbb{R}$ is player i's payoff function where $H^T \subset H$ is the subset of terminal histories. Note that $H^T = H \setminus H^P$.

A natural way to solve this game would be using backward induction or some similar concept. Subgame perfect and trembling hand perfect equilibria are guaranteed to exist, for example. There is a problem, however. It may be psychologically unreasonable, if H is large, to assume that informational assumptions implicit in such models are consistent with reality. In such cases it may be preferable to analyze the game using a limited foresight model.

The first specific example of an LFM will be of the horizon-based subclass. A value ω_i will be assigned to each player as a representation of how many moves ahead that player can foresee. We will then assign the values of visually-terminal nodes (that are intermediate in the original game) as some value β_{ji} (which may depend on certain features of the game) meant to represent each player j's belief about player i's expectation of the value of the game. These additional parameters will allow us to create visual games from each history. Play of the visual games will be assumed to be in accordance with backward induction.

We assign $\omega_i \in \mathbb{N}$ as player i's foresight parameter, and β_{ji} as player j's belief about the value to player i of visually terminal intermediate histories. Furthermore, let the sequence $h = (a^1, a^2, ..., a^{k-1}, a^k)$ be an element of H, where a^k is the action taken at $(a^1, a^2, ..., a^{k-1})$. We say that h' "follows" h (and that h "precedes" h') if the sequence h' begins with the sequence h, and define $H|h \subset H$ as the subset of histories that follow h. For any such $h' \in H|h$, we define $D|h(h')$ (what we'll call the "distance" of h' from h) as the number of actions in h' less the number of actions in h. If $h' \in H|h$ and $D|h(h') = 1$, we say h' "immediately follows" h.

With this additional notation, we now go on to define how to create a visual range. The visual range will be the set of histories included in a particular visual game.

So, for any intermediate history h (where $L(h) = j$), we define j's visual range, $H^V|h$, as the subset of $H|h$ having distance from h less than or equal to ω_j. Since only one player moves at any information set, we leave j out of the notation, but it is important to keep in mind that $H^V|h$ exists only for the player moving at h. We also define a set of "visually-terminal" histories $H^{VT}|h$ as histories in $H^V|h$ that are terminal with respect to the set $H^V|h$ (not necessarily terminal with respect to the set H).

We can now define a visual game for j at h by using the set of visual histories $H^V|h$, and by defining a new utility function that maps from the visually-terminal histories of this set.

We define this particular visual game as $G^V|h = (I, H^V|h, L, (u_i^V|h; i \in I))$. Here $H^V|h \subset H$ as defined above, and $u_i^V|h : H^{VT}|h \to \mathbb{R}$ is player i's payoff function as presently perceived by player j where $H^{VT}|h \subset H^V|h$ is the subset of visually-terminal histories. This payoff function has the following properties:

-if $h^{VT} \in H^T$ of the original game G, then $u_i^V|h(h^{VT}) = u_i(h^{VT})$

-otherwise, let $u_i^V|h(h^{VT}) = \beta_{ji}$.

We have now created a game that player j uses to mentally represent the situation he is in when making a decision at h. We now make the assumption that, given this mental representation, player j uses backward induction to choose an action to take at h.

So, let us now consider all paths of the game $G^V|h$ that are consistent with backward induction. Any history $h'' \in H^V|h$ reached along any of these paths of $G^V|h$ where $D|h(h'') = 1$[2] is then said to be "consistent" with this model and the specific values of ω_i, and β_{ji}.

By looking at all the histories in H, and solving their respective visual games as outlined above, a prediction for the game G can be made. Specifically, a prediction of our model is any h^T for which all preceding histories and h^T itself are consistent.

1.2.1 An Example Game

Consider the game presented in Figure 1.1.

[2] $D|h(h'') = 1$ means h'' immediately follows h.

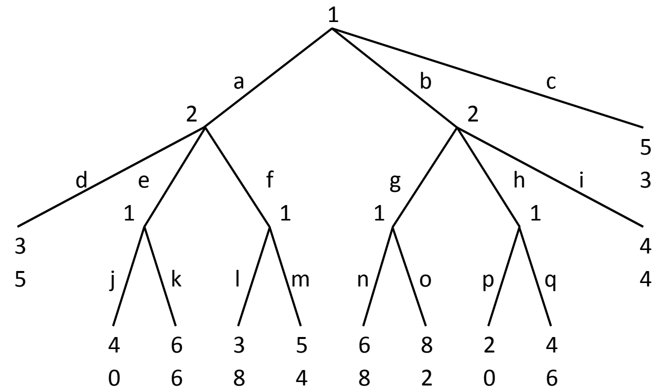

FIGURE 1.1. An Example Game (G)

In Figure 1.2 we show the profile of play generated by backward induction. Note that the path generated by the profile is (a,e,k).[3]

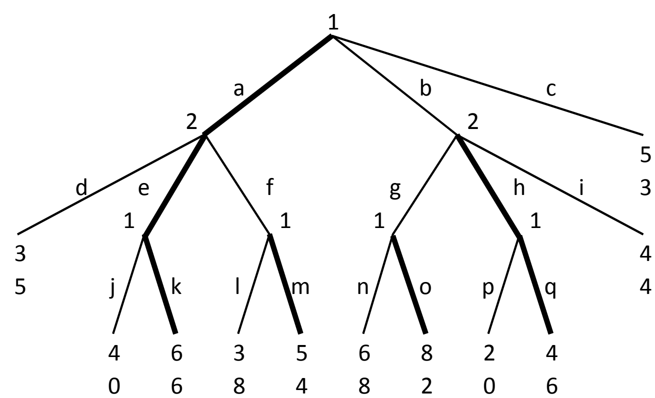

FIGURE 1.2. Subgame Perfection in G

Now let's look at the model presented in this section (HBLFM#1). Assume that $\omega_i = 2$ for both players, in which case the relevant visual games are as shown in Figure 1.3:

[3] Player 1 will take 6 (k) over 4 (j). Player 2 will take 6 (e,k) over 5 (d) or 4 (f,m). Player 1 then will choose 6 (a,e,k) over 5 (c) or 4 (b,h,q).

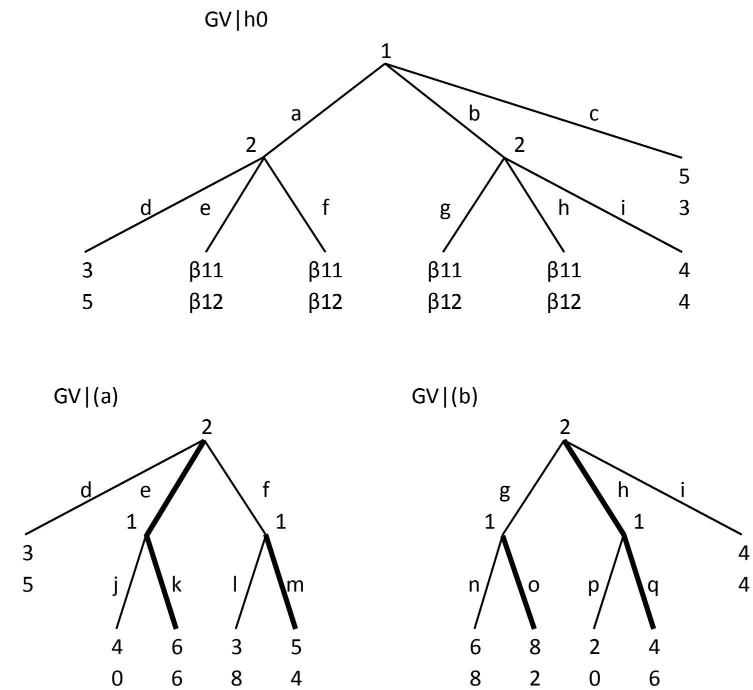

FIGURE 1.3. HBLFM#1

A prediction of the model then depends on β_{11} and β_{12} in the following way:

$$\text{Predicted Paths} \begin{cases} (c) & \text{if } \beta_{11} \leq 5 \text{ or } \beta_{12} \leq 4 \\ (a,e,h) & \text{if } \beta_{11} \geq 5 \text{ and } \beta_{12} \geq 4 \\ (b,h,q) & \text{if } \beta_{11} \geq 5 \text{ and } \beta_{12} \geq 5 \end{cases}$$

Note that the model makes multiple predictions at certain values of β_{11} and β_{12}.

1.2.2 An Application: The Game of 21

In their 2010 paper, Epiphany in the Game of 21, Dufwenberg, Sundaram, and Butler look at the game of 21.[4] The game can be described as two players (A and B) taking turns adding either one or two to a sum, until the sum equals 21.[5] The player who

[4]Cardella also looks at the game of 21 and was kind enough to loan us his experimental data.
[5]Play starts froma sum of 0, and players are not allowed to increase the sum to 22 from 20.

brings the sum to 21 wins (his payoff is π_{win}) while the other player loses (his payoff is π_{lose} where $\pi_{lose} < \pi_{win}$). Although the Dufwenberg et. al. paper is primarily focused on learning (epiphany) and the Cardella paper is interested in imitation learning, the results are interesting in the context of our simple model. The game of 21 is a good candidate for HBLFM#1 because it is a finite game of perfect information with only two possible payoffs. There exists an intuitively natural range for β_{ji}'s (the inclusive set between the two payoffs). The game is also long enough that most players are unlikely to be able to foresee all events, so limited foresight is likely to have an effect.[6] Figure 1.4 shows the game of 6, a shortened version of 21 used in the Dufwenberg et. al. paper where players take turns counting to 6 instead of 21. This game gives the intuition behind 21, which has a prohibitively large game tree to show here. In this version of the game of 6 $\pi_{win} = 1$, and $\pi_{lose} = 0$.

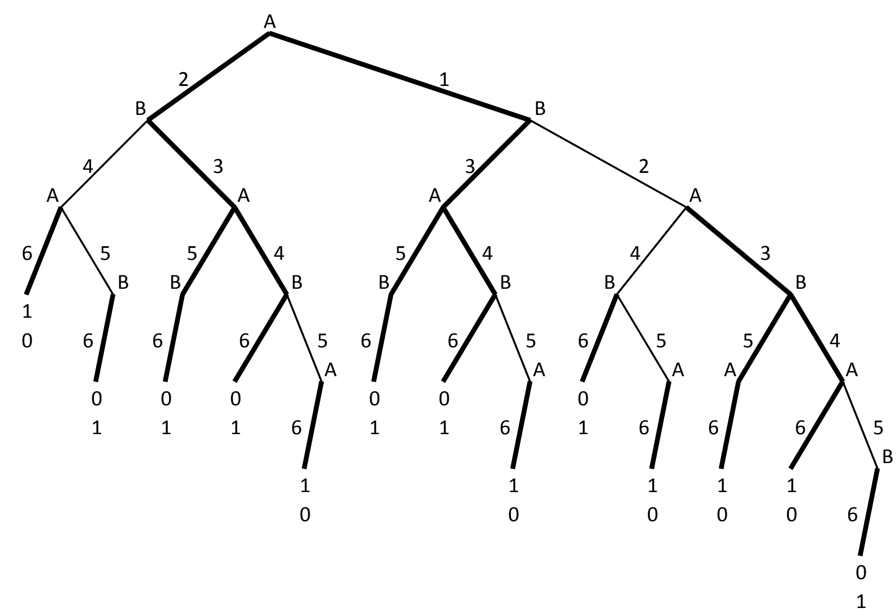

Figure 1.4. The Game of 6

[6]We note that since the game can be solved algorithmically that perfect foresight is not needed to ensure victory.

The game has backward induction solutions that suggests the players should always bring the sum to multiples of 3 when possible. It follows that any strategy where a player brings the sum to a multiple of 3 at some history h weakly dominates the near-identical strategy where the player does not bring the sum to a multiple of 3 at h, and the rest of the strategy is identical. However, experimental evidence from Cardella (2010) shows that inexperienced players[7] chose weakly dominated actions nearly as often as weakly dominant ones when acting near the beginning of the game of 21. Towards the end of the game, players in general were more likely to make the weakly dominant actions. This result is summarized in Table 1.1.

	3	6	9	12	15	18	21
Percent	0.5714	0.5	0.5946	0.5294	0.7273	0.6486	0.9688
Fraction	16/28	18/36	22/37	18/34	24/33*	24/37*	31/32*

* - Significantly different from 1/2 at 5% level

Table 1.1. Rate Subjects Chose Respective Numbers When Possible in 21

This type of behavior is consistent with the simple model whenever $\pi_{lose} < \beta_{ji} < \pi_{win}$ for both players and foresight ranges from 1 to 6 for most players. If we were to reasonably assume that indifference leads to equiprobable mixing, then this behavior would not just be consistent but predicted. Looking at the visual games for the game of 6 with $\omega_i = 2$ for both players (Figure 1.5), we can see the above intuition more clearly. The game tree on the left represents the first move of the game of 6. Note that (1) and (2) are both consistent. The middle tree represents player B's visual game at (1). Again both $(1,3)$ and $(1,2)$ are consistent. The game tree on the right represents player A's visual game at $(1,2)$ at this point we see that only $(1,2,3)$ is consistent. So in the game of 6 with foresight of 2, players would not necessarily choose "correctly" at (1) but would at $(1,2)$ or (2) and afterward.

[7] Inexperienced players refers to those playing the game for the first or second time.

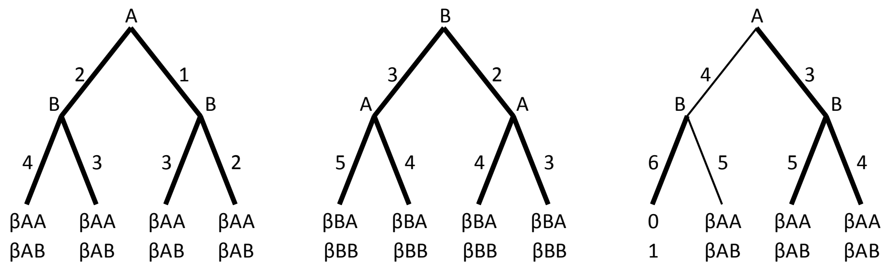

FIGURE 1.5. Visual Games of 6

1.3 Alternative Specifications

HBLFMs share the trait that they all define a visual range of some distance. Thus all histories within that distance of the current history will be part of the visual game. Because HBLFMs assume that players consider all possible actions at all histories within visual range, they are most appropriate when looking at games with small action sets at each history. When action sets get larger, it would make less sense to assume players could consider every contingency, even for a relatively short horizon.

Although the histories contained within the visual range are easily defined by the horizon, the visual game must have values associated with visually terminal nodes. True terminal histories already have a value, but visually terminal histories that were intermediate in the original game must be assigned values in some other fashion (in our first model the values are determined by the β_{ji}'s). There are countless ways to do this, and that is one of the criteria that separates one HBLFM from another.

1.3.1 Visually Terminal Valuation

First remember that a history is visually terminal if it is a terminal node of a visual game. This includes both intermediate histories of the original game that are visually terminal, as well as any terminal nodes of the original game that are within visual range. In the latter case the value of such a history is simply the value of that history in the original game. For the former, which we call visually-terminal intermediate histories or VTI histories, a value must be specified in some other way. One possibility would be to assign all these histories the same value, as we did in the first model. This value could be the median payoff, the minimax value, an exogenous psychological parameter like an aspiration level, or some other value. Another possibility would be to assign a value based on possible outcomes reachable from that history. One could take a weighted average of possible outcomes, use the median, draw a random value, etc. In fact, if one were to use some backward inducted value of that history, then models would generally make predictions consistent with backward induction.[8] Using backward inducted values near the end of games, but not near the beginning, could capture the idea of people using their cognitive resources to backward induct from possible end nodes, as opposed to using it on foresight.

