IHS Economics Series
Working Paper 12
September 1995

Hard Bargains and Lost Opportunities

Ken Binmore
Chris Prouix
Larry Samuelson
Joe Swierzbinski

Impressum

Author(s):
Ken Binmore, Chris Prouix, Larry Samuelson, Joe Swierzbinski

Title:
Hard Bargains and Lost Opportunities

ISSN: Unspecified

1995 Institut für Höhere Studien - Institute for Advanced Studies (IHS)
Josefstädter Straße 39, A-1080 Wien
E-Mail: office@ihs.ac.at
Web: www.ihs.ac.at

All IHS Working Papers are available online:
http://irihs.ihs.ac.at/view/ihs_series/
This paper is available for download without charge at:
https://irihs.ihs.ac.at/id/eprint/852/

Institut für Höhere Studien (IHS), Wien
Institute for Advanced Studies, Vienna

Reihe Ökonomie / Economics Series No. 12

HARD BARGAINS AND LOST OPPORTUNITIES

Ken Binmore, Chris Proulx, Larry Samuelson, and Joe Swierzbinski

Hard Bargains and Lost Opportunities

Ken Binmore, Chris Proulx, Larry Samuelson, Joe Swierzbinski

Reihe Ökonomie / Economics Series No. 12

September 1995

Ken Binmore
Department of Economics
University College London
Gower Street
London WC1E 6BT, UK

Chris Proulx
Department of Economics
University of Michigan
Ann Arbor
Michigan 48109 , USA

Larry Samuelson
Department of Economics
University of Wisconsin
Madison
Wisconsin 53706, USA

Joe Swierzbinski
Department of Economics
University College London
Gower Street
London WC1E 6BT, UK

Institut für Höhere Studien (IHS), Wien
Institute for Advanced Studies, Vienna

Abstract

We examine an *Outside Option Game* in which player *I* submits a claim for a share of a cake while player *II* simultaneously either makes a claim or chooses to opt out. If player *II* opts out, then she receives an opt-out payment while player *I* receives nothing. If player *II* opts in and if the claims total less than the cake, then each player receives his or her claim plus half of the surplus. If the claims total more than the cake, both players receive zero. Tension arises in this game between player *I*'s desire to seek as large a share of the cake as possible and the necessity of providing player *II* with a sufficiently large payoff to ensure that she will opt in. Economic theories that stress efficiency predict that player *II* will opt in. We argue that trial-and-error learning processes can teach the competitive skills needed to secure large shares of the cake more effectively than the cooperative skills needed to ensure that the cake is available to be divided. As a result, outcomes will arise in which player *II* opts out, especially when the payment from doing so is attractive. We conduct experiments in which player *II*s commonly opt in when their opt-out payment is small, but frequently opt out for larger opt-out payments.

Keywords

Bargaining, Coase Conjecture, Evolutionary Games, Drift

JEL-Classifications

C70, C72

Comments

First draft September, 1993. The support of National Science Foundation is gratefully acknowledged, especially for funding the experiments under Grant SES-8821521. We also thank the University of Michigan for founding the Michigan Economic Laboratory, where the experiments were conducted. The work was completed with the assistance of funding from the British Economic and Social Research Council `Beliefs and Behaviour Programme´- During its progress, Binmore and Samuelson enjoyed the hospitality and support of the Department of Economics at the University of Bonn (Deutsche Forschungsgemeinschaft, Sonderforschungsbereich 303) and the Institute for Advanced Studies at the Hebrew University of Jerusalem. We are grateful to both.

1 Introduction

Orthodox economic theory assumes that the entrepreneurial spirit will call forth individuals or institutions to exploit gains from trade whenever they exist. Implicit in this assumption is the understanding that all parties to any potential deal will receive a share of the surplus that is adequate to ensure their cooperation. This will be the case, for example, if they can count on distributional issues being settled by some variant of the split-the-difference principle, as commonly assumed in the labor economics literature (e.g. Macdonald and Solow [31]).

But why would an optimizing agent be satisfied to split the difference if he has the power to force a less equitable deal? Aggressive bargaining may put the whole deal at risk, but brinkmanship is what hard bargaining is all about. It is precisely by making credible threats to delay or withhold agreement that the strong are able to extort concessions from the weak. Indeed, in a world of perfectly rational agents, a hard bargainer will judge his tactics so finely that he extracts the last drop of surplus from his opponent consistent with her continuing in the partnership. Equilibrium offers in the Rubinstein bargaining model [37, 41], for example, are chosen to make the respondent indifferent between accepting and refusing. In equilibrium, she accepts and the resulting outcome is efficient. But only a small misjudgment on either side would be enough to delay agreement for at least one period. Rubinstein's perfectly rational agents are therefore permanently poised on the edge of a disagreement.

The modeling of agents as perfectly rational is often justified by arguing that learning, imitation, or the discipline of the market will cause low-payoff behavior to be supplanted over time by high-payoff behavior (e.g. Alchian [1], Friedman [22]). If enough money is involved, the school of hard knocks may indeed eventually teach its graduates to approximate the behavior of the idealized agents of orthodox economic theory, but there is no guarantee that they will succeed in precisely reproducing all the fine details of perfectly rational bargaining strategies. Instead of reaching a compromise just short of disagreement, one must consequently expect that sometimes a bargainer will overreach himself just enough to push his opponent over the edge. The chance of reaching an efficient deal will then be lost. Even in a world that is only marginally imperfect, a tension therefore arises between the intuition that agents will learn to drive a hard bargain and the claim that economic opportunities will seldom remain unexploited.

The tradition of Coase [17] resolves this tension between efficiency and distribution in favor of efficiency. Transaction cost economics, which views new forms of property rights and contracts as arising in response to existing inefficiencies, seconds this choice (Williamson [47]). The experimental work of Harrison and McKee [25] and Hoffman and Spitzer [27, 28] also points in the same direction.

This paper provides experimental and theoretical evidence for the contrary claim that interesting situations exist in which income-maximizing agents may fail to learn to exploit all gains from trade. For this purpose, we study a simple bargaining game in which one and only one of the two bargainers can opt out and so precipitate an inefficient outcome. Standard game-theoretic arguments applied to this Outside Option Game predict that the outcome will be efficient, but our experiments show that subjects frequently learn instead to opt out.

Critics of game-theoretic reasoning in a bargaining context see nothing paradoxical in its failing to predict how actual people behave in the laboratory (Thaler [44]). Their usual line is that people do not optimize when resolving distributional questions, but simply operate whatever social norm happens to be focal. We do not doubt that inadequately inexperienced or poorly motivated subjects behave more or less as such critics maintain. On the other hand, there is much evidence that the behavior of adequately motivated subjects in simple games often shifts in the direction predicted by optimizing theories—provided that they have ample opportunity to benefit from trial-and-error learning. When an optimizing theory fails to perform well under such circumstances, we therefore prefer not to abandon the theory but to reconsider its theoretical basis.

In previous work on the Ultimatum Game (Binmore *at al* [5]), we argued that apparently anomalous experimental results can sometimes be explained by paying close attention to the adjustment process by which the players learn to play a laboratory game. In the Ultimatum Game, the mechanics of such processes can lead subjects who are essentially income-maximizers to Nash equilibria that the literature on refinements of Nash equilibrium is unanimous in rejecting.

We pursue a similar line in the current paper by studying a simple model of adaptive learning in our Outside Option Game. We find that subjects are taught the competitive skills needed to secure a large share of the cake much faster than the cooperative skills needed to ensure that a cake is available for splitting. As a result, agents learn that it is not worth trying to cooperate when their outside opportunities are sufficiently attractive. Potential gains from trade then remain unexploited.

The following section introduces the Outside Option Game. Section 3 discusses the relationship of our work to previous results. Section 4 analyses the behavior of agents who must learn to play the Outside Option Game in an imperfect world. Section 5 reports the results of an experiment whose instructions appear in an appendix. Section 6 briefly summarizes our conclusions.