Interestingly, researchers looking into artificial intelligence[9] have looked at intermediate valuation. There is a significant amount of literature devoted to creating an artificial intelligence that can play games like chess. A chess playing computer faces many of the same issues that real players face. In particular, unlimited backward induction is not possible and so intermediate histories must be valued in some fashion. Although many of the ways that computers can make these evaluations would be implausible for humans, there is certainly value in taking a short detour to take a closer look at this literature.

Chess playing computers and more generally artificial intelligence have been studied for roughly fifty years. Early work focussed separately on symbolic manipulation largely based on logic, and sub-symbolic methods largely based on reinforcement and

[8]Predictions are "generally" consistent because inconsistencies can arise when there is indifference.

[9]See Block, et. al. (2008), Brooks (1990), Ramanujan, et. al. (2010), and Thrun (1995)

survival. Now, the two methods are often combined. One common method is to use simulation to update a value function.[10] This type of process is an obvious way of adding learning to limited foresight models, as we will discuss later. However, these methods do not provide intuition as to how a first-time player might approach a situation, nor do they explicitly attempt to explain how real economic agents would play given limited experience. Since these issues are of primary importance to LFMs, AI literature can only provide perspective as opposed to prescription when modeling intermediate valuation.

Regardless of the manner in which intermediate valuation is assigned, the visual game can now be completed, and the model is now almost ready to make predictions. To do so, the model must specify how this visual game is played.

1.3.2 Playing the Visual Game

A natural way to define play of the visual game would be to use some form of backward induction, and indeed this is how we have and will define play in our example models. However, doing so means assigning to the players a certain ignorance beyond limited foresight, a sort of perpetual misconception. backward induction would assume that players believe that actions at future histories will be taken with respect to the current visual game. This ignores the obvious fact that the visual game changes at each history. For this reason, it may instead be wise to define future play in some probabilistic fashion. One possibility would be to assume some sort of probability matching. Another would be to use backward induction but add noise proportional to the distance of the relevant history. Again the possibilities are countless, but for simplicity's sake, we have chosen to define play according to backward induction.

Once future play is determined, there still remains the question of how the initial

[10]Methods of simulation and forms of value functions often separate one specific method from another.

move of the visual game will be made. While the problem of changing visual games is not relevant, and thus there is no intrinsic uncertainty about following through on a predicted action, it is not clear that a deterministic prediction should be made. On the one hand, a deterministic prediction would be nice as it would generate a deterministic series of actions throughout the game, which could be easily compared to backward induction and other deterministic solution concepts. On the other hand, such a deterministic prediction would most likely fail to predict experimental results. Again, in this paper we stick to backward induction for the sake of simplicity and comparability to other solution concepts.

The choice of backward induction also implicitly imposes a psychological assumption, that all players believe that other players have roughly similar foresight ranges to their own. The implication is due to the fact that for backward induction to make psychological sense, a player should assume that each subgame of his visual game is the best representation he can presently make of the visual game to be faced by the player acting at the root of that subgame. For example, consider an acting player (A) with a wealth of foresight, who believes opposing players have little foresight. He might want to take advantage of other players in ways that are not consistent with backward induction is his visual game. Specifically, player A might want to construct the visual games he believes other players will see, and use those games to determine other players' choices. If those visual games are not subgames of player A's visual game, choices determined from those visual games may not be consistent with backward induction in player A's visual game.

1.3.3 An HBLFM with Equiprobable Valuation

Using the same value for all visually-terminal intermediate histories seems a little simplistic, but not enough work has been done looking at how real people evaluate intermediate histories to point to single specific algorithm. That said, if you believe that people have limited foresight, you almost have to believe that people evaluate immediate histories in some systematic way, or else all early play of games such as chess and go would be generally random. We will therefore take a stab at one possible means of evaluation. The model is intended to rank histories that "often" lead to better outcomes higher than histories that "often" lead to worse outcomes. In particular, it ranks VTI histories by assuming equiprobable mixing at all histories outside visual range. In this context, it might be better to think of visual range as the range to which strategic contemplation can take place, and not necessarily the range at which histories can be identified. The model, which we will call HBLFM#2, can be defined in a manner similar to the previous model with a couple of minor changes. First, we don't need β_{ji}, and second, we define the payoff function in the visual game as follows:

-if $h^{VT} \in H^T$ of the original game G, then $u_i^V | h(h^{VT}) = u_i(h^{VT})$

-otherwise, let $u_i^V | h(h^{VT})$ be the payoff to player i given equiprobable mixing among all possible actions by all players in all histories that follow h^{VT}.

The model is otherwise identical to the simple model presented earlier. To see the similarities and differences, let's look at the example game from section 2 and apply the new model (HBLFM#2) again with $\omega_i = 2$ for both players (Figure 1.6).

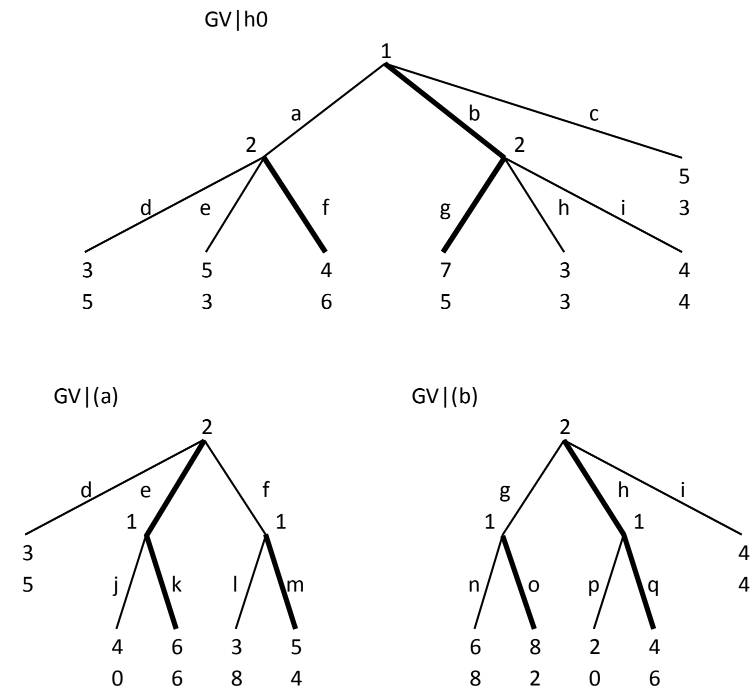

FIGURE 1.6. HBLFM#2

Notice that now each visually terminal intermediate history in $G^V|h_0$ has its value determined according to $u_i^V|h$. Also notice that the prediction of (b, h, q) is not consistent with backward induction.

1.3.4 An Application: The Centipede Game

In their 2009 paper, Field Centipedes, Palacios-Huerta and Volij look at centipede games played by various combinations of students and chess players. The centipede game shown in Figure 1.7 is a good candidate to analyze with HBLFM#2 because it is a game of perfect information with large payoff variation. Because the range of payoffs changes drastically as a function of decisions early in the game, HBLFM#1 is unlikely to give good predictions. HBLFM#2 should do better as its visually terminal intermediate valuations are sensitive to the range of payoffs possible from various intermediate histories.

Palacios-Huerta and Volij find that chess players, who are likely to be skilled at backward induction, tend to end centipede games earlier than students, who are likely to have average backward induction skill. The authors note that common knowledge of rationality may explain the differences. However, results are also consistent with HBLFM#2, assuming that chess players on average have a longer foresight horizon. To show this we look at a 4 step centipede game[11] and observe the outcomes with different values of ω_i, using HBLFM#2.

If $\omega_1 = \omega_2 \geq 4$, then the players perfectly backward induct and the outcome is as shown in Figure 1.7. Note that play begins at the top left with player 1. Moving to the right represents the action "continue" and moving down represents the "end" action.

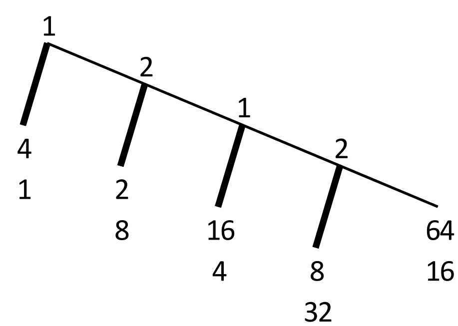

Figure 1.7. A Centipede Game

If $\omega_1 = \omega_2 = 3$, then player one will continue the game at h_0 (the root history) and player two will end the game at $h = (continue)$ as shown in Figure 1.8. Note that the top line of figures represent $G^V|h$ for various h's, while at the bottom is the path of G. We see that at h_0, player one views $h = (continue, continue, continue)$ as a visually-terminal history and assigns it a value of $(36, 24)$.[12]

[11] Palacios-Huerta and Volij looked at a 6 step version but the intuition is the same.
[12] $(64 + 8)/2 = 36$. $(16 + 32)/2 = 24$.

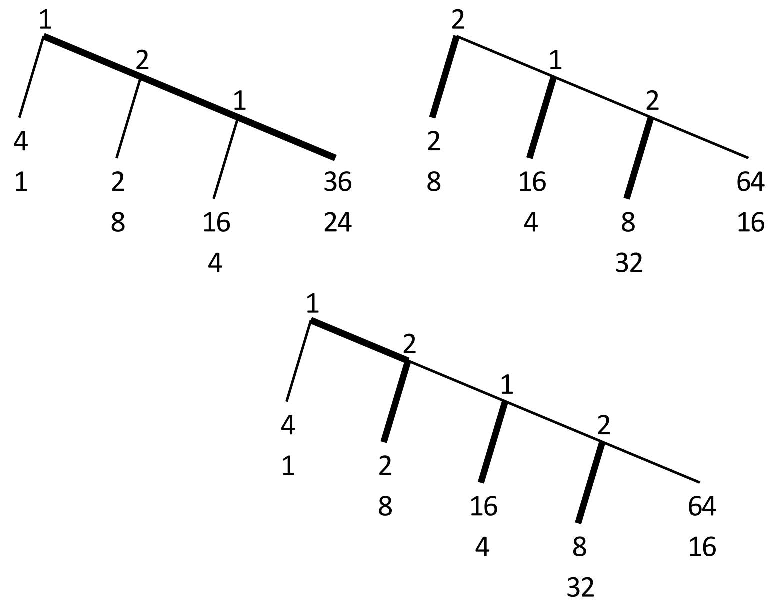

FIGURE 1.8. Centipede with Foresight of 3

If $\omega_1 = \omega_2 = 2$, then the game is ended at $h = (continue, continue)$ as shown in Figure 1.9.

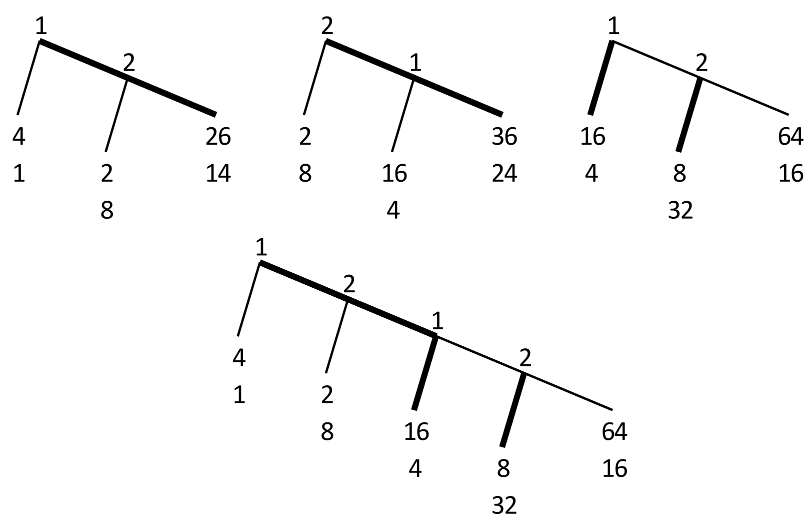

FIGURE 1.9. Centipede with Foresight of 2

If $\omega_1 = \omega_2 = 1$, then the game is ended at $h = (continue, continue, continue)$ as shown in Figure 1.10.

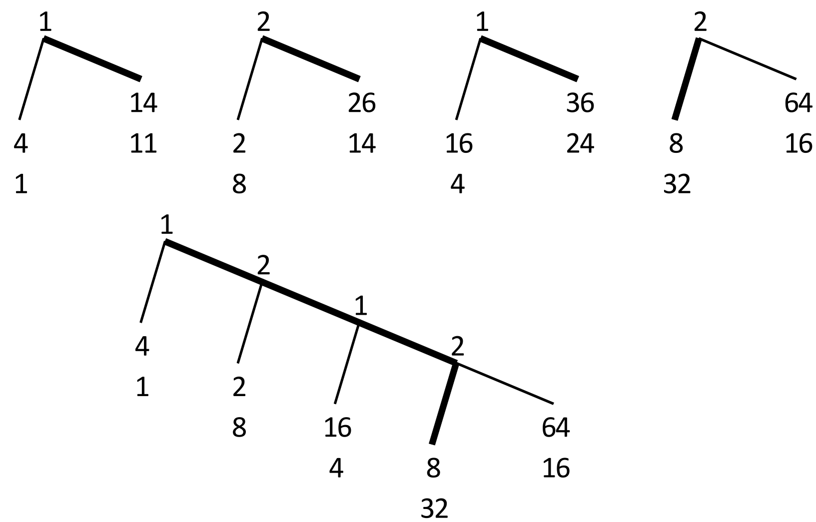

FIGURE 1.10. Centipede with Foresight of 1

Our figures show cases where $\omega_1 = \omega_2$, but it is easy to see the results if ω_i's are different. The punch line in the centipede game is that according to this model the first person to "see" the end of the game will end the game immediately. Obviously, if chess players have longer foresight horizons on average, then the model predicts they will end the game sooner on average.

Notice that as foresight decreases, the game is continued longer. In this example better foresight is actually reducing the players' payoffs on average. This is perhaps counter intuitive as foresight is generally regarded in society as a beneficial trait. What's more, this can be true even in zero sum games. Better foresight does not always lead to better outcomes in our model.

Before moving on, we must call attention to a forthcoming paper from Levitt, List, and Sadoff that attempts to replicate results from the Palacios-Huerta and Volij study and fails to do so. However, Levitt et. al. do not test a mixed subject pool of chess players and students, they only test chess players. More importantly, they do not inform their subjects that they are playing against other chess players. This brings up an important feature of our models which is the implicit assumption by players that other players have foresight ranges close to their own. Without this assumption, backward induction is a poor modeling choice for how subjects play the game.

1.3.5 An Application: Bargaining

There are many economic models that predict that when bargaining can produce profitable results, and transaction costs are low, agreements can and should be reached quickly and painlessly. However, observation of real world bargaining scenarios shows that often times negotiation can be long and messy, often ending in payoffs to both sides that are lower than would have been possible if an agreement could have been made sooner. Strikes for example often end up costing both sides. This seems to be a violation at least of the spirit of backward induction, but using our second horizon based limited foresight model, we can see at least one possible explanation for this phenomenon.

Consider the following stylized bargaining game. Two players must split \$512. The game begins with player 1 choosing either a 50/50 split of the money, or an 80/20 split in his favor. Player 2 must then choose to either accept the split proposed by player 1, or to reject the split and counter-offer either a 50/50 or an 80/20 split in player 2's favor. However, if the split is rejected by player 2, the total amount to split is reduced by half. This procedure continues back and forth until a split is accepted, or until there is one dollar left. When there is one dollar left, the proposed split of that dollar can either be accepted, or rejected in which case both players receive nothing. One possible end game is illustrated in Figure 1.11.

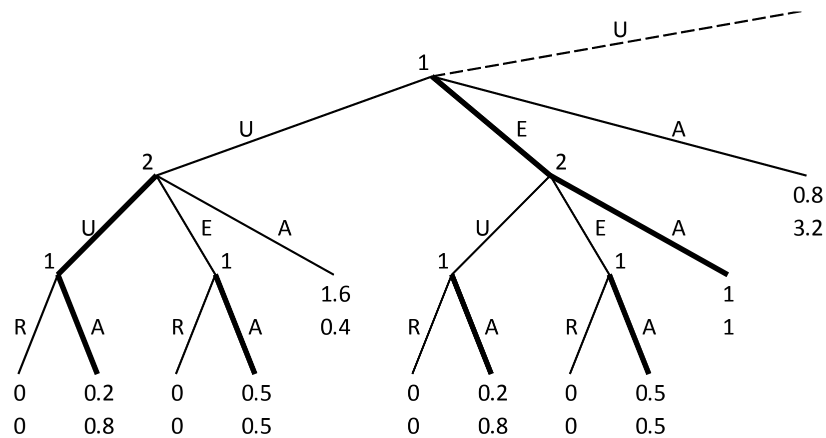

Figure 1.11. Bargaining Game (End Game)

Some may note that the rate of decay and uneven split percentages are rather extreme, but the intuition here extends to similar games with smaller rates of decay and less lopsided splits. These similar games are simply not as clean numerically. The specifics of the game were also chosen so as to highlight a difference in predicted play when the game is analyzed using backward induction and when analyzed using our HBLFM with equiprobable valuation. The relationship between the decay must be such that it is always profitable for either player to reject an uneven split and counter with an equal split. Also it must be profitable to accept an equal split rather than reject it and make some counter offer. When these two conditions are true then the unique backwards induction solution of the game is for player 1 to offer the 50/50 split and for player 2 to accept it.

Alternatively, when analyzing the game using HBLFM#2, with a foresight horizon of 2, the uncertainty of decisions at future histories makes offering an uneven split more attractive. This can destroy the backward induction solution as players successively counter with unfair offers at the beginning of the game. This result is consistent with behavior that is often witnessed in the real world bargaining contexts we mentioned earlier.

Figure 1.12 presents the difference in the first move of the game when analyzing it with backward induction and HBLFM#2. We see here that player 1 makes a lopsided offer, assuming that player 2 will accept it. However, as players will always assume that a fair offer will be accepted, player 2 will reject the lopsided offer.[13] The bargaining will continue, with neither player able to make more than if player 1 had simply offered a fair split to start, and player 2 had accepted.

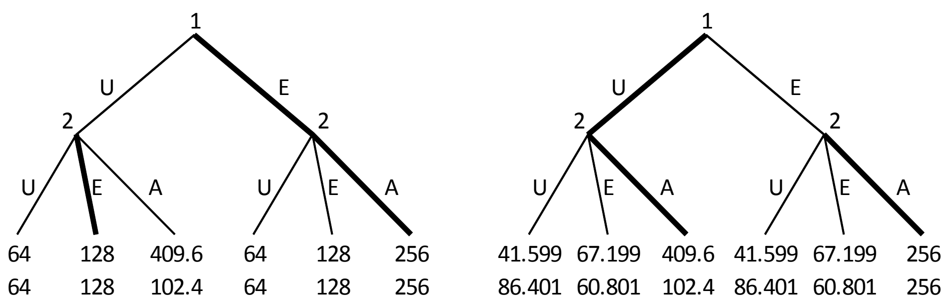

FIGURE 1.12. Backward Induction (left) vs. HBLFM (right)

With the current rate of decay, additional foresight will cause the predicted path of HBLFM#2 to match up with the path predicted by backward induction. This is because unless you can counter with an uneven offer and have it be accepted, the amount of money decayed is prohibitive. However, if the rate of decay were lessened and the uneven split option were more equal, additional foresight horizons could still result in multiple uneven offers.

1.4 Directed Limited Foresight Models

Directed limited foresight models are relatively more complicated than their horizon-based counterparts, but are arguably more psychologically accurate. Multiple possible futures are not seen in the same way one looks out over the countryside. A possible

[13] Player 2 can make at least 128 by countering with a fair offer vs. 102.4 if he accepts the lopsided offer.

future must generally be envisioned independently of alternatives. Therefore, possible futures would be considered one at a time. For this reason, it is likely that foresight in strategic situations would be directed, that is to say that players would choose to fully analyze certain possible paths of play, while essentially ignoring others. DLFMs take this intuition into account by allowing an element of freedom in the way visual ranges are determined.

The DLFM in this paper will assume that players have a limited endowment of foresight, at each history, with which to allocate over possible future histories. The foresight allocation will be used to determine a visual range. While this type of model leaves us with the same open-ended modeling questions as HBLFMs (evaluating visually-terminal intermediate histories, and play of the resulting game), it also adds an additional question: How do we allocate the foresight endowment? Again, as with the other modeling issues, there is not a great deal of research into how such a question should be answered. Still, as the general intuition makes sense, it is worth taking a stab at.

1.4.1 Directing Foresight

Directing (allocating) foresight is an interesting and complicated problem in itself. One's first instinct may be to assume that foresight is random, and in reality, this may be entirely reasonable when thinking of unsophisticated players playing a game none of them have ever seen before. This would be a very similar assumption as assuming that all visually-terminal intermediate histories are evaluated as being equal. In fact, if the latter assumption is made, even a more sophisticated foresight allocation rule could likely end up looking rather random. Intuition is consistent as if a player is assumed to have no particular clue about the value of being at any intermediate history, that player also probably has little clue about how to allocate foresight over those same histories. However, if we assume players have some intuition about the value of these VTI histories, it would be intuitive to direct foresight as a function of these values. In particular, an acting player would be more likely to allocate foresight to histories that he values highly given his current perspective. Still there are obviously alternative methods by which foresight can be allocated in an intuitive way.

To illustrate a directed limited foresight model, we will adapt the horizon-based LFM with equiprobable valuation. A player will begin his decision making process with no foresight allocated and so will have a very limited view of the game. Given this limited view, he will allocate his foresight to the history that "seems" to lead to the highest payoff. We will make this notion more clear in the coming section. With this new unit of foresight allocated, his view of the game will change. A player will then allocate the next unit of foresight given his new view of the game. Foresight will be allocated in this fashion, on the margin, until all units have been allocated or the entire remaining game is revealed.

1.4.2 A Specific Directed Limited Foresight Model

To begin, let $G = (I, H, L, ((u_i); i \in I))$ be a finite extensive form game of perfect information, where I is the set of players, H is the set of histories, $L : H^P \to I$ is the mapping of who moves at each partial history where $H^P \subset H$ is the set of intermediate or partial histories, and $u_i : H^T \to \mathbb{R}$ is player i's payoff function where $H^T \subset H$ is the subset of terminal histories. Let $\omega_i \in \mathbb{N}$ be player i's foresight endowment. Also, as it will become useful later, let $v_i(h) : H^P \to \mathbb{R}$ be the payoff to player i given equiprobable mixing among all possible actions by all players in all histories that follow h.

As was the case in the HBLFMs, our first task will be to define the visual range. Because a player's foresight endowment can potentially be allocated in different ways, there is actually a set of visual ranges at each history. We will define the set of foresight allocations with respect to a particular history and foresight endowment. A visual range will then be defined as the foresight allocation as well as all immediate followers of any history in the foresight allocation. This implies that only intermediate histories need to be investigated, and that once an intermediate history is reached, all of its immediate followers become visible. We'll start with zero units of foresight and then build from there.

We define $\eta^0|h$ as the set of possible foresight allocations at h (where $L(h) = j$) when $\omega_j = 0$. An element of $\eta^0|h$ is a set of histories that we will denote as $H^{\eta^0}|h$. Note that an element of $\eta^0|h$ is not simply any single history that is in any possible foresight allocation. This means that $\eta^0|h$ is a set of sets. However, since no investigation takes place when $\omega_j = 0$, $\eta^0|h$ has only one element (h, which should be thought of as "free").

Allocating zero units of foresight is of course rather simple, but in order to allocate additional units, the model assumes players allocate foresight to histories one unit at a time. To see how this is done,we will now define $\eta^1|h$ as the set of possible foresight allocations when $\omega_i = 1$. To do this we start at $\eta^0|h$.

For each $H^{\eta^0}|h \in \eta^0|h$ we can define a game, $G^{\eta^0}|h = (I, H^{V\eta^0}|h, L, (v_i|h; i \in I))$. Where $H^{V\eta^0}|h$ is the set of histories in $H^{\eta^0}|h$ as well as all immediate followers of those histories, less all histories that are in H^T, and less all histories that don't have at least one follower that (1) is not in $H^{\eta^0}|h$ and (2) is not in H^T. Then $v_i|h : H^{V\eta^0 T}|h \rightarrow \mathbb{R}$ is player i's payoff function, as defined earlier, where $H^{V\eta^0 T}|h \subset H^{V\eta^0}|h$ is the subset of visually-terminal histories. Terminal histories of $G^{\eta^0}|h$ can be thought of as possible intermediate histories to investigate.

Let us now consider all paths of the game $G^{\eta^0}|h$ that are consistent with backward induction. For one of these paths, consider the last history reached. This history may be investigated.[14] Add it to $H^{\eta^0}|h$ to create some $H^{\eta^1}|h$. Do this for each visually terminal history of each path to create different $H^{\eta^1}|h$'s. Each of these $H^{\eta^1}|h$'s is an

[14]The history will be investigated if there is only one path.

element of $\eta^1|h$.

We can define $\eta^1|h$ as the set of all sets of histories created in such a way from all elements of $\eta^0|h$. Similarly, we can define $\eta^m|h$ as the set of all sets of histories created from all elements of $\eta^{m-1}|h$. In this way, we can create the set of possible foresight allocations for any value of ω_j. Then, for any of these possible foresight allocations, where $L(h) = j$ and $\omega_j = m$, a visual game $G^V|h = (I, H^V|h, L, (u_i^V|h; i \in I))$ can be created where $H^V|h$ is the set of histories in some $H^{\eta^m}|h$ as well as all immediate followers of all histories in $H^{\eta^m}|h$. $H^{VT}|h$ is again the set of visually terminal histories. I and L are as defined before, and $u_i^V|h$ is defined as in our second model:

-if $h^{VT} \in H^T$ of the original game G, then $u_i^V|h(h^{VT}) = u_i(h^{VT})$

-otherwise, let $u_i^V|h(h^{VT})$ be the payoff to player i given equiprobable mixing among all possible actions by all players in all histories that follow h^{VT}.

Then, any history $h' \in H^V|h$ reached on any backward induction path of any $G^V|h$ where $D|h(h') = 1$ is then said to be consistent with this model. A prediction of our model is then any consistent h^T for which all preceding histories are consistent.

Let's look once more at our example game from Figure 1. Figure 1.13 shows the progression of $G^\eta|h_0$ as ω_1 increases from 0 to 5.

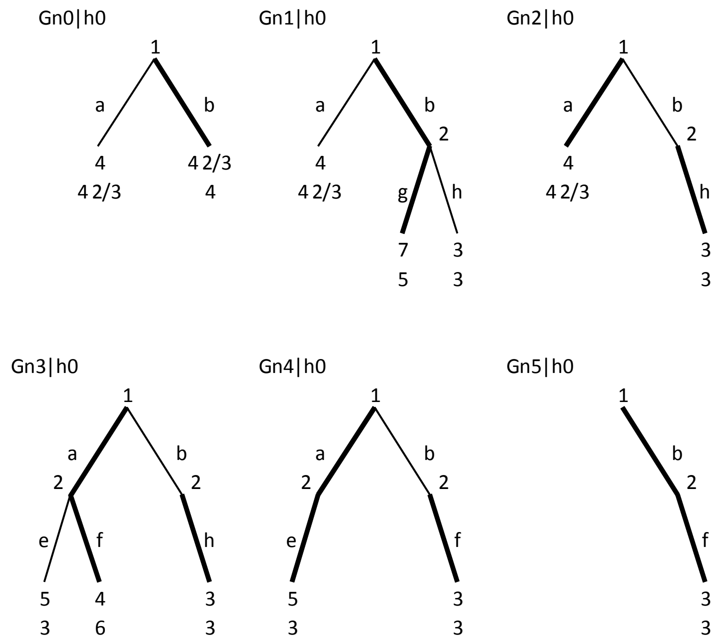

FIGURE 1.13. DLFM Foresight Allocation Procedure

Notice that with six units of foresight the entire game would be investigated by player 1 at h_0.

Figure 1.14 shows the visual game generated from $H^{\eta^2}|h_0$.

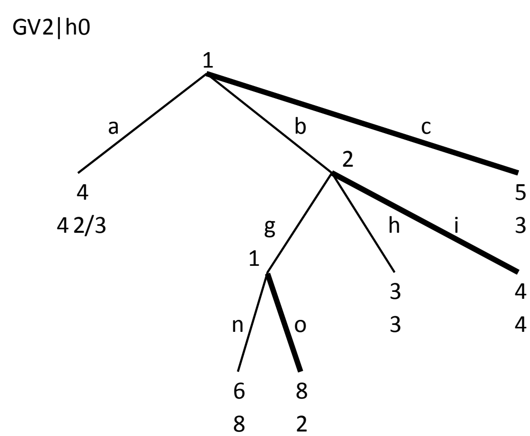

FIGURE 1.14. DLFM Visual Game

The prediction here would be simply (c).

Figure 1.15 shows the visual game generated from $H^{\eta^5}|h_0$.

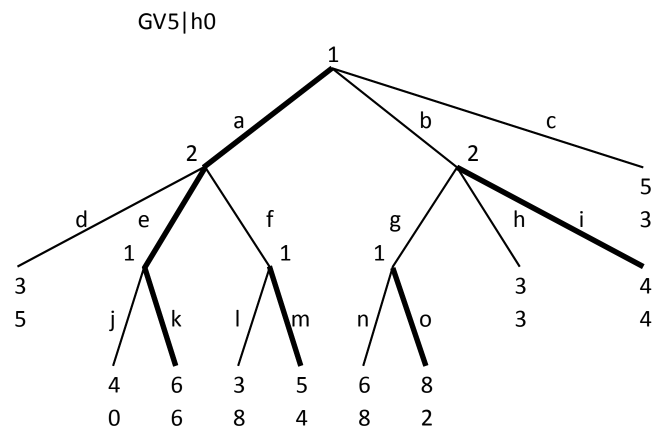

FIGURE 1.15. Another DLFM Visual Game

In this instance, assuming player 2 would see the entire subgame at (a), the prediction is (a, e, k).

1.5 Results

Because there are so many possible ways of constructing a limited foresight model, there are few results that would apply to the class as a whole. However, by making use of backward induction in the specific models presented here, there are a couple of results that apply to those models. Also, we will show that when restricting attention to general games of perfect information, we can say more about the relationship between backward induction, and our models.

1.5.1 The Existence of Limited Foresight Models

Theorem 1. *For each model presented in this paper, at least one predicted path exists for all finite games of perfect information.*

Proof. For all of our models,since G is a finite game of perfect information, G^V is a finite game of perfect information. All finite games of perfect information have at least one backward induction solution. Therefore, for each history in H^P, at least one following history is consistent. If each partial history has at least one consistent follower, then at least one h^T must exist where all histories preceding h^T are consistent. So, a prediction (h^T) exists. □

Note that generally predictions are not necessarily unique, even for specific realizations of exogenous variables such as ω_i or β_{ij}.