2 The Outside Option Game

Our point of departure is the work of Binmore *et al* ([6, 13]), who conducted experimental studies of Rubinstein bargaining models in which players alternate in making offers as often as they please, with the sum of money available for division shrinking fractionally after each refusal. Each player could abandon the negotiations in favor of his best outside option whenever he had just refused an offer. The outside options were inefficient in that the sum of opt-out payments was always smaller than the current surplus to be divided.

When both players are patient, two contending principles for dividing the surplus in such bargaining problems are commonly considered:

Split-the difference: This is the outcome obtained by applying the symmetric Nash bargaining solution after placing the *status quo* at the pair of outside options. It assigns each player his or her outside option plus half the remaining surplus.

Deal-me-out: This is the outcome to which one is led by applying the Rubinstein bargaining theory in the presence of outside options (Binmore *et al* [7, 8, 13]). Players split the surplus fifty-fifty regardless of their outside options as long as both are less than half the surplus. If player I has an outside option smaller than half the surplus and player II has an outside option larger than half the surplus, then player II receives her outside option while player I receives the rest of the surplus.[1]

Everybody agrees in predicting a fifty-fifty split when outside options are zero (and hence split-the-difference and deal-me-out both predict fifty-fifty). Advocates of split-the-difference argue that player II's share should increase from this base point as we strengthen her bargaining position by increasing her outside option. But as long as her outside option falls short of half the surplus, it is not at all clear that her bargaining position is indeed improved. Given the opportunity, she could threaten to opt out if not offered better than a fifty-fifty deal, but why should player I pay any heed? He is already offering her more than she will get if she opts out. Nor is it clear how she can bring pressure to bear on player I to secure more than her outside option when the latter exceeds half the surplus.

In the experiments of Binmore *et al* [6, 13], deal-me-out performs well when compared with split-the-difference as a predictor of player II's share of the money available at the time a deal is struck. Intriguingly, however, deal-me-out's prediction that player II will never opt out is often wrong. In fact, she frequently opts out when her outside option is sufficiently high. Hard bargaining over how the surplus from an agreement is to be distributed therefore leads to potential gains from trade remaining unexploited.

Although Rubinstein bargaining games are reasonably realistic, their complicated structure makes it difficult to explore the tension between distribution and efficiency revealed by the experimental results. This paper therefore introduces a static bargaining game in which only player II has a positive outside option. We call this game the Outside Option Game. An extensive set of experiments is reported in Section 5 which confirms that deal-me-out continues to predict player II's share in this game very much better than split-the-difference, but that player II still opts out frequently

[1] The Rubinstein theory predicts split-the-difference only if breakdown is *involuntary* (Binmore *et al* [6]). It predicts the commonly-used variant in which the symmetric Nash bargaining solution is replaced by an asymmetric version but the *status quo* remains at the pair of outside option payoffs only if breakdown is involuntary and the probability of a breakdown varies with the identity of the most recent proposer.

when her outside option is sufficiently high. We therefore regard the Outside Option Game as an ideal vehicle for a study of how the pressures in favor of learning to strike a hard bargain can win out over pressures that call for restraint lest the available surplus be lost altogether.

The Outside Option Game. The bargaining game we consider involves two players who can keep a ten dollar bill if they can agree on how to divide it. In the classic Nash Demand Game (Nash [34]), players I and II simultaneously announce take-it-or-leave-it demands, x and y. If $x+y \leq 10$, both players receive their demands. Otherwise each gets nothing.

The Outside Option Game modifies the rules of Nash's game in two ways. We first add an opt-out strategy "OO" to player II's list of strategies. If she opts out instead of making a demand, she receives a payoff of α ($0 < \alpha < 10$) and player I gets nothing, whatever he may have demanded. Second, if player II opts in and $x + y \leq 10$, then each player gets half the unclaimed surplus on top of his or her claim. Thus player I gets $x + \frac{1}{2}(10 - x - y)$ and player II gets $y + \frac{1}{2}(10 - x - y)$. Inefficiencies can then only occur if player II opts out or the two players make incompatible claims. The game is intended to model a negotiation over the division of the profits from a partnership worth ten dollars, where player II (only) incurs an opportunity cost of α dollars in joining the partnership.

Rubinstein bargaining games usually have many Nash equilibria, but just one subgame-perfect equilibrium. Since the Outside Option Game has no proper subgames, all its many Nash equilibria are subgame-perfect. We ignore the mixed-strategy equilibria (which are all inefficient) and sort the pure-strategy Nash equilibria into two classes:

> **Efficient equilibria:** Any pair $(10 - y, y)$ with $y \geq \alpha$ is an efficient Nash equilibrium.
>
> **Inefficient equilibria:** Any pair $(10 - y, \text{OO})$ with $y \leq \alpha$ is an inefficient Nash equilibrium.

Split-the-difference selects the Nash equilibrium in which player II opts in and receives $\alpha + \frac{1}{2}(10 - \alpha)$ and player I receives $\frac{1}{2}(10 - \alpha)$. Deal-me-out again calls for player II to opt in, but she gets only α when $\alpha > 5$, leaving $10 - \alpha$ for player I. When $\alpha \leq 5$, the surplus is divided fifty-fifty so that each player receives a payoff of 5.

Our basic questions are: *What bargaining convention will be used to split the surplus and compensate player II for the opportunity cost she incurs in* *joining the partnership? Will attempts to drive a hard bargain sometimes lead to inefficient outcomes in which player II opts out?*

Falling off the edge. To be consistent with the experimental results, a game-theoretic analysis should answer *deal-me-out* to the first of the previous two questions and *yes* to the second. However, game theory commonly joins the Coasian tradition in answering *no* to the second question. Cooperative game theories as well as noncooperative theories based on payoff dominance simply assume that equilibria must be efficient. When efficiency is not taken for granted, an argument along the following lines is often proposed. Player I reasons that player II would do better to opt out than to make a claim that results in her getting less than α. He therefore argues that he cannot hope for more than $10 - \alpha$ and so makes a claim for this amount or less, thereby making it safe for player II to opt in with a claim of α.[2]

When such "forward induction" arguments are proposed, it is usually a maintained hypothesis that there is common knowledge that everybody is perfectly rational. But the arguments become less plausible in more realistic worlds that retain even quite low levels of irrational behavior. In such cases, we shall argue that the logic of forward induction may fail to be learned at all.

To see how a period of trial-and-error adjustment may to an inefficient outcome in the Outside Option Game, consider a case in which players initially operate the split-the-difference norm. At the outset, it will therefore be optimal for player II to opt in. However, one must anticipate that there will be occasional coordination failures when other strategies are played by agents who are confused or who are not party to the implicit agreement to abide by the split-the-difference norm. Such perturbations would have only an ephemeral effect at a Nash equilibrium that is deep inside its basin of attraction (relative to the adjustment dynamics), and so is isolated from other equilibria of the game. But Nash equilibria in the Outside Option Game are packed close together. As a result, even small perturbations may be enough to shift the system from an equilibrium to one of its neighbors.

[2] Notice, however, that this result is not produced by common forward induction criteria such as the iterated elimination of weakly dominated strategies. Nöldeke and Samuelson [35] present an evolutionary model that would eliminate equilibria in which the outside option is chosen in this game and in the Dalek Game considered below. In the terms of Binmore *et al* [11], Nöldeke and Samuelson's is an ultralong-run theory, while we are working here with long-run theories that we believe more relevant to experimental data.

We follow the biological literature in referring to perturbations that have this effect as drift. In the Outside Option Game, we argue that drift is likely to have a centralizing tendency, meaning that it pushes players toward the fifty-fifty outcome. At this point, we simply explore the impact that such centralizing drift is likely to have on equilibrium selection in the Outside Option Game. Section 4 explains how and why perturbed adjustment in the Outside Option Game can generate such drift.

Split-the-difference lies inside the set of efficient equilibria of the Outside Option Game. Close enough to this set of equilibria, the adjustment pressures pushing the players towards income-maximizing behavior become weak (because their behavior is already close to optimal). Drift then dominates. If it has a centralizing tendency, it will move the system toward the fifty-fifty outcome, where it will be stabilized if $\alpha < 5$. But the fifty-fifty outcome is not an equilibrium when $\alpha > 5$. In this case, the drift will move the system to the endpoint $(10 - \alpha, \alpha)$ of the cluster of efficient equilibria. Player *II* then receives only the minimal rational compensation needed to persuade her to opt in. This endpoint is therefore less stable than interior points of the cluster. Once it has been reached, there is a risk that any further perturbations will lead to the system "falling off the edge" and being carried away by renewed selection pressure to the cluster of inefficient equilibria in which player *II* opts out.

3 Other Experiments and Alternative Explanations

The experimental results of Binmore *et al* [6, 13] on Rubinstein bargaining models with outside options need to be compared with the experiments of Hoffman and Spitzer [27, 28] and Harrison and McKee [25], who studied the behavior of subjects in free-form, face-to-face bargaining sessions. In these latter experiments, one of the subjects was designated as the controller. The controller could choose either to receive α dollars (leaving the other subject with nothing) or to split a larger sum of money with the other subject. The two subjects discussed which choice the controller should make and then signed a binding agreement specifying how the total payment should be split between them. In contrast to the results of Binmore *et al* [6, 13], the deals reached under such circumstances were commonly efficient. In particular, the controller rarely exercised his or her capacity to opt out.

We have doubts about the extent to which the results from such face-to-face bargaining experiments involving relatively small amounts of money are likely to generalize to a wider economic context, because the intimacy generated while the subjects fraternize is likely to inhibit the hard bargaining that we are interested in studying. Since we do not think we can offer experimental subjects enough money to overcome this intimacy in face-to-face encounters, we have them communicate through a computer so that both parties to the negotiation remain anonymous throughout.[3] Our doubts about hard bargaining in face-to-face experiments are seconded by the willingness of the controller in some of Hoffman and Spitzer's treatments to agree to deals close to a fifty-fifty split, even though he or she could have obtained several dollars more by foregoing an agreement altogether and opting out.[4] Such behavior is virtually absent in the experiment reported in this paper.

Results like those of Hoffman and Spitzer [27, 28] have led many authors to argue that income-maximizing models of human behavior have little or no predictive power in bargaining situations. Sometimes the optimizing paradigm is not abandoned, but the subjects' behavior is rationalized by attributing complex motivations to the subjects which may include a "taste for fairness" (Bolton [14]). However, it is common to reject the optimizing explanation altogether and argue that subjects simply operate whatever fairness norm happens to be triggered by the manner in which the experiment is framed.

We do not deny that people are often ruled by social norms in their day-to-day lives, nor that subjects bring such social norms into the laboratory with them. However, we think it is a mistake to proceed as though social norms are fixed and immutable. Real-world social norms have presumably evolved to coordinate human behavior on equilibria in the game of life. But when one of these norms is triggered in the laboratory, it is unlikely that the behavior it engenders in the subjects will be adapted to the artificial game invented by the experimenter. The subjects' initial behavior may then be hard to reconcile with income-maximization. But ample evidence exists that new social norms can sometimes evolve during the experiment that eventually succeed in coordinating the behavior of the subjects on an equilibrium of the laboratory game.

[3]The suspicion that anonymity matters in such circumstances has been confirmed directly in a number of studies, notably that of Hoffman *et al* [29].

[4]Ochs and Roth [36] report that their subjects also often make disadvantageous offers when anonymously playing alternating-offers bargaining games.

If the game resembles situations the subjects have had the opportunity to learn about outside the laboratory, they may only have to fine-tune the social norm triggered at the outset of the experiment. In other cases, they may have to evolve an entirely new social norm from scratch. We suspect the Outside Option Game represents a case nearer the first of these two extremes. However, we need not commit ourselves to a view on this issue, since our learning story applies both to learning within the laboratory and to learning outside the laboratory.

The evidence that experimental subjects change their behavior over time is overwhelming (Andreoni and Miller [2], Binmore *et al* [4, 6, 12, 13], Crawford [19, 20], Miller and Andreoni [32], and Roth and Erev [40]). The change of behavior over time is especially dramatic in Binmore *et al* [12]. In a two-stage bargaining game, totally inexperienced subjects played much like subjects are reported to behave in the much-replicated experiments on the Ultimatum Game (Güth, Schmittberger and Schwarze [23], Güth and Tietze [24], Roth [39]).[5] In particular, the modal opening offer was fifty-fifty. But after experiencing just one game as player II, subjects occupying the role of player I mostly switched to the backward induction offer. Neither an explanation that attributes exotic utility functions to the players nor an explanation in terms of fixed fairness norms readily accommodates such data. It is, however, consistent with models that treat the subjects as simple but imperfect income-maximizers who need time to learn how best to play a game.

Binmore *et al* [4] offers some insight on the evolution of fairness norms in the laboratory. In a smoothed version of the Nash Demand Game, it proved easy to condition the subjects to coordinate on a variety of Pareto-efficient focal points by having them play against suitably programmed robots. But when the subjects began to play against each other, their behavior adapted until each group established a tight consensus on an exact Nash equilibrium of the game. Different groups converged on different equilibria, but when asked what was fair in a computerized debriefing, there was a strong tendency to identify whatever consensus was achieved in their own group with the fair outcome of the game. Similar results are reported in Binmore *et al* [13] and the current paper.

[5] In the Ultimatum Game, player I makes on offer to player II for the division of the sum of money. The money is split in the manner specified if she accepts. If she refuses, both players get nothing.

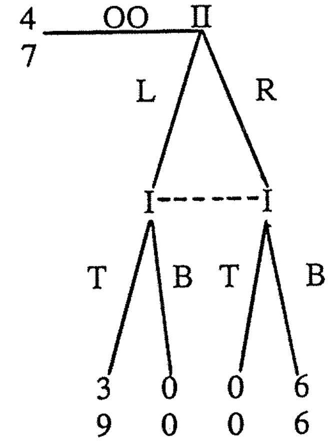

			II	
		OO	*L*	*R*
I	*T*	4,7	3,9	0,0
	B	4,7	0,0	6,6

Figure 1: The Dalek Game

4 Falling Off The Edge

This section examines the role of drift in equilibrium selection. We first consider a stripped-down version of the Outside Option Game. Figure 1 shows a game whose basic form we borrow from Kohlberg and Mertens [30] and which is sometimes called the Dalek Game.

The Dalek Game has two subgame-perfect equilibrium outcomes: one in which players choose (T, L) and so obtain the payoff pair $(3, 9)$, and one in which player II takes her outside option (strategy OO) because player I plays B with probability at least 2/9. Only the first of these two possibilities satisfies the forward induction criteria that are usually defended in this context by an appeal to the iterated elimination of weakly dominated strategies.[6] In particular, R is strictly dominated by OO for player II. She

[6]The iterated elimination of weakly dominated strategies is identified by Kohlberg and Mertens [30] as one of the basic desiderata that an equilibrium concept should satisfy. Dekel and Fudenberg [21, p.245] argue that the iterated elimination of weakly dominated strategies "clearly incorporates certain intuitive rationality postulates", while Nalebuff and Dixit [33, p.86] offer it as one of their four basic rules for playing games. In the Dalek game, the equilibrium (T, L) is also selected by the never-weak-best-response criterion (Kohlberg and Mertens [30]) and van Damme's [45] forward induction criterion