1.5.2 Unlimited Foresight

Theorem 2. *For any finite game of perfect information G, and where K is the distance from the root history of the farthest terminal history in H^T, all backward induction paths of G are predicted by our horizon based models when $\omega_i \geq K$ $\forall i$.*

Theorem 3. *For any finite game of perfect information G, and where M is the number of histories in H^P, all backward induction paths of G are predicted by our directed model when $\omega_i \geq M$ $\forall i$.*

Proof. If $\omega_i \geq K(M)$ $\forall i$, then all visual games G^V are subgames of G. Any backward induction solution to G, and in particular any h^T reached on a backward induction path, projects to any subgame of G (i.e. any G^V). Therefore, any history reached along a backward induction path of G must be consistent. Since all histories preceding h^T, and h^T itself, are consistent, h^T is a prediction of our models. □

This notably does not rule out the possibility that additional paths of play that are not along a backward induction path may be predicted, even when foresight is effectively unlimited. For example, consider the game in Figure 1.16. There are two backward induction paths of this game ((out) and (in, In, b)). However, with effectively unlimited foresight, all four of the terminal nodes are predicted by our models. This is because each decision is made independently. Player 1 might "believe" that player 2 will choose (in, In, b) over (in, In, a) but that does not mean that player 2 actually will. Furthermore, player 1, need not have consistent "beliefs" about player 2 between h_0 and (in). At h_0 he might think player 2 will chose b, but at (in) he may change his mind and think player 2 will choose a. This type of behavior may be appropriate in some situations, like if there was a long wait between decisions at h_0 and (in), but not in others. For modelers who wish to rule out this sort of behavior, the models presented can be modified so as to enforce consistent "beliefs" whenever player i's visual games is a subgame of player i's previous visual game.

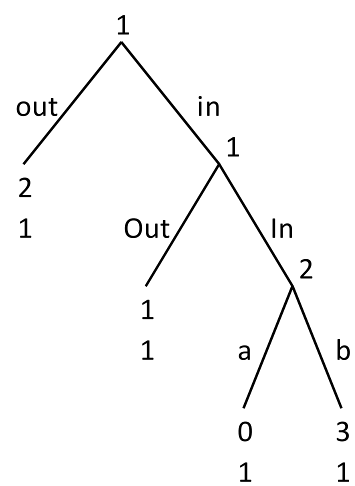

FIGURE 1.16. An Illustrative Example

In generic games of perfect information, the indifference in the previous example is eliminated and we don't get the same ambiguity.

Corollary 4. *If G is a* ***generic*** *finite game of perfect information, with unlimited foresight, our models all uniquely predict the backward induction solution.*

Proof. Proceeding from our previous proof, as there is a unique backward induction profile to any generic game, then for any visual game G^V there is only one consistent history. There is therefore a unique prediction of our models, and, by the preceding theorems, it must be the backward induction path. □

1.5.3 More on Generic Games of Perfect Information

Although our models do not always predict the backward induction path when foresight is limited, we can say that they generally do in the following sense:

Proposition 5. *For any finite game tree of perfect information, where payoffs are assigned to terminal histories according to a continuous single-valued CDF. Then, before these payoffs are realized, the probability that our horizon-based models predict the backward induction solution is greater than* $1/T$*, where* T *is the number of histories in* H^T*.*

Proof. See Appendix. □

1.6 Further Thoughts.

We now look at how LFMs can be extended and built upon in future literature (the future in this case meaning only a step or two ahead). We'll discuss extending the models to a wider set of games, as well as how elements of learning can fit into the model in very interesting and intuitive ways.

1.6.1 Extending the Models

The models presented in the previous sections are defined for finite games of perfect information with discrete action sets, but the intuition that these models are based on applies to all games. It is therefore reasonable to assume that the limited foresight models could be generalized to infinite games and games of imperfect information.

The main issue when dealing with infinite games relates to evaluating VTI histories. Some of the models above involve calculations that would not be well defined with an infinite game. That said, our first LFM would be defined for infinite games, and there are many other LFMs that would be as well.

Continuous action sets are not a big problem when defining HBLFMs, though they can arguably be a problem intuitively. DLFMs are more difficult, as when foresight is defined as an integer endowment, it is not clear how one would deal with continuous action sets. It is certainly possible to think of choosing a random point within the continuous range, and this may even be close to what most people do when confronted with such a decision. The main problem is when an action set is both continuous and discrete. For example, if an action set consisted of the number 10, and all real numbers between 1 and 5, then there is a question about how to allocate weight on the number 10. This would not be all that difficult to define in a model, but the intuition is less clear.

Various forms of imperfect information can be a problem in a number of ways, especially when it comes to defining visual ranges. For example, if there are multiple paths to the same information set, would a player recognize this when investigating only one of these paths? Again, this question is far from impossible to answer, but it adds complexity to the model. Also, with simultaneous move games, or generally games with information sets containing multiple histories, visual games might look very different from the original game, as depicted, in that unresolved uncertainty always "looks" as though it is in the future, even if the uncertainty is about an action earlier in the game tree. Theoretically this should still be possible to deal with, but the cost of added complexity outweighs benefits of generality for the current paper.

The bottom line with all these extensions is that they are possible, and so LFMs as a class should be thought of as capable of describing play in any game, but as this paper is meant as an introduction, I leave specific definition of more general models to future papers.

1.6.2 Learning

The intuition of LFMs leads naturally to the concept of learning. In particular, we can think of the process of evaluating visually-terminal intermediate histories as one that involves significant learning. After all, the rules of the game do not give instruction as to how such histories should be valued, and so determining these values must generally be learned.

To illustrate, think of the game of chess. Great chess playing is not merely a function of having a wealth of foresight. It also involves understanding how positioning on the board affects your ability to win, understanding the value of different pieces, etc. Yet you do not get points for having certain pieces when the game ends, nor for being in certain positions (other than check-mate). Therefore these concepts are not defined by the actual payoffs of the game, and instead seem to be directly related to valuing VTI histories.

This type of learning is not touched on in the current economic learning literature, and yet it is incredibly prevalent in society. Understanding the value of being in some intermediate position is what allows economic agents with limited foresight to make long-term decisions. An unknown and complex future creates a fundamental uncertainty that we all face. Dealing with it requires significant learning.

Obviously this is an area where knowledge of learning by artificial intelligences can relate closely. Many chess-playing programs use outcomes of simulated games to generate value functions over possible histories. It is unlikely that even grandmaster chess players use these exact algorithms, but it is plausible that they do something similar. And while it is likely that a chess playing computer with enough processing power and simulation time would be able to beat the grandmaster chess player at chess, the grandmaster chess player would likely be able to apply his experience with chess to similar strategic games in a way that the computer program could not. Understanding this difference could help shed light on the nature of human vs.

artificial intelligence.

It would seem natural then to study how real people evaluate intermediate positions, whether these evaluations change with experience, etc. Again, these are all topics for future papers, but we hope that starting from a more psychologically intuitive model will result in answers that can be applied broadly, as answers gleaned from models that are not psychologically intuitive tend only to be applicable to those models and/or those situations being modeled themselves.

1.7 Conclusions

We've presented a number of models based on the idea that people have limited ability to backward induct, or foresee future events. These models present a psychologically intuitive framework for describing how individuals might actually play complex strategic games. Generally, this intuitivity comes at the cost of complexity. However, if models of decision making are to be applied broadly, it is reasonable to assume that they must be based on realistic assumptions about the cognitive abilities of actual decision makers.

Although the class of limited foresight models is broad and so specific predictions of play are not made by the class as a whole, experimentation can serve to tease out which specific LFMs have good predictive power for various games. Indeed there is room for much more work on LFMs, in both the directions alluded to in the final sections of the paper, and undoubtedly in directions we cannot currently foresee.

Chapter 2
A Mechanism to Elicit Continuation Values in Experiments

2.1 Introduction

There are many situations where experimenters might like to know experimental subjects' continuation value in a game. Continuation values, the value a subject would accept in order to forgo continuing the game and receiving his payoff in the game, can be used to more directly test beliefs about payoffs than can actions. Therefore, in situations where theory makes predictions about beliefs and payoffs, knowledge of continuation values would be valuable.

For the most part, in such cases, experimenters have tried to infer these values from recorded actions. However, even in well designed experiments these inferences are often made with some amount of uncertainty. In order to resolve this uncertainty, we propose a mechanism for games of discretized time and bounded payoffs. The mechanism builds off those proposed by Becker, DeGroot, and Marschak ('64), Allen ('87), Grenther ('92), Karni ('09), and Mobius, et al. ('10) in that it elicits truthful beliefs. We show that such a mechanism can be implemented repeatedly throughout the course of a game in such a way that play of the underlying game is unchanged. As motivation for our mechanism, the following examples readily come to mind.

Much mixed strategy experimentation uses various distributional implications of independent mixing to estimate if players were playing according to a mixed Nash equilibrium. However, knowing subjects' continuation values would allow a more direct check of theory as one could compare the continuation value to the expected payoff of a subject's chosen action in equilibrium. In other words, we could check if subjects were trying to exploit each other or if their continuation value was for example the minimax value.

Various models of bounded rationality implicitly make predictions about subjects' expected payoffs in certain games. Comparing continuation values could help test the validity of these theories and could be used to compare these models to models of full rationality. Talk of bounded rationality naturally leads us to models of learning as well. Continuation values can give a clearer picture of learning than actions alone.

As a specific example, in Dufwenberg, et. al. ('10) the authors categorized players as having had a certain epiphany if they did not play any dominated actions for the remainder of the experiment. However, it is possible that some number of non-dominated actions were simply chosen by luck. Comparing subjects' continuation values to the dominant strategy payoff, on the other hand, could give an independent and perhaps more precise measure of the timing of this epiphany.

These examples are by no means exhaustive. Indeed there are a great number of potential applications for a mechanism designed to elicit continuation values in experiments.

In the coming section we will introduce a simple game for which we might want to know continuation values, and show why a number of simple mechanisms are insufficient for eliciting continuation values. In section 2.3 we discuss our mechanism and show that under certain conditions the mechanism can elicit truthful continuation values without changing the underlying play of the game.

2.2 What Not to Do

This section outlines some of the intuition of the mechanism using informal language and an example. For those who would prefer a formal treatment of our mechanism to start, feel free to skip to the next section.

The first challenge in designing a mechanism to elicit continuation values is to incentivize subjects to report accurately. This is accomplished using a Becker, De-

Groot, Marschak (BDM) type procedure. Subjects are asked to give a valuation conditional on being at a certain point in a game. For a risk neutral agent this value can be thought of as how much they expect to make in the game given their current position. Sometime after reporting the value, a random number is drawn.[1] If that number is greater than the valuation given, the subject will be paid the random number. Otherwise, the subject will be paid the winnings of the game. This type of procedure has nice properties. First, reporting what he expects to make in the game maximizes a subject's payoff in expectation. Second, the game can continue even when the subject is paid the random number.

Although at first it may seem as though one could simply repeat the above procedure at various nodes without much thought, the above procedure actually changes the payoffs of the underlying game. Keeping incentives aligned creates a second challenge. If we do not align incentives properly, then we can not apply our results to the original game or, by extension, any real world scenarios. To see how things can break down, consider the game represented by the following extensive form:

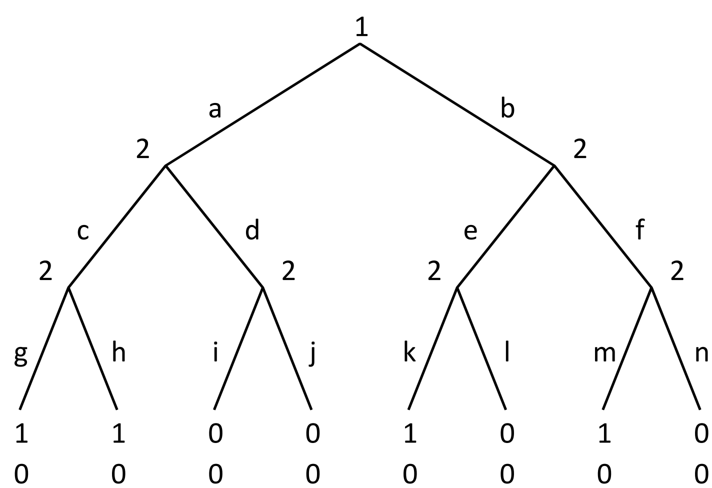

FIGURE 2.1. The Following Extensive Form

[1]The random number should be drawn from a support that encompasses the maximum and minimum possible earnings for that subject in the game (at least conditional on the current position).

For the purposes of this thought exercise lets assume we are only interested in player 1's continuation value at every penultimate node ((a, c),(a, d),(b, e),(b, f)). It would be nice if we could simply ask for a continuation value when we reach one of the four nodes of interest. We cannot. Imagine that in the original game, player 1 believes player 2 will mix equiprobably at every opportunity. As a result player 1 is indifferent between a and b. Now we implement the BDM-style procedure. Player 1 would no longer be indifferent. If he chooses a, then when eventually reporting a continuation value (at (a, c) or (a, d)), he will know for certain what his "payoff" will be in the underlying game. On the other hand, if he chooses b his "payoff" in the underlying game will be uncertain when reporting a value. Assume we draw our random number from a uniform distribution between 0 and 1, player 1's expected payoff (in the modified game) for choosing a would be .75.[2] Choosing b would only net him .625.[3] By adding the BDM-style procedure, we have changed the incentives, and thus play, of the underlying game.

There are many examples of similar games and similar procedures that will alter incentives. Generally payoffs with the procedure depend on play of the game in a fashion that is difficult to control. We get around this problem in the following way. First, we randomly select a single valuation and pay according to a BDM-style procedure. Second, the valuation selected is not a function of the path of play.

One may immediately wonder how to obtain a valuation if it is not on the path of play. This is achieved by engaging in a second play through of the game from the node of the selected valuation if that node was not on the original path of play. A successful process requires an assumption that beliefs about play from that node do not change in the second play through when compared with beliefs had that node

[2] Half the time he will be at (a, c), report a continuation value of 1 and get 1 for sure. Half the time he will be at (a, d) he will report 0 and thus always receive the random number, which will be .5 on average. $(.5 * 1 + .5 * .5 = .75)$

[3] At both (b, e) and (b, f) he will report .5. Then, half the time the random number will be above .5 and he'll get .75 on average. The other half the time he'll get his game payoff which will be .5 on average. $(.5 * .75 + .5 * .5 = .625)$

been reached in the first. In some games this requirement is innocuous, in others it is very strict. After presenting the mechanism in greater detail, we will discuss this requirement with respect to various classes of games.

2.3 The Mechanism

For simplicity we present the mechanism here with perfect information. We discuss imperfect information in a later section. We also assume agents are risk neutral. Discussion of robust risk preferences we also save for later in the paper. We will also be making a strong assumption about agents' beliefs that we elaborate on later.

Consider some game G, where $G = (N, H, L, A, (u_i)_{i \in N})$, N is the set of players, H is the set of histories, $h \in H$ is a set of actions, $L : H \to N$ represents who moves at each history. $A(h)$ is the set of actions available to $L(h)$ at h, and $u_i : H^T \to \mathbb{R}$ is player i's payoff function where H^T is the set of terminal histories. $L(h)$, $A(h)$, and $u_i(h)$ are of course defined similarly but at the individual history level. Further, let $\underline{u_i}$ be the minimum possible payoff to i, and $\overline{u_i}$ be the highest. We also say h' precedes h if $h' \subset h$.