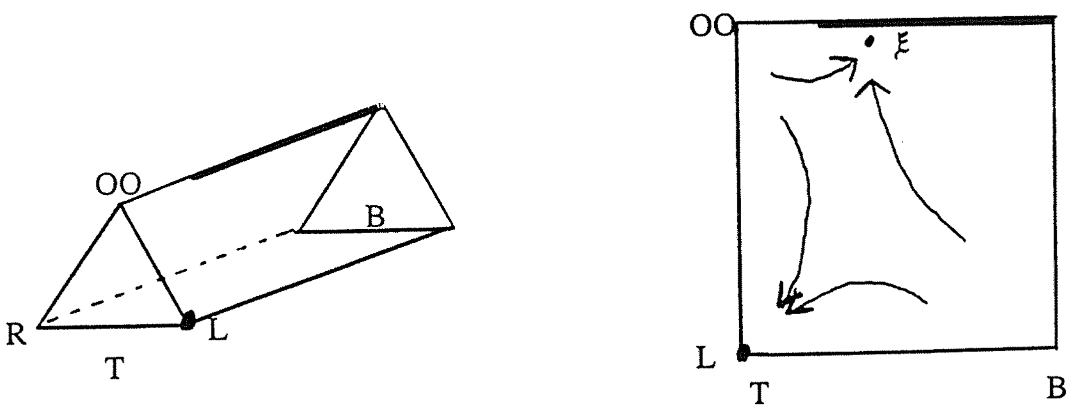

2a 2b

Figure 2: Adaptation in the Dalek Game

will then never opt in to the game and play R because OO gives a higher payoff. Player I, realizing this, eliminates R from consideration and then finds that T weakly dominates B. Once B is accordingly eliminated from consideration, L dominates OO for player II, causing her to opt in. The result is the equilibrium (T, L).

This argument is appealing, but experimental studies of the Dalek Game by Balkenborg [3] show that player II virtually always opts out. Cooper *et al* [18] find similar results in related games. How can such data be reconciled with a theory that treats players as income-maximizers?

Our analysis of the Dalek Game begins by supposing that the game is played repeatedly by pairs of agents chosen at random from two infinite populations. We take the fraction of agents in population I using strategy i at time t to be x_i $(i = T, B)$. The proportion of agents in population II using strategy j at time t is taken to be y_j $(j = OO, L, R)$. The expected payoff to an agent from population I using strategy i who is chosen to play at time t is f_i. The average payoff to agents in population I at time t is then $\overline{f} = x_T f_T + x_B f_B$. The expected payoffs g_j and $\overline{g}$ for population II are defined similarly.

The agents in our two populations learn how to play the game over time. Our model of learning is taken from Binmore, Gale and Samuelson [5], who show that an aspiration and imitation-based learning process leads to the following form of the replicator equations,

$$\begin{aligned} \dot{x}_i &= x_i(f_i - \overline{f})/\Delta \qquad (i = T, B) \\ \dot{y}_j &= y_j(g_j - \overline{g})/\Delta \qquad (j = OO, L, R). \end{aligned}$$

where Δ is a constant that reflects how sensitive agents' strategy choices are to their payoffs and that could be eliminated by rescaling our measure of time. Observe that the rate at which the fraction of a population increases depends on how large the fraction is already and how much better it is currently doing than the average. Other papers examining learning models that lead to some version of the replicator dynamics include Binmore and Samuelson [9], Börgers and Sarin [15], Cabrales [16], Schlag [43], and Weibull [46]. While we must specify some learning model in order to do the simulations reported below, and the replicator dynamic is convenient, it is important to note that the qualititative nature of our results would continue to hold as long as learning causes strategies to respond smoothly to payoffs.

Our learning model is unlikely to be so perfect as to have captured all of the forces shaping agents' strategy decisions, though we hope to have captured the important ones. Omitted considerations appear in the form of "drift", which we capture by examining the following perturbed version of the replicator dynamics:

$$\begin{aligned} \dot{x}_i &= (1-\delta_1)x_i(f_i - \overline{f})/\Delta + \delta_1(\tfrac{1}{2} - x_i) \qquad (i = T, B) &\quad (1) \\ \dot{y}_j &= (1-\delta_2)y_j(g_j - \overline{g})/\Delta + \delta_2(\tfrac{1}{3} - y_j) \qquad (j = OO, L, R). &\quad (2) \end{aligned}$$

The parameters δ_1 and δ_2 measure how much drift there is in the two populations. As in Binmore *et al* [5], we model drift as switching an agent from one strategy to another independently of the selection pressures, where such a switch leads to each possible strategy being employed with equal probability. We are interested in cases in which δ_1 and δ_2 are small because large influences on strategy choices have presumably been incorporated in the basic learning model.

Figure 2a shows the state space for the perturbed replicator dynamics in the Dalek Game. The strategy R is strictly dominated, and will be eliminated by the replicator dynamics (Samuelson and Zhang [42]).[7] Figure 2b accordingly shows the bottom face of the state space, where strategy

[7] A small proportion of R will survive in the *perturbed* dynamics because drift continually introduces B into the population.

R is never used, along with a phase diagram for the perturbed replicator dynamics. The details of this phase diagram will depend upon the precise specification of drift.[8] The vertex (T, L) corresponds to the population state in which all members of both populations play the forward induction solution. There is an asymptotically stable state nearby that approaches (T, L) as δ_1 and δ_2 approach zero (for any specification of drift). However, we focus attention on a second asymptotically stable state ξ that Binmore and Samuelson [10] show can arise close to the component of Nash equilibria on the edge of the prism that corresponds to player II's choosing her outside option OO.

Efficiency is therefore not guaranteed. Nor need the final outcome respect the deletion of weakly dominated strategies. Nor can equilibria in mixed strategies safely be neglected. Where the system ends up depends on the basin of attraction in which it begins, and the basin of attraction for ξ is substantial.

What drives these results? Unless the process is too noisy, we expect the adaptive learning process to converge on whatever equilibrium lies at the heart of the basin of attraction in which the system begins.[9] At the outset, adjustment will be quick as the more extreme types of irrationality are eliminated. But as an equilibrium is approached, convergence will slow down as the driving force behind the dynamics, payoff differences between strategies, gets small. When rival equilibria are far apart, such slowing down in the rate of convergence creates no problem for equilibrium selection. In the Dalek game however, the component of equilibria in which player II plays OO consists of a large number of equilibria packed arbitrarily close together. Near this component, the learning forces modeled by the replicator dynamics become arbitrarily weak and the motion of the system is almost entirely driven by drift. It then becomes crucial whether drift pushes the system toward or away from this component. If drift pushes the system toward the component by continually introducing strategy B as well as T for player I, then an asymptotically stable state near the component can appear, as in Figure 2. This asymptotically stable point remains no matter how small are the drift parameters δ_1 and δ_2.

In the Dalek Game, there is no counterpart of the deal-me-out solution and so no way to examine the role of drift in determining whether split-

[8]Binmore and Samuelson [10] explore this dependence.

[9]This obviously ignores a host of nonconvergence problems. Fortunately, we work with such simple games in this paper as to not have to worry about such problems.

the-difference or deal-me-out is selected, as well as no way to study the relationship between this choice and whether player II opts out.[10] We accordingly turn to the Outside Option Game.

The Outside Option Game. The key feature of drift in the Outside Option Game is its potential to have a centralizing tendency, much as the key feature of drift in the Dalek game is whether it pushes the system toward or away from the component of equilibria in which player II plays OO. In this section, we explore the forces behind this centralizing tendency. To more effectively study these forces, we consider the Outside Option Game *without* the outside option.

The sum of money to be divided in our experimental version of the Outside Option Game is \$10.00, with claims made in increments of a dime. Hence, the set of possible claims for both players (in the absence of an outside option) is $X = \{0, 0.1, 0.2, \ldots, 10\}$. The perturbed replicator dynamics are given by:

$$\dot{x}_i = (1-\delta)x_i(f_i - \bar{f})/\Delta + \delta(\tfrac{1}{101} - x_i) \qquad (i \in X) \tag{3}$$
$$\dot{y}_j = (1-\delta)y_j(g_j - \bar{g})/\Delta + \delta(\tfrac{1}{101} - y_j) \qquad (j \in X). \tag{4}$$

Since nothing hangs on asymmetries in the drift levels, we take $\delta_1 = \delta_2 = \delta > 0$.

We are unable to solve this system analytically. Instead, we compute numerical solutions. Figure 3 reports some of the calculations.[11]

The first column describes the initial condition. In a uniform initial condition, 1/101 of the players in each population initially make each of the possible claims. In a "spike" initial condition, described with a pair of numbers written as x/y, 95 percent of population I players initially make claim x while 95 percent of population II players initially make claim y. If $x + y = 10$, these claims exactly exhaust the surplus, while claims with $x + y > 10$ are incompatible. The remaining five percent in each population is equally distributed among the remaining claims. The second column in the table reports the value of the noise parameter δ.[12]

[10] The outside option in the Dalek Game might also be interpreted as having fairness virtues that are absent in the Outside Option Game.

[11] Proulx [38] contains a more extensive report of the numerical calculations.

[12] Strategies and payoffs in the numerical calculations where measured in dimes, and hence ranged from 0 to 100. We set $\Delta = 100$. The value of Δ again affects nothing other than the units in which we measure time and rates of drift.

Initial Condition	δ	x_f $(f(x_f))$	y_g $(g(y_g))$	$F(10-y_g)$	$G(10-x_f)$
1 Uniform	.0001	5.0 (.99900)	5.0 (.99900)	.99990	.99990
1 Uniform	.010	5.0 (.89889)	5.0 (.89889)	.98981	.98981
3 Uniform	.023	5.0 (.75620)	5.0 (.75620)	.97514	.97514
4 Uniform	.058	5.0 (.22772)	5.0 (.22772)	.91638	.91638
5 Uniform	.160	5.0 (.01664)	5.0 (.01644)	.77378	.77378
6 8.0/8.0	.010	5.0 (.89889)	5.0 (.89889)	.98981	.98981
7 2.0/2.0	.010	5.0 (.89889)	5.0 (.89889)	.98981	.98981
8 5.0/8.0	.010	5.0 (.89889)	5.0 (.89889)	.98981	.98981
9 4.0/8.0	.010	4.0 (.89918)	6.0 (.89718)	.98478	.99317
10 7.5/9.5	.010	7.5 (.89131)	2.5 (.89240)	.99653	.96977
11 9.0/1.0	.010	9.0 (.86406)	1.0 (.86406)	.99877	.91093
12 9.0/1.0	.015	9.0 (.74107)	1.0 (.78061)	.99803	.86632
13 9.0/1.0	.023	7.9 (.68101)	2.1 (.74175)	.99287	.90894
14 9.0/1.0	.054	5.5 (.24444)	4.5 (.35281)	.93995	.91020
15 5.5/4.5	.015	5.5 (.84498)	4.5 (.84751)	.98724	.98109
16 5.5/4.5	.023	5.5 (.75201)	4.5 (.75893)	.97953	.96977
17 5.5/4.5	.054	5.5 (.24444)	4.5 (.35281)	.93995	.91020
18 5.5/4.5	.063	4.9 (.13192)	4.9 (.13192)	.90746	.90746

Figure 3: Numerical Calculations

The next two columns present the modal claim when the system has reached equilibrium, with x_f being the modal claim in population I and y_g the modal claim in population II. The modal claims are followed by numbers $f(x_f)$ and $g(x_g)$ (in parentheses) indicating the proportion of the population playing this claim. The final columns report the fraction $F(10-y_g)$ of population I making a claim less than the surplus left by the population II modal claim, and the fraction $G(10-x_f)$ of population II making claims less than the surplus left by the population I modal claim. These provide some guide as to how much disagreement survives in equilibrium, with higher numbers indicating less disagreement.

The first five lines of Figure 3 show that the smaller is δ, the more concentrated is the distribution of claims in each population around its mode. As the level of drift increases claims tend to become more dispersed, though this dispersion leads primarily to players claiming less instead of more than the surplus left by the opponent's modal claim. In each of these cases, the modal claim is the fifty-fifty outcome. This is a product of the symmetry of the uniform initial condition. Lines 6–8 show that other initial conditions can also give the fifty-fifty outcome, including claims that are initially incompatible, claims that initially leave a large surplus, and asymmetric claims. Lines 9–10 show, however, that not all initial conditions lead to the fifty-fifty outcome.

Lines 11–18 of Figure 3 reveal the centralizing force of drift. In each of these lines, we start with an initial condition in which the agents have very nearly coordinated on an asymmetric division of the cake. A division that provides 90% of the surplus to player I survives for a drift level of .015 but not for drift levels of .023 and higher. In contrast, if the initial division provides only 55% of the surplus to player I, then these modal claims survive for drift levels as high as .054.

Drifting toward the center. Why does drift yield a centralizing force? Let $f(x,\delta)$ be a density describing the proportions of population I making the various claims in the strategy set X. Let $F(x,\delta)$ be the cumulative distribution, so that $F(x,\delta) = \sum_{x' \leq x} f(x',\delta)$. Let $g(y,\delta)$ similarly be a density describing population II's claims and let $G(y,\delta) = \sum_{y' \leq y} g(y,\delta)$ be the corresponding cumulative distribution. We will be interested in cases in which $f(x,\delta)$ and $g(y,\delta)$ describe stationary states of (3)–(4).

Let x_f and y_g be the modes of f and g. From (3), it is clear that, in a stationary state, the mode x_f is the player I claim that has the highest expected payoff. Because there is drift, $f(x_f) < 1$. Similarly, y_g is the payoff-maximizing claim for player II and $g(y_g) < 1$. Just as the game has many Nash equilibria, the system (3)–(4) has many stationary states when drift is small. For any claims x_f and y_g that exactly exhaust the surplus and give each player positive surplus, there is a stationary state of (3)–(4) with modes x_f and y_g if drift is sufficiently small:[13]

Proposition 1 *Let $0 < x_f < 10$ and $x_f + y_g = 10$. Then for sufficiently small δ there exists an asymptotically stable stationary state of (3)–(4) with modes x_f and y_g.*

[13]This contrasts with the Dalek game, where the effects of drift appear for arbitrarily small values of δ. In the Dalek game, the key effect of drift is to move the system along a continuum of weak Nash equilibria. In the Outside Option Game, we are dealing with a large number of strict Nash equilibria. These equilibria lie close together in the state space, so that relatively small levels of drift suffice to move the system between Nash equilibria. However, any pure strategy Nash equilibria in which each player gets positive surplus will survive if δ is made sufficiently small.

Proof: Let $0 < x_f < 10$ and $x_f + y_g = 10$. Then it is a strict Nash equilibrium for players I and II to make claims x_f and y_g. There is thus an asymptotically stable state of the unperturbed replicator dynamics, denoted by x^*, in which x_f and y_g are played with unitary probability. But then for sufficiently small δ, the replicator dynamics with drift have an asymptotically stable state arbitrarily close to z^* (Hirsch and Smale [26, Theorems 1–2, p. 305]). □

We do not expect drift to always be arbitrarily small. If not, some equilibria will no longer be stationary states. To determine which equilibria are eliminated, we require some more information about the nature of stationary states. Our approach here is to observe that the stationary states in our numerical calculations satisfy three properties. Let $(f_1(x,\delta_1), g(y_1,\delta_1))$ and $(f_2(x,\delta_2), g_2(y,\delta_2))$ be stationary states of (3)–(4). Then:

Property 1 *Let $x_{f_1} = 10 - y_{g_1}$. Then $f_1(x,\delta_1)$ and $g_1(y,\delta_1)$ are quasiconcave. In addition, $(\hat{f}(x,\delta), \hat{g}(y,\delta))$ is a stationary state, where $\hat{f}(x,\delta) = f(10-x,\delta)$ and $\hat{g}(y,\delta) = g(10-y,\delta)$.*

Property 2 *Let $x_{f_1} = x_{f_2} \equiv x_f = 10 - y_{g_1} = 10 - y_{g_2} \equiv 10 - y_g$ and $\delta_1 > \delta_2$. If $x_f > 5$, then*

$$\frac{g_1(y_g + 0.1, \delta_1)}{G_1(y_g, \delta_1)} > \frac{g_2(y_g + 0.1, \delta_2)}{G_2(y_g, \delta_2)}. \tag{5}$$

A symmetric condition holds for f_1 and f_2 when $x_f < 5$.