Now, suppose we would like to know the value that players place on continuing the game as the game progresses. We can ask each player at each history on the path of play, as it is reached, to choose a number x_{hi} in the set $[\underline{u_i}, \overline{u_i}]$ (x_{hi} should be thought of as player i's continuation value at h). At the end of the game, a single history is chosen according to a lottery over all histories in H, regardless of the path of play, and according to probabilities $\beta_h \in [0, 1]$.[4] Players are informed of these probabilities before play begins. If the chosen h was not reached while playing the game, then each player chooses another x_{hi}. The game is then played through to completion starting at the chosen history. If the chosen h has been reached in the first play through, then no additional play is required. We will then draw a number r_{hi} (one for each player)

[4] $\sum_H \beta_h = 1$

uniformly from the set $[\underline{u_i}, \overline{u_i}]$. Each player's payoff is then determined according to a function $u_i(h^T) + \delta_{hi} * u_i(h^{T'}) + \delta'_{hi} * w_{hi}(x_{hi}, h^{T'})$ with the following properties: $\delta_{hi}, \delta'_{hi} \in \mathbb{R}_+$, h^T is the terminal history reached in the original play-through of the game G, $h^{T'}$ is the terminal history reached in the second play-through of G (if the chosen history was reached in the original play-through, $h^{T'} = h^T$) and $w_{hi}(x_{hi}, h^{T'})$ is defined as:

$$w_{hi}(x_{hi}, h^{T'}) = \begin{cases} r_{hi} & \text{if } r_{hi} \geq x_{hi} \\ u_i(h^{T'}) & \text{if } r_{hi} < x_{hi} \end{cases}$$

Theorem 6. *Let $v_i(h) \in [\underline{u_i}, \overline{u_i}]$ be player i's continuation value at h (with respect to u_i). Choosing $x_{hi} = v_i(h)$ maximizes player i's payoff in expectation.*

Proof. The value that maximizes player i's payoff in expectation is:

$$\arg\max_{x_{hi}} \frac{\beta_h * \delta'_{hi} * ((x_{hi} - \underline{u_i}) * v_i(h) + (\overline{u_i} - x_{hi})(\overline{u_i} + x_{hi})/2)}{\overline{u_i} - \underline{u_i}}$$

The first order condition of this maximization problem is $\beta_h * \delta'_{hi} * (v_i(h) - x_{hi}) = 0$. So, $x_{hi} = v_i(h)$ solves the first order condition, and since the function is concave everywhere, it is a maximum.[5] □

Now, with the addition of the mechanism we are inducing a new game. The payoffs in this game can be written as $u'_i(h^T, \overrightarrow{h}^{T'}) = u_i(h^T) + W_i(X_i, \overrightarrow{h}^{T'})$, where $W_i(X_i, \overrightarrow{h}^{T'}) = \sum_{j \in H} \beta_j * (\delta_{ji} * u_i(h^{T'}(j)) + \delta'_{ji} * w_{ji}(x_{ji}, h^{T'}(j)))$. Here the payoff represents the expected payoff when reaching h^T, given a vector of outcomes $\overrightarrow{h}^{T'}$, one for each j such that each j is associated with a specific terminal history, $h^{T'}(j)$. Although the game has changed, we note that the elements of G remain in G'.[6] We refer to these elements as the "underlying" elements of G' (in contrast with x_{hi} choices). We also use the term "underlying play" to mean the portions of a strategy

[5] Notice that if β_h or δ^2_{hn} are zero then any value of x_{hn} maximizes the equation. This could correspond to the situation where you don't care about beliefs at h_n and so you don't ask

[6] For consistency we assume that the names and labels of these elements remain that same in G'.

profile in G' that are not x_{hi} choices. We will show that, under certain assumptions about beliefs, the incentives related to the elements of G are unchanged in G' when using the proposed mechanism, and that therefore, a player who uses some strategy σ in G would want to use a projection of σ (onto the common underlying elements in G') when playing G'.

Our assumptions about beliefs are embodied in a distribution Γ_i. The notation for Γ_i is as follows. $\Gamma_i(h)$ is a probability distribution over h^T's representing player i's beliefs about the probability of reaching those h^T's when at h. $\Gamma_i(a(h))$ is a probability distribution over h^T's representing player i's beliefs after taking action $a(h)$ Note that there is a distribution for each h or $a(h)$. $\Gamma_i(h^T|h)$ (and similarly $\Gamma_i(h^T|a(h)))$ is the probability placed on h^T.

Definition 7. *We define a "projection" of play in the following way. Consider the games G and G' as defined above. Now consider a strategy σ' in G'. Notice that some elements of σ' relate to continuation value choices (x_{hi}), and the rest relate to underlying action choices which correspond to actions choices in G. Now consider a strategy σ in G. σ' is a "projection" of σ if every underlying action chosen in σ' is the same as the corresponding action chosen in σ.*[7]

We note here that beliefs about play can be projected in a similar fashion.

Theorem 8. *Let σ_i be a strategy in G for player i that maximizes his expected payoff under some set of beliefs about future play Γ_i. Let σ' be a strategy in G' for player i with the property that all x_{hi} choices are made to maximize expected payoffs. If σ' is a projection of σ, then σ' maximizes player i's expected payoff in G' when beliefs about underlying play in G' are a projection of Γ_i.*

A formal proof is available as an appendix. For now we give an intuitive defense of Theorem 8.

[7] This needs to be true for underlying actions in both play throughs.

As alluded to previously, Theorem 8 relies heavily on an assumption that beliefs about play in both play throughs of G' are the same as those about play in G. This assumption perhaps seems strong, but if play is consistent when beliefs are the same, then assuming beliefs are in fact the same seems more reasonable. This means we essentially show existence but not uniqueness.

To start, think of the second play through of G' (if necessary). Whatever x_{hi} you give, $u_i(h^{T'})$ enters your payoff with weight $\delta_{hi} + \delta'_{hi} * \frac{x_{hi} - \underline{u_i}}{\overline{u_i} - \underline{u_i}}$. The remaining part of your payoff is not a function of play. Clearly then, if beliefs are the same in the second play through of G' and in G, then incentives should be proportional in G and in the second play through of G'. Therefore, so long as indifference resolves in the same direction in both, play is consistent between both.

Now, since the lottery over histories is not a function of the first play through of G', the additional payoff of the mechanism $(W_i(X_i, \overrightarrow{h}^{T'}))$ is only a function of the first play through when the lottery selects a history on the first path of play. In expectation, then, similar as before, $u_i(h^T)$ enters with weight $1 + \sum_{j \subset h^T, j \in H} \beta_j * (\delta_{ji} + \delta'_{ji} * \frac{x_{ji} - \underline{u_i}}{\overline{u_i} - \underline{u_i}})$. Again this suggests if beliefs about play in the first play through of G' and beliefs about play in G are the same, then direct incentives are proportional. However, there is one additional issue of indirect incentives (how h^T might affect play in the second play through). We must also require that players believe (when making decisions in the first play through of G') that beliefs in the second play through of G' will be the same as current beliefs (beliefs in the first play through). This ensures that subjects believe the payoffs related to a second play through are truly independent of the first play through. So direct incentives are proportional and indirect incentives are ruled out. Therefore, so long as indifference is again resolved in the same direction in all cases, play is consistent between all play throughs.

2.4 Strategy Method

The mechanism is easily extended to the strategy method. In fact there is no need to play through the game a second time. Since the strategy method allows us to map an action to each history already, a lottery over those histories will always choose a history with which a terminal history has been associated.

The procedure for the strategy mechanism is, for any history of interest, to ask players their continuation value using the same x_{hi} procedure as before, and to ask for an action at each history where a player acts. After these decisions are made, there is a lottery over histories and the players are paid according to the terminal histories reached according to their choices, and the random term associated with x_{hi}'s and r_{hi}'s, as before.

The proofs work largely as before since we are merely asking subjects their continuation values and action choices before the lottery instead of after. As a result, the payoffs are only scaled by the probabilities associated with the lottery, and preferences over actions remain unchanged at any given history (since the payoff to each action, at a given history, is scaled by the same probabilities).

In this case the mechanism is actually much cleaner, and experimenters can potentially get much more information about off-path play. Therefore, in instances where continuation values are desired and the experimenter believes the strategy method appropriate, this procedure is likely the best option. Of course in many games the strategy method may not be appropriate and so the more general version of the mechanism is preferable.

2.5 Imperfect Information

Imperfect information creates a more difficult problem for obtaining continuation values without altering the incentives of the game. First, the information structure of the game must be more detailed than in standard imperfect information games. Since players need to give valuations at nodes where other players are acting, histories must be partitioned into information sets for each player individually as in Battigalli and Dufwenberg (2005). When dealing with a history that has perfect information for all players, the mechanism applies as before. However, when one or more players has imperfect information, the mechanism has to compensate. Lotteries must still be over histories, but players should be informed only of the probability the lottery selects each of their information sets.[8] Then, to maintain incentives, we must assume that the beliefs about the probability of being at each history in an information set are the same whether the information set was reached in the first or second play through. This is a very strong assumption since in the second play through, the probability of being at a given history may be a function of the lottery in addition to other players' choices.

Further, we still can not ensure that incentives align unless we require a second play through for all realizations of the lottery. If we do not require a second play through when choosing a history reached in the first play through, then the act of playing a second time may reveal information to players that contradicts their information sets.[9] Also, if subjects are playing from a history they've already reached, one would expect any deviation from expected play that took place in the first play through to cause some sort of updating of beliefs about the second play through. We must assume this does not happen.

With these modifications, the intuitions for the perfect information proofs extend to imperfect information. The strategy method also extends to imperfect information. In fact, when using the strategy method, no second play through is required, so

[8] So if there are two histories in an information set for some player, one is chosen by the lottery 5% of the time, the other 10%, then the player is only informed that the lottery will choose a history in the information set 15% of the time.

[9] This is to say that players may be able to distinguish between histories that, according to the structure of the information set, they are not supposed to be able to distinguish.

some of our concerns are mitigated. This makes the strategy method a particularly attractive option when dealing with games of imperfect information.

2.6 Further Discussion of Assumptions

Although the mechanism theoretically maintains incentives of the underlying game, there are cases where the appropriateness of our assumptions must be questioned. We require that players believe that play in the second play through of G' will be identical to play in G and the first play through of G'. In this section we discuss a non-exhaustive number of situations in which it is natural to question the appropriateness of that requirement.

First, consider dealing with indifference and multiplicity of equilibria. The consistency of play is based on resolving indifference in the same direction in both games, but this need not be accurate. Similarly, coordination on one equilibrium in the first play through of G' does not generally require coordination on the same equilibrium, or the projection of that equilibrium onto a subgame, in the second.

Risk preferences are also an issue, as games tend to be inherently risky. Unfortunately the additional payoffs related to the mechanism fundamentally change the risks of the game. For example a risky action in G may seem less so in G' as the different components of payoffs in G' ($u_i(h^T)$, $u_i(h^{T'})$, and $w_{hi}(x_{hi}, h^{T'})$) act much like diversification of a portfolio.

There are also issues of information gathering. There is an implicit assumption that the game is well understood. If players do not understand the game (or the other players in the game), then knowing that parts of the game may be played a second time creates a potential incentive to experiment in the first play-through that is impossible to account for.[10] This issue becomes particularly acute when thinking about repeated games and learning. Therefore, it may be preferable to implement the

[10] We have not proven this conjecture and we would happily be proven wrong

mechanism on a per-stage-game basis when dealing with repeated games.[11] We also suggest random rematching to reduce the incentive to learn a particular opponent's proclivities. To illustrate information issues, consider the following game with utility payoffs:

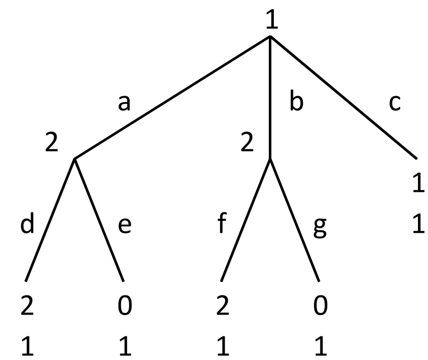

FIGURE 2.2. The Following Game with Utility Payoffs

If player 1 were indifferent between a, b, and c initially, and then the mechanism were introduced, then he would have an incentive not to choose c as both a and b, because they are similar choices, potentially provide information about player 2 that may increase payoffs in the second play through. For example, assume player 1 chooses a and player 2 then chooses d. If the lottery then selects (b), player 1 can use the information about player 2's choice at (a) to inform his choice of x_{hi} at (b).

Similarly, there are issues when the lottery chooses a history that for some players belongs in a non-singleton information set. Here players must form beliefs about their specific history given they are at a particular information set. To guarantee incentive alignment, we need to assume that these beliefs are the same whether the information set was reached naturally or was selected by the lottery. Generally this may be a stretch. One partial solution would be to run a pilot experiment and then

[11] Implementing the mechanism in this fashion would not eliminate the problem but it should reduce the incentive to experiment early on.

use the empirical probabilities within each information set to construct the lottery. Alternatively, this issue can be mitigated by taking lotteries only over nodes where all players have perfect information.[12] In many games this restriction actually does little to reduce the usefulness of the mechanism.[13] Unfortunately, in other games the ability to obtain values at relevant points may be lost.[14]

Social preferences are also an issue. Analysis in this paper assumes that payoffs are in terms of utility, but in experiments, observed payoffs are likely to be in some form of currency, and the mapping from currency allocations to utility is not known to the experimenter. In some cases (for example psychological games) such a mapping may not even exist.

Another issue arises when recognizing that being at a particular information set may not always be a function of previous play. Instead it can be a combination of previous play and a lottery draw. While future options in both cases are identical, the "history" of the game may differ in ways that are critical to certain solution concepts. For example consider the following game with monetary payoffs:

[12] This would also no longer require us to explicitly ask for mixtures in an experiment.

[13] For example, consider a game with two players. First player 1 moves, then player 2 moves without observing player 1's choice, then both players are informed of the other's choice. If we restrict ourselves to only asking players' continuation values before player 1 moves, we can still learn a significant amount about what each player expects the other player to do.

[14] For example, consider a game with three players. First player 1 moves, then player 2 observes this move and makes his own, then player 3 moves without observing anything, finally everyone is informed of all choices made. We would not be able to ask player 2 to report his continuation value after observing player 1's choice as player 3's information set at that point is not a singleton.

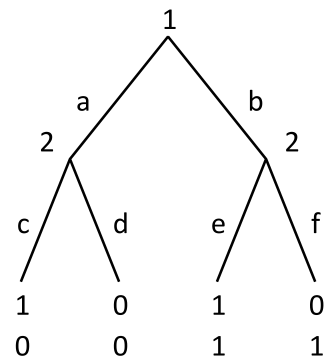

Figure 2.3. The Following Game with Monetary Payoffs

If player 2 found himself at history (a) due to the choice of player 1, he might likely choose d out of spite or revenge. However, if he found himself at (a) due to a lottery draw, such feelings would seem out of place. Indeed there are many situations where the true "history" is as relevant as future options. In these situations it is unlikely that our assumptions about beliefs would hold.

Due to these and related issues we suggest that anyone interested in applying our method run a clean treatment, without asking for continuation values, to use as a comparison with underlying play when implementing the mechanism. If underlying play in the clean treatment, and both play throughs of the mechanism treatment are statistically similar, then experimenters can be more confident in the accuracy of reported continuation values.