Property 3 *Let $\delta_1 = \delta_2 \equiv \delta$ and $x_{f_1} = 10 - y_{g_1} > x_{f_2} = 10 - y_{g_2} > 5$. Then:*

$$\frac{g_1(y_{g_1} + 0.1, \delta)}{G_1(y_{g_1}, \delta)} > \frac{g_2(y_{g_2} + 0.1, \delta)}{G_2(y_{g_2}, \delta)}. \tag{6}$$

A similar condition holds for f_1 and f_2 when $x_{f_1} = 10 - y_{g1} < x_{f_2} = 10 - y_{g_2} < 5$.

Property 1 indicates that the distributions $f(x,\delta)$ and $g(y,\delta)$ increase monotonically as x and y approach the modes of f and g. This is simply the statement that expected payoffs increase as one approaches the expected-payoff maximizing strategy. Property 1 also notes that the game is symmetric. Properties 2–3 address a particular measure of the dispersion of the densities f and g. This measure concerns the agent receiving the smaller share of the surplus, and is the ratio of the probability attached to a claim just higher than the modal claim to the probability attached to claims less than or equal to the modal claim. Properties 2–3 state that this measure is higher when drift is higher and when the mode is further from 5.

A necessary condition for equilibrium is that the player with the larger share receive a higher payoff from the modal claim than from the next lower claim. When $x_f > 5$, we thus need $\pi_I(x_f, g) > \pi_I(x_f - .01, g)$, where $\pi_I(x, g)$ is player I's payoff from claim x given density g, or:

$$\begin{aligned}
\pi_I(x_f, g) &- \pi_I(x_f - .01, g) \\
= \sum_{y=0}^{10-x_f} \left(x_f + \frac{10 - y - x_f}{2}\right) &- \sum_{y=0}^{10-x_f+0.1} \left(x_f - 0.1 + \frac{10 - y - x_f + 0.1}{2}\right) \\
= \tfrac{1}{2}G(10 - x_f, \delta) &- (x_f - 0.1)g(x_f - 0.1, \delta) > 0. \qquad (7)
\end{aligned}$$

Comparing (7) with Properties (2) and (3), we immediately obtain:

Proposition 2 *Let $(f_i(x, \delta), g(y_i, \delta))_{i \in I}$ be a collection of strictly positive densities satisfying Properties 1–3.*

(2.1) Fix $\delta_i = \delta$ for all i. If there exists an i with $x_{f_i} = 10 - y_{g_i} > 5$, then then we can add a pair $(f_j(x, \delta), g_j(y, \delta))$ that preserves Properties 1–3 such that $x_{f_j} = 10 - y_{g_j}$ for any $x_{f_j} \in [10 - x_{f_i}, x_{f_i}]$, but may be unable to do so for $x_{f_j} \notin [10 - x_{f_i}, x_{f_i}]$.

(2.2) If there exists an i with $x_{f_i} = 10 - y_{g_i}$, then then we can add a pair $(f_j(x, \delta_j), g_j(y, \delta_j))$ that preserves Properties 1–3 such that $x_{f_i} = x_{f_j} = 10 - y_{g_j}$ for any $0 \leq \delta_j < \delta_i$, but may be unable to do so for $\delta_j > \delta_i$.

This proposition reveals the centralizing tendency of drift. The necessary condition (7) for a stationary state with exactly compatible modal claims, given Properties 2–3, can be satisfied for an interval of modal claims that is centered around the fifty-fifty outcome, and which shrinks toward the fifty-fifty outcome as the drift level increases.

To see the forces behind this result, notice that abandoning the modal claim to make the next lowest claim reduces the surplus from agreements when the opponent makes her modal claim or less, but achieves some new agreements (arising from cases in which the opponent makes a larger claim). As the drift rate increases, the relative probability of such a larger claim increases and hence it is less likely that (7) is satisfied, so that higher drift levels yield a stronger centralizing force. More importantly, the new agreements that can be realized by reducing a player's claim are most lucrative when the player is already receiving a large share of the cake. Hence, condition (7) is most likely to fail for extreme divisions of the surplus and most likely to hold for more equitable divisions. The collection of possible exactly-compatible modal claims will thus be centered around the fifty-fifty outcome, and higher drift levels will give a smaller collection of such claims. This is the centralizing force created by drift.

Why are we interested in Proposition 2 when (7) is only a necessary condition for a stationary state, showing that a single alternative claim is not a better reply than the modal claim? We have reason to believe that the particular alternative considered in (7) is the most likely superior reply to the modal claim. In particular, we expect a better reply to be a lower rather than a higher claim, since a higher claim sacrifices any chance at an agreement when the opponent makes her modal claim. Among lower claims, we expect the claim just below the modal claim to be the most profitable, since (given the quasiconcavity of f and g noted in Property 1) it is here that the largest increase in the probability of an agreement is achieved. In addition, this next lower claim is most likely to be a better reply for the player receiving the largest share of the surplus, since this is the player for whom the newly achieved agreements are most valuable. This is precisely the claim addressed by (7).[14]

Finally, what does this have to do with the Outside Option Game? Suppose $\alpha > 5$. An efficient outcome requires an agreement giving player II a payoff larger than α. Suppose this is produced a stationary state of the learning dynamics in which $x_g = 10 - x_f > \alpha$. If the drift level is sufficiently small, this is no difficulty. For somewhat higher drift levels, however, such a stationary state does not exist. Instead, player II finds it more profitable to reduce her claim in order to secure more agreements with the noisy player Is who are claiming more than $10 - x_g$. Because x_g is relatively large, it is worth claiming somewhat less in order to secure these agreements. A stationary state thus requires modal claims closer to $10 - \alpha$ (for player I) and α (for player II). If there is enough drift, a stationary state in which player II does not opt out will require a modal claim for player II less than α and will give player II a payoff less than α. But then player II will opt out, and an inefficient outcome appears.

[14] None of our numerical calculations produced a counterexample to this belief.

5 Experimental Results

This section examines the extent to which the theory of the previous section matches the outcomes of experiments in the Outside Option Game.

Experimental design. The experiment was conducted at the Michigan Economics Laboratory with undergraduates of the University of Michigan. Each experimental session involved 12 subjects who sat at networked microcomputers that were screened from each other. The subjects were asked to read the written instructions (reproduced in the Appendix) and given an interactive demonstration of how claims were registered, payoffs determined, and so forth.

Following the demonstration, subjects participated in a series of bargaining sessions. At the beginning of each session, subjects saw on their video displays the outline in white against a black background of a tall, hollow, rectangular "cake". To the left of the cake, in blue print, the number 10 together with brackets reaching from top to bottom reminded subjects that the total height of the cake represented an amount of money that was always nominally worth 10 dollars. Almost as wide and slightly inside the rectangular cake was a second, smaller, hollow rectangle which began at the bottom of the cake and whose height represented the amount of money that player *II* could obtain unilaterally by opting out. The numerical value of the opt-out payment was also indicated. To avoid suggesting focal points, only these two numbers were indicated on the display.

Player *I* indicated his claim by moving a small red cursor that pointed to the right side of his cake up or down using the computer's up and down arrow keys. As the red cursor moved down from its initial position at the top of the cake, the area of the rectangle between the top and the cursor filled with red to represent the amount of player *I*'s claim. When player *I* was satisfied, he registered his claim by pressing ENTER. At that point the numerical value of his claim was indicated and he had the opportunity to revise the claim by pressing the function key F10 or confirm the claim by pressing SPACE BAR.

The procedure by which player *II* indicated a claim was similar except that player *II* moved a small green cursor that was initially positioned at the bottom right hand side of his cake. As the green cursor moved up or down, the area between the bottom of the cake and the cursor filled with green to indicate the amount of player *II*'s claim. Player *II* could also indicate a decision to opt out by pressing the BACKSPACE key, at which time the area of the rectangle indicating the opt-out payment filled with white to indicate the amount that player *II* could gain by opting out. As with player *I*, player *II* was given the opportunity to revise his choice by pressing F10 or confirm it by pressing SPACE BAR.

After both player *I* and player *II* confirmed their choices, the choice of each player's counterpart for the session was displayed by overlying the appropriate red, green, or white region on the player's own display. If the red and green claims of the two players overlapped, then the total claimed by both players was more than 10 dollars, and neither player received anything. The area of overlap was shown in yellow. If the red and green claims did not overlap, then each player received his claim together with half the unclaimed cake. A white line dividing the surplus (i.e., the remaining dark region in the middle of the cake) was displayed together with the numerical value of the player's total payoff. Finally, if player *II* opted out, then he received his opt-out payment while player *I* received nothing. Along with a graphical display of the players' choices and payoffs, a brief written summary of the outcome was displayed. For example, if player *I* claimed 2 dollars and player *II* claimed 4 dollars, so that a surplus of 4 dollars remained, then the following message was shown:

Player II has opted in and the claims can be met. Player I gets \$4.00. *Player II gets* \$6.00. *Either player could have gotten more by claiming more.*

Subjects did not know with whom they had been paired in each session and communicated anonymously through the computer as described above. After each session, subjects were paired with a new partner who was chosen randomly subject to constraints discussed at the end of this section. Whether a player was player *I* or player *II* in a given session was also determined randomly subject to the constraint that no subject was the same type of player for more than two sessions in a row.[15]

Players participated in twenty "practice" sessions followed by ten "real" sessions. The cake in each session was always nominally worth ten dollars. However, subjects were paid the amounts they succeeded in obtaining in only two of each set of ten sessions. Moreover, for the first ten practice sessions, the subjects were paid at the rate of one *dime* for each dollar they earned. In the second set of practice sessions, subjects were paid at the

[15]If we had strictly alternated the players' types, then any given subject could have participated in bargaining sessions with only half of the subjects.

rate of one *quarter* for each dollar they earned. Only in the final set of ten real sessions were subjects paid at the full rate of one dollar per dollar earned. After each set of ten sessions, a "roulette wheel" appeared on each subject's screen, and the two sessions for which that subject would be paid were randomly selected.[16]

The opt-out payment which a player *II* could receive varied from session to session. There were two types of experiments. In experiments that received the "up" treatment, the opt-out payments for each set of ten practice or real sessions were in the following ascending order: { 0.90, 0.90, 2.50, 2.50, 4.90, 4.90, 6.40, 6.40, 8.10, 8.10 }. In experiments that received the "down" treatment, the opt-out payments for each set of ten sessions were in the opposite, descending order.

After the bargaining sessions were over, subjects were asked to complete a computerized questionnaire. For each opt-out payment, subjects were asked the question: "What do you feel would be a fair amount for player *II* to get?" by moving a green cursor to indicate a claim on the rectangular cake precisely as in the actual bargaining sessions as player *II*. The opt-out payments in the questionnaire were presented in the same (ascending or descending) order that was used in the bargaining sessions.[17]

Our expectation before undertaking the experiment was that player *II*s would not take the outside option, with the primary question of interest being how much compensation they would receive for not doing so. Foregoing a large outside option is potentially risky, however, especially if there are player *I*s who make such large claims as to not allow player *II* a payoff at least equal to the outside option. When designing the experiment, we attempted to isolate the effect of this risk and attain conditions under which player *II* would not opt out. This motivated our presenting the outside options in both an up and a down treatment. In addition, subjects were "filtered" in the practice bargaining sessions.

Our motivation for the filtering was a suspicion that "irrational" behavior by player *I*s would be correlated with a larger frequency of disagreements and, consequently, lower profits and greater risk to player *II* from not taking the outside option. Hence, after the first ten practice sessions, the four subjects with the lowest total profit in these sessions were "filtered out",

[16]We believe that "intermittent reinforcement" like that which we provided increases the subjects' interest in the experiment. Such effects are widely reported by psychologists.

[17]For each opt-out payment, subjects were also asked to indicate their best guess of the median of the claims that the other subjects in their group designated as fair for a player *II*. The subject whose guess was closest to the actual median was awarded $2.

and, in subsequent practice and real sessions, these subjects were matched only with others in their group. After the second ten practice sessions, the four subjects of the remaining eight who had the lowest cumulative profit in all twenty practice sessions were also grouped. Thus, at the start of the real bargaining sessions there were three groups of four subjects who had been selected by their profits in the practice sessions and who bargained in real sessions only with subjects in their own group. Subjects were not informed of this filtering procedure.

Somewhat to our surprise, the behavior of the average subject in each group during the real sessions and the average responses to the questionnaire did not differ much from group to group. For example, pooled over all subjects who experienced the same treatment, the differences between the frequencies with which subjects in different groups opted out were always within 0.11 for opt-out payments greater than 5.00 dollars. For each opt-out payment, the median player-I claims of each group never differed by more than \$.40, and the medians of the claims indicated as fair for player II in the questionnaire were identical for all three groups. As a result, Figures 4–5 summarize the data pooled across all three groups.

Experimental Results. Figures 4 and 5 summarize the results from the real bargaining sessions of the experiments and the questionnaire. The data are reported separately for each opt-out payment and each treatment. There were a total of 9 experiments where subjects were presented with the opt-out payments in ascending order (the up treatment) and 19 experiments where the opt-out payments were presented in descending order (the down treatment). There were 12 subjects in each experiment and each opt-out payment was presented for two real bargaining sessions. Since half the subjects were player I and half were player II in each session, for every opt-out payment there were a total of 108 choices by each type of player in real sessions with the up treatment and 228 choices by each type of player for the down treatment.[18]

We summarize the results as follows:

> **Division of surplus:** Deal-me-out is in many respects a good predictor of subjects' behavior. As player I, the median subject made claims that were only slightly less than those predicted by the deal-me-out

[18]Since every subject responded once to each questionnaire item, there were also a total of 108 responses to each question for the up treatment and 228 responses for the down treatment.

outcome: for each value of α the median claims depart from the deal-me-out claims of $\min\{5, 10 - \alpha\}$ by less than 50 cents. The median claims of those player *II*s who chose not to opt out were within 10 cents of the deal-me-out claims of $\max\{5, \alpha\}$. With one exception, the expected-profit-maximizing claim for a player *II* who chose not to opt out was within 30 cents of $\max\{5, \alpha\}$. Finally, after the bargaining sessions were over, subjects were asked what would be a fair claim for a player who could opt out. For each opt-out payment, the median claim designated as fair was $\max\{5, \alpha\}$.

Opting out: When α was large, player *II* frequently chose to opt out, yielding an inefficient outcome. The opt-out frequencies for $\alpha = \$4.90$, $\alpha = \$6.40$, and $\alpha = \$8.10$ were .33, .61, and .83, respectively.