Furthermore, it is recommended that subjects be encouraged to play in a risk-neutral fashion. This can be done in practice by using the following methodology, or many similar alternatives. Subjects win white poker chips during the experiment that, at the end of the experiment, will be placed in a bag with some number of black chips. A chip is drawn randomly from the bag and; if it is white, the subject receives $\$M$; if it is black the subject receives less than $\$M$. This procedure is then repeated for each subject. A computerized version of this procedure may be preferable to allow for fractional payoffs. We recommend against using a single lottery for multiple subjects as doing so can create competitive incentives in games that are unintended.[15]

2.7 Concluding Remarks

The mechanism presented here allows experimenters to gather more information from subjects without theoretically changing underlying incentives. However, the assumptions necessary to guarantee the consistency of underlying incentives are strong and in some cases prohibitive. For this reason, empirical research on the validity of our assumptions would be valuable. In cases where assumptions hold, information gathered by the mechanism can be useful in understanding learning processes, overconfidence, bounded rationality, and various belief-dependent processes. The mechanism is fairly general in that it allows experimenters the freedom to look at many information sets in most games, to ignore information sets where continuation values are not desired, and to use the strategy method to learn about values off the path of play.

[15] Consider a lottery where one subject is selected to receive a large prize, and the probability a given subject is chosen is that subject's experimental winnings divided by the total of all subjects. Subjects would still be incentivized to maximize their own winnings, but also to minimize other subjects' winnings, turning the experimental session into a kind of zero-sum game.

Chapter 3

Dealing with Eminent Domain[1]

3.1 Introduction

Eminent domain (ED) is the right of the state to acquire property from a seller in exchange for just compensation. In the United States, ED, is built into the Bill of Rights, as part of the 5th Amendment's takings clause. The clause is intentionally vague, and its interpretation has been largely left to the judicial system. In 2005, ED was given a slightly new interpretation by the US Supreme Court's ruling on the Kelo vs. New London case. Prior to the decision, private companies could not obtain land via the takings clause. Kelo vs. New London shifted this interpretation so that private companies can use eminent domain for development if it can be shown that the project provides public benefits.

This decision has sparked interest in areas of the economic literature that have been dormant for some time. In effect, land assembly is a problem of anti-commons, a buyer must assemble multiple inputs that are perfectly complementary in order to produce output. Failure to obtain one of these inputs prevents production. The Coase Theorem suggests that negotiations in which both parties benefit should be successful in the absence of prohibitive transactions cost. Yet, in a private negotiation setting, many projects that are profitable to a developer may not be realized because of perceived or real difficulties in assembly. Assembly may involve prohibitively high transactions cost due to the length of negotiations or due to private information, which is costly to discover credibly. Additionally, there is a problem of holdup where later sellers can demand high prices for their properties that may make the project unprofitable for a buyer, even though there are positive gains from trade. ED is

[1] Joint work with Carl Kitchens

meant to overcome this problem of holdup. If a buyer has a valuation for the final output in excess of the sum of court determined prices and legal fees, the project may go forward regardless of the sellers' valuations. An alternative solution to the holdup problem is the use of contingent contracts. Hawkins (2011) shows that full refundability eliminates the holdup problem. Both of these solutions have unpleasant properties. In the case of ED, deals may occur when gains from trade are negative, and contingent contracts result in an inequitable distribution of surplus.

In this paper we apply a sequential Nash Bargaining model (Moresi, Salop, and Sarafidis 2008) where a buyer is tasked with purchasing N identical perfectly complementary goods from N queued sellers with exogenous ordering. We examine the scenario using two different mechanisms. In one scenario, contracted prices must be agreed upon between the buyer and each seller. However, these contracts are contingent on successful negotiation with all sellers. In the next protocol, if a buyer and seller are unable to agree on a contracted price, the buyer has the additional option to obtain the property at a predetermined price (under eminent domain) but must incur a transaction cost to use this feature. The Nash Bargaining solution in the first mechanism predicts that conditional on equal bargaining weights, sellers later in the queue obtain lower prices than those earlier in the queue. Our second model predicts that all sellers should obtain equal prices in equilibrium. We test these predictions using a laboratory experiment. Our results show that when using contingent contracts, behavior of subjects in the laboratory was not consistent with Sequential Nash Bargaining. However, when ED was available to the buyer, subject behavior aligned with Sequential Nash Bargaining prediction.

In what follows, we discuss the existing literature on takings and eminent domain, step through our model, formally develop predictions, describe how the theory is tested using a laboratory experiment. We wrap things up by reviewing our results, giving special attention to why Sequential Nash Bargaining failed to predict subject behavior with contingent contracts.

3.2 Related Literature

Recently, ED has gained attention in the economic literature on the heels of the Kelo vs. New London decision. The ruling extended the power of ED to include private companies when development is in the public interest. Kelo was received with varied responses across the states. Many states acted quickly to change their legislation to either clarify ED procedures, or protect property owners. Lopez, Jewell, and Campbell (2009) examine the responses of state legislatures based on the demographics and political tendencies of the state. They find that conservative states tended to enact legislation that would protect land owners, while more progressive states were less likely to change their current legal code.

ED, and the greater problem of land assembly is rooted in the tradition of the anticommons literature. The tragedy of the anticommons, a term coined by Heller (1998) occurs when projects that would be profitable to undertake are avoided because assembly becomes prohibitively expensive. In a sequential land assembly game without contingent contracts the price paid for land to a previous seller is sunk. Because the costs are sunk to the buyer, later sellers can extract the full value of the project from the developer. This results in developers going forward with a suboptimal number of projects. Cournot was the first to analyze how a buyer may face difficulty when acquiring perfectly complementary goods from monopoly sellers. Hawkins (2011) examines the ways to overcome the anticommons problem through refundability and stand alone value. Hawkins shows that full refundability overcomes the anticommons problem. We use this result when developing our contingent contracts model.

One way to think of ED, is that it provides a legal mechanism for overcoming the tragedy of the anticommons. Micili and Segerson (2007) apply a simple Nash bargaining model to the ED setting. Their study utilizes a two period design, in which bargaining can happen today or tomorrow. The authors find that without ED, bargaining will yield two possible sub game perfect equilibria, either all parties agree to sell in the first period, or they wait until the later period. Perhaps the more interesting result of their paper is that when ED is allowed to be used by a developer, all property holders decide to sell their property in the first period. This implies that in equilibrium ED will never be used, and that the credible threat is enough to overcome the holdup problem in perfect information settings.

Models of heterogeneous types have been introduced by Munch (1976) who empirically studies ED during urban renewal in Chicago, IL. Munch predicts that property owners with high values will obtain higher prices, conditional on observables. One difficulty in the empirical estimation of the paper was that privately negotiated prices were unobserved and had to be predicted from observables. Kitchens (2011) also develops a model of multiple seller types. The model predicts that when there is a sufficiently low number of high value sellers, it is optimal for a developer to blanket offer all sellers a low price, letting high value types refuse to sell and thus select into court proceedings. The author then tests this prediction empirically using data from land assembly at a Tennessee Valley Authority reservoir in 1936. The author finds evidence that high type individuals that selected into court proceedings obtained higher prices, while controlling for the selection process.

Several studies have examined sequential bargaining theoretically and experimentally. Moresi, Salop, and Sarafidis (2008) create a model with Nash bargaining occurring sequentially. Their results indicate that if the utility function defined is strictly concave, then bargaining sequentially will yield a lower payoff for the buyer than simultaneous bargaining. The position of the sellers in the queue does not affect the buyer's payoff. Because the model includes sunk cost, seller payments are increasing as the sellers appear later in the queue.

Experimentally, Swope and Schmitt (2008) adopt the model from Miceli and Segerson (2007) and play a take it or leave it ultimatum game. In their experimental design, there is no possibility for the buyer to obtain the property if an agreement cannot be made privately. The primary aim of their model is to test how subjects respond to delay costs in a multi-period bargaining game. They find that as the delay becomes more costly, parties are more likely to accept offers in earlier rounds of negotiation.

Goswami, Noe, and Wang (2005) experimentally explore sequential bargaining. Their work is designed to have a land developer acquire a set of properties from the sellers. The value of the project is common knowledge, and has a simple backward induction equilibrium. Their results indicate that even with a simple structure, subjects have difficulty in reaching the equilibrium due to two primary causes. First, subjects had difficulty with sunk costs, so that a buyer may incur a negative payoff. Second, subjects had difficulty accepting offers that did not equally distribute surplus to the sellers.

George Ng (2010) examines a series of bargaining rules to see which lead to greater efficiency and success when ED is not available to a developer. Each of the rules selected by Ng does not allow for contingent contracting, so holdup cannot be eliminated by design. Ng instead examines which rules mitigate holdup, leading to higher rates of success for the developer. Ng finds that without contingent contracts, a sealed equal-split bargaining rule leads to the highest rate of success.

Our study adds to this literature by including contingent contracts in our baseline. Further, it extends the ED literature by explicitly examining its effects on buyer surplus. While not the main purpose of this study, we are also able to see how contingent contracting mitigates holdup.

3.3 Nash Bargaining

We consider two versions of the following bargaining scenario. A developer (buyer) must acquire property from a number (N) of land owners (sellers) in order to go forward with a project. If all the land is acquired, the developer receives the benefits of the project (V_B) minus the sum of the purchase prices of the land. If a land owner sells his property he receives some transfer payment (P_i) from the developer. If a land owner does not sell his property, he receives his use value of the land(V_s). If and when the developer exercises eminent domain to acquire a parcel of land, the land owner receives a court-determined transfer payment (C) from the developer, and the developer must pay court fees (F). These values are all common knowledge. We assume that sellers have homogenous bargaining weights (α).

3.3.1 Sequential Bargaining

We begin by presenting the solution for the special case where $N = 2$. We solve for the optimal prices for each seller and then give a simple parameterization of the model to provide intuition. Afterward, we generalize to N sellers.

First we solve the following maximization problem for seller 2:

$$\arg\max_{P_2}(V_B - P_1 - P_2)^{(1-\alpha)}(P_2 - V_S)^{\alpha}$$

The relevant first order condition of this maximization problem is:

$$FOC : \frac{\alpha}{P_2 - V_S} - \frac{1-\alpha}{V_B - P_1 - P_2} = 0$$

The value that satisfies this first order condition is then:

$$P_2^* = (1-\alpha)V_S + \alpha(V_B - P_1)$$

Now that we have a value for the optimal P_2 for seller 2, we can solve for P_1 using the following maximization problem:

$$\arg\max_{P_1}(V_B - P_1 - ((1-\alpha)V_S + \alpha(V_B - P_1))^{(1-\alpha)}(P_1 - V_S)^{\alpha}$$

This simplifies to:

$$\arg\max_{P_1}((1-\alpha)(V_B - P_1 - V_S)^{(1-\alpha)}(P_1 - V_S)^{\alpha}$$

Seller 1's first order condition is then:

$$FOC: \frac{\alpha}{P_1 - V_S} - \frac{1-\alpha}{V_B - P_1 - V_S} = 0$$

The optimal value of P_1 is then:

$$P_1^* = (1-\alpha)V_S + a(V_B - V_S)$$

We now assign values to V_B, V_S, and α, as follows:

$$\begin{aligned} V_B &= 1 \\ V_S &= 0 \\ \alpha &= .5 \end{aligned}$$

This parameterization results in the following optimal prices:

$$\begin{aligned} P_1^* &= 1/2 \\ P_2^* &= 1/4 \end{aligned}$$

More generally, to find the Nash bargaining solution when N sellers bargain using contingent contracts, we solve the following maximization problem for the n^{th} seller:

$$\arg\max_{P_n}(1-\alpha)\ln[(1-\alpha)^{(N-n)}(V_B - \sum_{i=1}^{n-1} P_i - P_n - \sum_{j=n+1}^{N} V_S)]$$
$$+\alpha\ln[P_n - V_S]$$

$$FOC \quad : \quad \frac{\alpha}{P_n - V_S} - \frac{(1-\alpha)}{(V_B - \sum_i P_i - P_n - \sum_j V_S)} = 0$$
$$P_n^* \quad = \quad (1-\alpha)V_S + \alpha(V_B - \sum_{i=1}^{n-1} P_i - \sum_{j=n+1}^{N} V_S)$$

Note that P_i's must be at least as large as V_S, and as n increases, V_S terms are replaced by P_i terms. Therefore, P_n^* is weakly decreasing in n. The intuition for this result stems from the fact that any surplus given to the earlier sellers reduces the amount of surplus available to split with subsequent sellers. Therefore, paying an earlier seller an extra dollar reduces the buyer's final payoff by less than a dollar. This creates a pecuniary externality that earlier sellers can exploit to secure higher prices for themselves. This result is driven by the exogenously determined queue. If for example buyers and sellers could endogenously determine their ordering, this result would not hold. Furthermore, if contracts were not contingent, they could fall prey to the holdup problem.

3.3.2 Sequential Bargaining with Eminent Domain

Again, we begin by presenting the solution for the special case where $N = 2$ before generalizing to N sellers. When ED is available to the buyer, the developer can exercise this option if private negotiations stall. In this case the land owner receives a court-determined transfer payment, C, from the developer, and the developer must pay court fees, F.

First we solve the following maximization problem for seller 2:

$$\arg\max_{P_2}(V_B - P_1 - P_2 - (V_B - P_1 - C - F))^{(1-\alpha)}(P_2 - V_S - (C - V_S))^{\alpha}$$
$$\arg\max_{P_2}(-P_2 + C + F)^{(1-\alpha)}(P_2 - C)^{\alpha}$$

The relevant first order condition of this maximization problem is:

$$FOC: \frac{\alpha}{P_2 - C} - \frac{1-\alpha}{P_2 + C + F} = 0$$

The value that satisfies this first order condition is then:

$$P_2^* = C + \alpha F$$

Now that we have a value for P_2, we can solve for P_1 using the following maximization problem:

$$\arg\max_{P_1} (1-\alpha)\ln(V_B - P_1 - (C + \alpha F) - (V_B - C - F - (C + aF)))$$
$$+\alpha \ln(P_1 - V_S - (C - V_S))$$

This simplifies to:

$$\arg\max_{P_1} (1-\alpha)\ln(-P_1 + C + F) + \alpha \ln(P_1 - C)$$

Our first order condition is then:

$$FOC: \frac{\alpha}{P_1 - C} - \frac{1-\alpha}{P_1 + C + F} = 0$$

The optimal value of P_1 is then:

$$P_1^* = C + \alpha F$$

We now assign values to V_B, V_S, and α, as follows:

$$\begin{aligned} V_B &= 1 \\ V_S &= 0 \\ \alpha &= .5 \\ C &= 0 \\ F &= .5 \end{aligned}$$

This parameterization results in the following optimal prices:

$$\begin{aligned} P_1^* &= 1/4 \\ P_2^* &= 1/4 \end{aligned}$$

Generalizing to N sellers, to find the Nash Bargaining solution when the buyer has the option of ED, we solve the following maximization problem for the n^{th} seller:

$$\begin{aligned} &\arg\max_{P_n}[(V_B - \sum_{i=1}^{n-1} P_i - P_n - \sum_{j=n+1}^{N}(C + \alpha F)) \\ &-(V_B - \sum_{i=1}^{n-1} P_i - C - F - \sum_{j=n+1}^{N}(C + \alpha F))]^{(1-\alpha)} \\ &*[(P_n - V_S) - (C - V_S)]^{\alpha} \end{aligned}$$

$$\begin{aligned} &= \arg\max_{P_n}(1-\alpha)\ln[C + F - P_n] + \alpha\ln[P_n - C] \\ FOC &: \frac{\alpha}{P_n - C} - \frac{(1-\alpha)}{(C + F - P_n)} = 0 \\ P_n^* &= C + \alpha F \end{aligned}$$

The solution no longer depends on positioning in the queue or seller private valuations. Each seller faces the same set of choices, either negotiate a private offer or receive an award from the court valued at C. In order for a seller to accept, the buyer only has to offer slightly above the court assessed price. However, going to court causes the buyer to incur a legal fee of F. Each seller in line is negotiating to capture as much of F as possible. Because all sellers face this choice, the equilibrium prices are predicted to be equal when bargaining weights are equal and are independent of previously contracted prices. This result is generally robust to other features of the protocol, such as the contingency of contracts.

3.3.3 Hypotheses

Using the models in this section, we derived three hypotheses to test experimentally.

Hypothesis 1: Total welfare is maximized using either bargaining protocol.

Hypothesis 2: Without eminent domain, prices paid to sellers later in the queue will be lower than those paid to earlier sellers.

Hypothesis 3: With eminent domain, seller prices should be not be correlated with seller position.

Hypothesis 1 is derived using the models in the previous section assuming there are gains from trade. Hypotheses 2 and 3 are derived assuming additionally that all agents have equal bargaining weight.

3.4 Experimental Design

We exploited a within- and between-treatment design to test the hypotheses in the preceding section. Subjects were randomly assigned into groups of five, with one group member assigned the role of a buyer and the rest as sellers. If assigned the role of a seller, the subject was also assigned a position in the bargaining queue.

We allowed for open bargaining as Nash bargaining does not specify the structure of bargaining, and we wanted to capture this feature. Specifically, in a bargaining process between the buyer and a seller, we imposed no order in which sides must make offers and counter offers. For any active buyer-seller pair, either side could propose offers and counter offers until an agreement was reached . An agreement was made if both the buyer and the current seller agreed on the most recent proposal. Once an agreement was made with the first seller, negotiations proceeded to the next seller in the queue.

During the experiment sellers waiting in the queue were locked out of the negotiations and were asked to quietly read the school newspaper. The screen was only active between the buyer and the i^{th} seller. Contingent contract prices between the buyer and previous sellers were visible to the current seller. There were two treatments investigated: a no eminent domain (NED) treatment and an eminent domain (ED) treatment.

In the NED treatment, sellers were given the option to end negotiations. Ending negotiations, or walking away, voided all previous contingent contracts made between buyers and sellers. In this scenario, each seller obtained a payoff equal to their usage value and the buyer received a payoff of zero.

In the ED treatment, the option to acquire properties by force was available to the buyer. If this option was used, there was a set court price, C, as well as known legal fees, F. If the buyer decided to use eminent domain, only the current property was obtained through the ED process. The buyer received the property at a cost of $C + F$ and the seller received C. This did not affect negotiations between the buyer and other sellers in the queue.

The experiment was parameterized as follows. Each buyer had a project value of \$50. In the NED treatment, each seller had a private use value for the property of \$4. In the ED treatment, the court award C, was set to be \$4, and the legal fees charged for the process were set at \$8.50. Given these parameters, neither treatment should have been prey to holdup, and in both treatments gains from trade were available. Furthermore, the "walk away" payoff for each subject was the same as if eminent domain were exercised on all sellers. In this way the outside options were comparable though not identical.

Subjects were randomly drawn from the Economic Science Laboratory database of subjects at the University of Arizona. We implement the experiment using computers running zTree experimental software. We held six sessions of 15-20 subjects each; three to run the contingent contract, NED treatment, three to hold the ED treatment. We had 12 groups of 5 subjects each in the NED treatment and 11 groups of 5 subjects each in the ED, totaling 23 groups (115 subjects). Subjects were awarded a five dollar show up fee in addition to their earnings in the experiment. Subjects were unable to identify the other members of their group. After being seated, subjects were given time to read over the instructions, which were then read aloud by an experimenter.

3.5 Experimental Results

Overall, our results are mixed. At first blush, our theory fails due to a lack of an error structure built into the framework. Qualitatively, on the other hand, the theory does a good job of predicting subject behavior. Efficiency is generally high in both treatments and here is little variation in the prices received by sellers in the ED treatment, as predicted. However, we find no evidence that seller prices are negatively correlated with position in the queue in the NED treatment. We more closely examine each hypothesis in the following sections. We also look at results dealing with empirical bargaining weights as an alternative to our price analysis, so that we may test more direct assumptions on the primitives of the bargaining model.

3.5.1 Hypothesis 1

Our first hypothesis is that welfare is maximized and should be the same in each treatment. Table 3.1 summarizes the welfare results of our experiment. Only one group (8.3%) failed to reach a deal in the NED treatment, and in the ED treatment, ED was only used to collect 3 of the 44 properties (6.8%). While a relatively low percentage of groups were unable to reach a deal, an exact binomial test leaves no room for error, predicting that all groups reach a deal. However, though we trivially reject our theory, we do find that the relative efficiency of the two protocols is very similar (91.67% vs. 93.18%). In this sense, part of the theory is confirmed in that both treatments had statistically identical levels of efficiency. Furthermore, we note that despite not aligning perfectly with theory, results of over 90% efficiency seem pretty good, especially for the laboratory setting. It should therefore be noted that from a non-statistically-strict standpoint, Hypothesis 1 actually does very well.

3.5.2 Hypotheses 2 and 3

Hypothesis two and three examine the relationship between seller position in the queue and treatment status. In the NED treatment, sellers who bargain later in the queue should receive lower prices. In the ED treatment; prices should be equal throughout the queue.

In the NED treatment, there is a great deal of variation in the observed data, while in the ED treatment, prices range over a smaller set of values. The first seller in the NED treatment obtained an average price of $6.88, with the maximum price received equal to $16 and the minimum equal to $2.50. Sellers later in the queue fared better, in contrast to the prediction of the theory. In the ED treatment, the first seller also averaged $6.88, with the prices ranging from $3.50 - $12. Table 3.2 summarizes the observed prices by seller position for each treatment. The prices observed are quite similar across the treatments. In fact, a Mann-Whitney Test fails to reject differences in the underlying price distribution of each treatment. This could be due to a variety of factors, which we will explore later. Because we do not see a downward trend in the prices, we turned to a series of pair wise comparisons of the prices by seller position. We perform a Mann-Whitney Test on the prices by seller position within each treatment. In the NED treatment we find no evidence of differences in the observed prices by seller position. The full results are displayed in Table 3.3. The table also shows the pair wise comparisons for the ED treatment. The ED model predicts that all sellers obtain the same price. Our results show that there are no significant differences for adjacent sellers; however, there are statistical differences between Seller 1 and Seller 4, as well as Seller 2 and Seller 4 (at the 10% level).

3.5.3 Empirical Bargaining Weight

The assumption of equal bargaining weights for all subjects is a rather strong one. Worse, when combined with the theory it suggests no error structure with which to design statistical tests. However, if we look at the one primitive of the theory that is not induced by experimental design (bargaining weight), and assume that all subjects bargaining weights are drawn from the same normal distribution, an error structure becomes available.

In this subsection we assume that the bargaining weights of agents are drawn from the same (normal) distribution. We need make no assumptions about beliefs about other sellers' bargaining weights as future sellers' bargaining weights have no impact on P_n^*. We show this using the following maximization problem for the n^{th} seller.

$$\begin{aligned}&\arg\max_{P_n}(1-\alpha_n)\ln[(\prod_{h=n+1}^{N}(1-\alpha_h))(V_B-\sum_{i=1}^{n-1}P_i-P_n-\sum_{j=n+1}^{N}V_S)]\\&+\alpha_{N-1}\ln[P_{N-1}-V_S]\end{aligned}$$

$$
\begin{aligned}
FOC \quad &: \quad \frac{\alpha_n}{P_n - V_S} - \frac{(1-\alpha_n)(\prod_{h=n+1}^{N}(1-\alpha_h))}{(\prod_{h=n+1}^{N}(1-\alpha_h))(V_B - \sum_i P_i - P_n - \sum_j V_S)} = 0 \\
P_n^* \quad &= \quad (1-\alpha_n)V_S + \alpha_n(V_B - \sum_{i=1}^{n-1} P_i - \sum_{j=n+1}^{N} V_S)
\end{aligned}
$$

We use this result, as well as the observed prices and experimental parameters to back out empirical bargaining weights (summarized in Table 3.4). We then use these empirical bargaining weights to construct pair-wise comparisons between sellers using the Mann-Whitney Test.

Table 3.5 summarizes the comparisons. We find that in the NED treatment, we must reject that the data generating process is from the same distribution for Sellers 1 and 3 and Sellers 1 and 4. In the ED treatment we must also reject the assumption, as we find significant differences between Sellers 2 and 4 and Sellers 1 and 4.

Due to these rejections, we turn to the chat logs for further insight. We recognize some participants may be shrewder than others when it comes to bargaining. We propose that a way to measure one aspect of this shrewdness is to examine the relationship between the number of messages sent by the buyer and seller on the contracted price. To explore this relationship formally, we regress the contract price on treatment and the number of messages sent by the buyer and the seller. The results from this specification are presented in Table 3.6. Sellers messages seem to have the largest effect, as an additional seller message in associated with a price that was almost 9 cents higher. This means that taking the time to communicate was very valuable for a seller. An additional buyer message has a negative relationship with price, but the relationship is not statistically significant.

We further run similar regressions, substituting the empirical bargaining weight for price. We find similar results as before. Seller messages increase bargaining power, for each additional message sent, bargaining power increases by .008. Buyer messages have the expected sign, but are less statistically significant. Overall, there seems to be a correlation between messages sent and surplus captured. The direction of causation, however, is not identifiable and suggests further study may be fruitful.

3.6 Conclusions

There does not seem to be a significant difference when using eminent domain in our setting. Contingent contracts seem to be enough to eliminate hold-up. Furthermore, we did not see the difference in prices as a function of queue position that sequential Nash bargaining theory predicts. This suggests that when sufficient gains from trade exist, eminent domain may not be necessary provided that buyers and sellers can contract prices that are contingent on agreements with sellers later in the queue. It remains to be seen what differences may arise in other settings, particularly when there are little or even negative gains from trade.

We also generally observe multiple deviations from theory in both treatments. Evidence suggests that these deviations may result from certain behavioral factors including bounded rationality and other-regarding preferences. Since our design is not specifically tailored to distinguish between these various factors, additional experimental study would likely be fruitful.

3.7 Tables

	% Max Welfare	
	Projected	Actual
NED	100	91.67*
obs.	-	12
ED	100	93.18*
obs.	-	44
NED-ED	-	-1.52

* - significantly different at any confidence level

TABLE 3.1. Welfare

	NED		ED	
	Price	Std. Dev.	Price	Std. Dev.
Seller 1	687.9	(434.0)	688	(291.2)
Seller 2	757.1	(339.3)	704.4	(226.7)
Seller 3	875.4	(361.2)	840.7	(202.8)
Seller 4	774.8	(271.7)	843.4	(223.2)

TABLE 3.2. Prices

	NED		ED	
	MW Stat	P Value	MW Stat	P Value
Seller 1 vs 2	-0.658	0.511	-0.417	0.677
Seller 1 vs 3	-1.316	0.188	-1.248	0.212
Seller 1 vs 4	-0.888	0.375	-1.649	0.099
Seller 2 vs 3	-0.363	0.717	-1.238	0.216
Seller 2 vs 4	-0.066	0.948	-1.671	0.095
Seller 3 vs 4	0.560	0.576	-0.725	0.469

Table 3.3. Within Treatment Price Comparison

	NED		ED	
	Weight	Std. Dev.	Weight	Std. Dev.
Seller 1	0.092	(0.131)	0.339	(0.343)
Seller 2	0.138	(0.114)	0.358	(0.267)
Seller 3	0.221	(0.187)	0.519	(0.239)
Seller 4	0.220	(0.153)	0.522	(0.263)

Table 3.4. Bargaining Weights

	NED		ED	
	MW Stat	P Value	MW Stat	P Value
Seller 1 vs 2	-1.149	0.251	-0.417	0.677
Seller 1 vs 3	-1.806	0.071	-1.248	0.212
Seller 1 vs 4	-2.003	0.045	-1.649	0.099
Seller 2 vs 3	-0.952	0.341	-1.238	0.216
Seller 2 vs 4	-1.248	0.212	-1.671	0.095
Seller 3 vs 4	-0.164	0.870	-0.725	0.469

TABLE 3.5. Within Treatment Bargaining Weight Comparison

Y Variable	Bargaining Weight		Price	
	Coef.	Std. Err.	Coef.	Std. Err.
Seller Messages	0.0087**	(0.0030)	8.90*	(4.87)
Buyer Message	-0.0041*	(0.0022)	-2.23	(5.08)
Treatment	0.2771**	(0.0681)	0.93	(100.57)
Constant	0.1350**	(0.0393)	740.18**	(92.84)
N	92		92	
R squared	0.3264		0.0502	

OLS clustered on group
* - significant at 10% level
** - significant at 5% level

TABLE 3.6. Messages

Appendix A
Proofs

A.1 Proof of Proposition 5

Proof. To prove the theorems in this section for our second HBLFM, we characterize the value of various payoffs conditioned on a backward induction path. We then show that, given this characterization, we can assign a probability that our models predict a given path, We then show that this probability is weakly increasing in each ω_i.

Let $G = (I, H, L, ((u_i); i \in I))$ be a finite extensive form general game of perfect information, where I is the set of players, H is the set of histories, $L : H^P \to I$ is the mapping of who moves at each partial history where $H^P \subset H$ is the set of intermediate or partial histories, and $u_i : H^T \to \mathbb{R}$ is player i's payoff function where $H^T \subset H$ is the subset of terminal histories. Without loss of generality, let each $u_i(h^T)\tilde{}U(0,1)$. Now let u_i^* refer to the payoff to player i on a unique backward induction path h^*. Now, before u_i's are realized, u_i^* and h^* simply refer to possible outcomes of random variables.

Notice that regardless of where h^* ends up on the game tree, for at least one player i, u_i^* will be larger than at least one other u_i' (a payoff player i can get by deviating from backward induction) unless there is only one terminal history, in which case our models trivially predict that history. In fact, if we knew and fixed the strategies of all other players, we can easily compute the number of possible deviation payoffs by counting the number of alternative terminal histories that could be reached considering all possible deviations by player i. We could call this number d_i, and we could say that u_i^* is the maximum of $d_i + 1$ standard uniform random variables, and so $u_i^*\tilde{}B(d_i + 1, 1)$. We can then say that on average $u_i'\tilde{}U(0, u_i^*)$ for all d_i u_i''s. We can further observe that for any possible value of d_i, $B(d_i + 1, 1) \succ U(0,1) \succ U(0, u_i^*)$

over the interval $(0, 1)$.

Turning to our second HBLFM, we will show that whenever h^* is in $H|h^{VT}$, the median of the distribution of $u_i|h(h^{VT})$ is weakly greater than the median of the distribution of any other $u_i|h(h^{VT\prime})$. The expected distribution of $u_i|h(h^{VT})$ is a weighted average of the distributions of some number of random variables. Since we don't know the position on the game tree of any of these random variables, the weighting drops out.[1] The distributions of the possible relevant random variables are:[2]

$$u_i^* \tilde{} B(d_i + 1, 1)$$
$$u_i^- \tilde{} U(0, 1)$$
$$u_i' \tilde{} U(0, u_i^*)$$

Note that the distribution of the average of one draw from $B(d_i + 1, 1)$ and d_i draws from $U(0, u_i^*)$ is the same as the distribution of the average of $d_i + 1$ draws from $U(0, 1)$. Also, since $U(0, 1)$ dominates $U(0, u_i^*)$, the distribution of the average of one draw from $B(d_i + 1, 1)$ and less than d_i draws from $U(0, u_i^*)$ dominates the distribution of the average of an equal total number of standard normal random variables. Since the distribution including u_i^* weakly dominates a distribution with a median of $1/2$, we know that it's median is weakly greater than $1/2$. Meanwhile, all medians of distributions composed solely of u_i^-'s and u_i''s must have medians weakly below $1/2$. This means that any $u_i|h(h^{VT})$ including u_i^* will be greater than some other $u_i|h(h^{VT\prime})$ at least half of the time.

This is not quite enough to say that any $u_i|h(h^{VT})$ including u_i^* will be greater than the maximum of n $u_i|h(h^{VT\prime})$'s with more than $\frac{1}{n+1}$ probability, since a number of high variance, low median draws are likely to have a high maximum. However,

[1] The expected weight on each random variable is the same, and is equal to one over the number of random variables.

[2] Here u_i^- is the payoff to player i at some history that player i cannot deviate to.

since positioning on the tree is random, the distribution of the number of random variables in $u_i|h(h^{VT})$ is the same as in $u_i|h(h^{VT\prime})$, and we showed $u_i|h(h^{VT})$ including u_i^* dominates any $u_i|h(h^{VT\prime})$ when the number of random variables is the same, we can say that for any combination of numbers of random variables in both $u_i|h(h^{VT})$ and some number of $u_i|h(h^{VT\prime})$'s, the probability of $u_i|h(h^{VT})$ being the maximum is weakly greater than the same probability if the random variables in $u_i|h(h^{VT})$ were all replaced with u_i^-'s. Since the latter probability, when summed over all possible distributions of the numbers of random variables in the averages, must be at least $\frac{1}{n+1}$, then the former probability when summed in the same fashion must be weakly greater than $\frac{1}{n+1}$.

Finally, since $u_i|h(h^{VT})$ including u_i^* is chosen with more than $1/T$ probability in expectation, actions leading to u_i^* are taken with more than $1/T$ probability, and so u_i^* is reached with more than $1/T$ probability. So, proposition 5 must be true for HBLFM#2.

For our first HBLFM, we note that the only time actual payoffs are relevant is when their terminal history is visible. This means we don't have to worry about combining payoffs in any way. Now, if we don't condition on the backward induction solution, for any values of the ω_i's and β_{ji}'s, each terminal history has some probability of being reached, and the average probability must be $1/T$. Now, for some terminal history, instead of being drawn from a standard uniform, let its payoff u_i^* be drawn from a beta distribution as described in the previous proof. Since such a beta distribution dominates a standard uniform, the probability that the terminal history in question will be reached must weakly increase.[3] Therefore the average probability over all possible terminal histories must be greater than $1/T$. So, proposition 5 must be true

[3] In all visual games where u_i^* is relevant, and it is reachable given some set of moves by player i, increasing its value will weakly increase the likelihood of its terminal history being on the backward induction path of such visual games, since the relevance and "reachability" of any terminal history are independent of the value of u_i^*. We can also ignore the effect of u_i' (vs. u_i^-) since if the terminal history in question is more likely to be reached among u_i^-'s it must necessarily be among some combination of u_i''s and u_i^-'s.

for HBLFM#1. □

A.2 Proof of Theorem 8

Proof. Here we prove Theorem 3 using a direct appeal to the payoffs of G' and G.

Let $v_i(h)$ be player i's continuation value at h with respect to u_i. Let $v_i'(h)$ be player i's continuation value with respect to u_i'. Let $v_i(a(h))$ be player i's continuation value after taking action $a(h)$.[4] Further, we assume that these values sufficiently represent player i's preferences over actions and in turn player i's play in G. Similarly, let $v_i'(a(h))$ be player i's continuation value with respect to u_i'. These v_i''s represent preferences and play in G'.

First we can look at the second play through of G' compared to G. Once a history has been selected, and x_{ki} has been chosen. The expected payoff of choosing some $a(h)$ is $\left(\sum_{H^{T'}}(\delta_h + \delta_h' \frac{x_i - \underline{u_i}}{\overline{u_i} - \underline{u_i}}) u_i(h^{T'})\Gamma_i(h^T|a(h))\right) + \frac{1}{2}\frac{(\overline{u_i} - x_i)^2}{\overline{u_i} - \underline{u_i}}$. In G the expected payoff would be simply $\sum_{H^{T'}} u_i(h^{T'})\Gamma_i(h^T|a(h))$. One being a linear transformation of the other, incentives are clearly aligned, and preferences should remain the same.

Now we will show that for any two possible actions at h ($a^1(h)$ and $a^2(h)$), if $v_i(a^1(h)) > (=)(<)\ v_i(a^2(h))$, then $v_i'(a^1(h)) > (=)(<)\ v_i'(a^2(h))$. This condition implies no preference between two actions will shift when moving from G to G' (in the first play through).

Recall that $u_i'(h^T, \overrightarrow{h}^{T'}) = u_i(h^T) + W_i(X_i, \overrightarrow{h}^{T'})$. We can say,

$$v_i(a(h)) = \sum_{H^T} u_i(h^T)\Gamma_i(h^T|a(h))$$

Similarly,

$$v_i'(a(h)) = \sum_{H^T}(u_i(h^T) + E[W_i(X_i, \overrightarrow{h}^{T'})])\Gamma_i(h^T|a(h))$$

[4] *We make this distinction as the histories that $a(h)$ leads to may also reveal information. We want the value after the action but before any new information is gleaned.*

The expectation term ($E[W_i(X_i, \overrightarrow{h}^{T'})]$)can be expanded as follows:

$$\sum\nolimits_{j\in H}\left(\beta_{ji} * \sum\nolimits_{H^{T'}} \delta_{ji} * u_i(h^{T'}) + \delta'_{ji} * w_{ji}(x_{ji}, h^{T'})\Gamma_i(h^{T'}|j, h^T)\right)$$

Now, notice that the only times h^T is relevant to $E[W_i(X_i, \overrightarrow{h}^{T'})]$ are when the j chosen by the lottery precedes h^T (in which cases $h^{T'} = h^T$). Further, the only times $a^1(h)$ and $a^2(h)$ are relevant are when j precedes h. This relies critically on our assumption that Γ_i does not change in the second play through as a function of what you did in the first. Otherwise $h^{T'}$ could be influenced by h^T in various ways.

Now suppose that $v_i(a^1(h)) > v_i(a^2(h))$ and that x_{hi}'s are chosen according to $v_i(a^2(h))$. We will call $v_i'^2(a^1(h))$ the value in G' of choosing x_{hi}'s according to $v_i(a^2(h))$ but playing according to $v_i(a^1(h))$. Clearly $v_i'(a^1(h)) > v_i'^2(a^1(h))$ since we are no longer maximizing with our choice of x_{hi}'s. We can then write the difference between $v_i'^2(a^1(h))$ and $v_i'(a^2(h))$ as:

$$\begin{aligned}&\sum\nolimits_{H^T}(u_i(h^T) + E[W_i(X_i^2, \overrightarrow{h}^{T'})])\Gamma_i(h^T|a^1(h))\\ -&\sum\nolimits_{H^T}(u_i(h^T) + E[W_i(X_i^2, \overrightarrow{h}^{T'})])\Gamma_i(h^T|a^2(h))\end{aligned}$$

(Where X_i^2 is the set of x_{hi}'s chosen according to $v_i(a^2(h))$)

Pulling the u_i terms out we get:

$$\begin{aligned}&v_i(a^1(h)) - v_i(a^2(h))\\ &+\sum\nolimits_{H^T} E[W_i(X_i^2, \overrightarrow{h}^{T'})]\Gamma_i(h^T|a^1(h))\\ &-\sum\nolimits_{H^T} E[W_i(X_i^2, \overrightarrow{h}^{T'})]\Gamma_i(h^T|a^2(h))\end{aligned}$$

The latter sums are mostly identical, the only differences arise when the information set chosen by the lottery precedes h. So, after cancelling, we get:

$$\begin{aligned}&v_i(a^1(h)) - v_i(a^2(h))\\ &+\sum\nolimits_{j\subset h, j\in H} \beta_{ji} * (\delta_{ji} + \delta'_{ji} * \frac{x_{ji}^2 - \underline{u_i}}{\overline{u_i} - \underline{u_i}}) * (v_i(a^1(h)) - v_i(a^2(h)))\end{aligned}$$

The remaining terms in the sum are when $j \subset h$, and when $r_{ji} < x_{ji}^2$. In such cases weighted $u_i(h^T)$ terms remain. The expectations of those terms are the v_i's.

or,

$$\left(1 + \sum\nolimits_{j \subset h, j \in H} \beta_{ji} * (\delta_{ji} + \delta'_{ji} * \frac{x_{ji}^2 - \underline{u_i}}{\overline{u_i} - \underline{u_i}})\right) (v_i(a^1(h)) - v_i(a^2(h)))$$

Since $v_i(a^1(h)) > v_i(a^2(h))$, $v'_i(a^1(h)) > v'_{i2}(a^1(h)) > v'_i(a^2(h))$.

The other two cases (=)(<) follow the same line of reasoning. Therefore, assuming ties are broken in the same direction in G' as in G, play in G' is a projection of play in G. $\square$

References

[1] Abreu, D. and Rubenstein, A. The Structure of Nash Equilibrium in Repeated Games with Finite Automata. *Econometrica*, 56:1259-1281, 1988.

[2] Block, M., Bader, M., Tapia, E., Ramirez, M., Gunnarsson, K., Cuevas, E., Zaldivar, D., Rojas, R. Using Reinforcement Learning in Chess Engines. *Research in Computing Science*, 35:31-40, 2008.

[3] Brooks, R. Elephants Don't Play Chess. *Robotics and Autonomous Systems*, 6:3-15, 1990.

[4] Cardella, E. Making Better Decisions: Playing vs Observing Skillful Decisions Makers. *Working Paper.*

[5] Chase, W. and Simon, H. Perception in Chess. *Cognitive Psychology*, 4:55-81, 1973.

[6] Dufwenberg, M., Sundaram, R., Butler, D. Epiphany in the Game of 21. *Journal of Economic Behavior & Organization*, 75:132-143, 2010.

[7] Ewerhart, C. Backward Induction and the Game-Theoretic Analysis of Chess. *Games and Economic Behavior*, 39:206-214, 2002.

[8] Gneezy, U., Rustichini, A., Vostroknutov, A. Experience and Insight in the Race Game. *Journal of Economic Behavior & Organization*, 75:144-155, 2010.

[9] Haller, H. and Holden, S. A Letter to the Editor on Wage Bargaining. *Journal of Economic Thoery*, 52:232-236, 1990.

[10] Jehiel, P. Limited Foresight May Force Cooperation. The Review of Economic Studies, 68:369-391, 2001.

[11] Johnson, E., Camerer, C., Sen, S., Rymon, T. Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining. *Journal of Economic Theory*, 104:16-47. 2002.

[12] Kochov, A. A Model of Limited Foresight. Job Market Paper, 2009.

[13] Levitt, S., List, J., Sadoff, S. Checkmate: Exploring Backward Induction Among Chess Players. *American Economic Review*, Forthcoming.

[14] Palacios-Huerta, I., Volij, O. Field Centipedes. *American Economic Review*, 99:1619-1635, 2009.

[15] Ramanujan, R., Sabharwal, A., Selman, B. On Adversarial Search Spaces and Sampling-Based Planning. *20th ICAPS*, pp. 242-245, Toronto, 2010.

[16] Rosenthal, R. Games of Perfect Information, Predatory Pricing and the Chain Store Paradox. *Journal of Economic Thoery*, 25:92-100, 1981.

[17] Rubenstein, A. Perfect Equilibrium in a Bargaining Model. Econometrica, 50:97-109, 1982.

[18] Rubenstein, A. Finite Automata Play the Repeated Prisoner's Dilemma. *Journal of Economic Theory*, 39:83-96, 1986.

[19] Selten, R. The Chain Store Paradox. *Theory and Decision*, 9:127-159, 1978.

[20] Thrun, S. Learning to Play the Game of Chess. *Advances in Neural Information Processing Systems 7*, 1995.

[21] Wichardt, P. Modelling Equilibrium Play as Governed by Analogy and Limited Foresight. *Games and Economic Behovior*, Forthcoming.

[22] Allen, Franklin, "Discovering personal probabilities when utility functions are unknown," *Management Science*, 33:542–544, 1987.

[23] P. Battigalli and M. Dufwenberg, "Dynamic Psychological Games," *Journal of Economic Theory*, 144:1-35, 2009.

[24] G. Becker, M. DeGroot and J. Marschak, "Measuring utility by a single response sequential method," *Behavioral Science* 9:226–236, 1964.

[25] Grether, David, "Testing Bayes rule and the representativeness heuristic: Some experimental evidence," *Journal of Economic Behavior & Organization*, 17:31–57, 1992.

[26] Karni, Edi, "A Mechanism for Eliciting Probabilities," *Econometrica*, 77:603–606, 2009.

[27] Mobius, Markus M., Muriel Niederle, Paul Niehaus and Tanya Rosenblat, "Self-Confidence Management: Theory and Experimental Evidence", Working Paper, May 2010.

[28] Buchanan, J., and Yoon, Y. A Smithean Perspective on Increasing Returns. *Journal of the History of Economic Thought,* 22: 43-48, 2000.

[29] Fehr, E., and Schmidt, K. A Theory Of Fairness, Competition, and Cooperation. *The Quarterly Journal of Economics,* 144:817-868, 1999.

[30] Fehr, E., and Gächter, S. Fairness and Retaliation: The Economics of Reciprocity. *The Journal of Economic Perspectives,* 14:159-181, 2000.

[31] Goswami, G., Noe, T., and Wang, J. Buying up the Block: An Experimental Investigation of Capturing Economic Rents through Sequential Negotiations. Working Paper, March 2005.

[32] Hawkins, J. Mitigating Hold-up Through Complementarities and Refundability. Working Paper, 2011.

[33] Heller, M., and Eisenberg, R. Can Patents Deter Innovation? The Anticommons in Biomedical Research. *Science,* 280:698-701, 1998.

[34] Kitchens, C. The Use of Eminent Domain in Land Assembly: A Case Study of the TVA. Working Paper, 2011.

[35] Levitt, S., List, J., and Sadoff, S. Checkmate: Exploring Backward Induction Among Chess Player. *American Economic Review,* forthcoming, (2010).

[36] Libecap, G. Chinatown: Owens Valley and Western Water Reallocation - Getting the Record Straight and What It Means for Water Markets. *Texas Law Review,* 83:2055-2089, 2005.

[37] Lopez, E., Jewell, R., and Campbell, N. Pass a Law, Any Law, Fast! State Legislative Responses to the Kelo Backlash. *Review of Law & Economics,* 5.1:5, 2009.

[38] Miceli, T., and Segerson, K. A Bargaining Model of Holdouts and Takings. *American Law and Economics Review,* 9:160-174, 2007.

[39] Moresi, S., Salop, S., and Sarafidis, Y. A Model of Ordered Bargaining with Applications. Working Paper, October 2008.

[40] Munch, P. An Economic Analysis of Eminent Domain. *The Journal of Political Economy,* 84:473-497, 1976

[41] Ng, G. Multilateral Bargaining with Incomplete Information in the Anticommons. Working Paper, 2010.

[42] Prasnikar, V., and Roth, A. Considerations of fairness and strategy: Experimental data from sequential games. *The Quarterly Journal of Economics,* 107:865-888, 1992.

[43] Rosenthal, R. Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox. Journal of Economic Theory, 25:92-100, 1981.

[44] Scheiber, H. Property Law, Expropriation, and Resource Allocation by Government: The United States, 1789-1910. *The Journal of Economic History,* 33:232-251, 1973.

[45] Swope, K., Schmitt, P., Cadigan, J., and Shupp, R. An Experimental Study of the Holdout Problem in a Multilateral Bargaining Game. Us Naval Academy Department of Economics Working Papers, 21, 2008.