Rows 1a and 1b of Figures 4 and 5 describe the player-*II* claims made by players who chose not to opt out, with the median, 5*th* percentile and 95*th* percentile claims indicated (the latter two being the first and second numbers in parentheses, respectively) in each case.[19] Rows 2a and 2b similarly report player *I* claims, while rows 3a and 3b report the subjects' estimates of what would be a "fair" claim in each case. The median claims of both player *I*s and player *II*s reported in rows 1a, 1b, 2a and 2b and the median claims indicated as fair for player *II* in rows 3a and 3b correspond well to the predictions of the deal-me-out outcome. Moreover, the 95th percentiles reported for the player *II* claims indicate that virtually no player *II* expected to receive much more than the deal-me-out claim. In addition, rows 4a and 4b of Figures 4 and 5 show that player *II* rarely opted in and made a claim less than α, much less made a disadvantageous offer.[20]

Player *I* behavior is also generally consistent with the deal-me-out outcome, though the 5th percentiles for the player *I* claims for opt-out payments \$6.40 and \$8.10 show that at least some subjects made claims as player *I* that were close to the split-the-difference outcome.

[19] For 108 observations, the 5th percentile is calculated as the mean of the 6th and 7th order statistics, that is, the 6th and 7th elements of a list of the observations sorted from lowest to highest. The 95th percentile is the mean of the 102th and 103th order statistics and the median or 50th percentile is the mean of the 54th and 55th order statistics. For 228 observations, the 5th, 50th, and 95th percentiles are given by the means of the 12th and 13th, the 114th and 115th, and the 216 and 217 order statistics, respectively.

[20] By "disadvantageous offer" we mean a claim making it impossible for player *II* to receive a payoff of at least α. Claiming less than α may yield payoffs higher than α as long as some player *I*s claim less than $10 - \alpha$.

Opt-out payments[21]	0.90	2.50	4.90	6.40	8.10
1a. Median claims of player IIs who did not opt out	4.95	5.00	5.00	6.40	8.10
	(4.15 – 5.20)	(3.95 – 5.10)	(4.90 – 5.20)	(5.10 – 6.80)	(5.05 – 8.40)
2a. Median claims of player Is	4.90	4.90	4.70	3.25	1.60
	(4.20 – 5.25)	(4.30 – 5.00)	(2.75 – 5.00)	(2.20 – 4.80)	(0.90 – 4.00)
3a. Median claims asserted as fair for player II	5.00	5.00	5.00	6.40	8.10
	(4.40 – 5.50)	(4.05 – 6.15)	(4.90 – 6.80)	(5.05 – 7.35)	(4.90 – 8.70)
4a. Frequency with which player IIs made claims less than their outside option	0.000	0.000	0.019	0.065	0.028
5a. Frequency of player I claims greater than 4.50	0.880	0.852	0.574	0.074	0.037
6a. Frequency of player I claims providing player II a payoff lower than the opt-out value	0.000	0.000	0.009	0.083	0.102
7a. Frequency with which player IIs opted out	0.000	0.019	0.343	0.556	0.750
8a. Mean profit of player IIs who did not opt out.[22]	4.43	4.63	4.72	5.79	7.54
	(0.00,5.00,5.27)	(0.00,5.00,5.27)	(0.00,5.15,5.65)	(0.00,6.60,7.28)	(3.35,8.25,8.53)
9a. Maximum expected profit of a player II who does not opt out.[23]	4.77	5.00	5.23	6.13	7.47
	(4.90)	(4.90)	(5.00)	(6.40)	(7.80)
10a. Maximum expected profit of a player I[23]	4.78	4.89	4.70	3.31	1.89
	(4.80)	(4.80)	(4.80)	(3.10)	(1.60)

Figure 4: **Summary Data for "UP" Treatment**

Opt-out payments[21]	0.90	2.50	4.90	6.40	8.10
1b. Median claims of player *II*s who did not opt out	5.00	4.90	5.00	6.40	8.10
	(3.85 – 5.20)	(3.90 – 5.20)	(4.90 – 5.70)	(4.95 – 7.00)	(4.85 – 8.35)
2b. Median claims of player *I*s	4.90	4.90	4.60	3.15	1.70
	(4.10 – 5.20)	(4.00 – 5.10)	(3.45 – 5.00)	(2.05 – 4.85)	(1.00 – 5.00)
3b. Median claims asserted as fair for player *II*	5.00	5.00	5.00	6.40	8.10
	(2.20 – 5.45)	(2.50 – 5.70)	(4.35 – 5.70)	(4.10 – 6.95)	(5.00 – 8.45)
4b. Frequency with which player *II*s made claims less than their outside option	0.000	0.000	0.013	0.105	0.061
5b. Frequency of player *I* claims greater than 4.50	0.846	0.759	0.535	0.083	0.154
6b. Frequency of player *I* claims providing player *II* a payoff lower than the opt-out value	0.000	0.000	0.009	0.140	0.237
7b. Frequency with which player *II*s opted out	0.000	0.000	0.325	0.640	0.868
8b. Mean profit of player *II*s who did not opt out.[22]	4.42	4.66	4.88	5.45	6.51
	(0.00,5.00,5.40)	(0.00,5.00,5.45)	(0.00,5.20,6.00)	(0.00,6.57,7.33)	(0.00,8.15,8.70)
9b. Maximum expected profit of a player *II* who does not opt out.[23]	4.85	4.96	5.21	5.82	6.37
	(4.80)	(4.80)	(5.00)	(5.00)[24]	(8.00)
10b. Maximum expected profit of a player *I*[23]	4.77	4.89	4.56	3.24	2.21
	(4.80)	(4.80)	(4.20)	(3.00)	(1.60)

Figure 5: **Summary Data for "DOWN" Treatment**

Rows 5a and 5b report the frequencies with which player Is made claims greater than 4.50 dollars, i.e., claims that were close to the fifty-fifty prediction. Rows 6a and 6b report the frequencies with which player Is made claims greater that 10 dollars minus the opt-out payment. Such claims give player II a smaller playoff than opting out. For both treatments, the frequencies reported in rows 5a, 5b, 6a and 6b are relatively small for opt-out payments that exceed half the cake, which is consistent with deal-me-out. The most noticeable differences in the data from the up and down treatments are the larger frequencies with which player Is in sessions with the down treatment made claims that did not leave player II with a payoff larger than α when the opt-out payment was \$6.40 or \$8.10.[25]

The deal-me-out solution thus matches player I's behavior reasonably well and matches player II's behavior reasonably well *when player II opts in*. Contrary to the deal-me-out prediction, however, player II frequently opts out. Rows 7a and 7b in Figures 4 and 5 reports the frequency with which player IIs chose to opt out.[26]

Why do player II's opt out? Rows 8a and 8b describe the profit achieved by player IIs who did not opt out.[27] The first number is mean profit. In addition to the 5th and 95th percentiles, the middle number reported in parentheses in rows 8a and 8b is the median profit obtained by the II player IIs who did not opt out. In each case, the median profit is the same as or slightly larger than the opt-out payment or half the cake, whichever is larger. On the other hand, the *mean* profit is always lower than the median profit, often by a substantial margin. The difference between the mean and

---

[21] Except where noted, the statistic reported is the median of the observations pooled over all subjects who participated in experiments with the same treatment, and the numbers in parentheses are the 5th and 95th percentile of the observations. Claims and profits are in dollars.

[22] Numbers in parenthesis are respectively the 5th percentile, median, and 95th percentile of the profits obtained by player IIs who did not opt out.

[23] See text for details. The number in parenthesis is the optimal claim for such a player.

[24] The expected profit function was not always a unimodal function of the subject's claim. In this case, for example, there was a second local maximum at 6.40. The expected profit obtained by making a claim of 6.40 was 5.76 dollars.

[25] As one might expect, the larger frequencies of such claims for the down treatment coincide with lower mean profits for player IIs as reported in rows 8a and 8b.

[26] The overall opt-out frequencies reported in the summary are the weighted average of the frequencies reported in Figures 4–5 for the up and the down treatments.

[27] For the larger opt-out payments, many player IIs chose to opt out; hence, for these opt-out payments, the numbers of observations summarized in rows 1a,1b,8a and 8b are much smaller than 108 or 228.

median profit is one measure of the risk of disagreement. This risk pushes the mean profit of player *II*s who did not opt out below the opt-out payment for the three largest opt-out payments.

The mean profit reported in rows 8a and 8b involves only those player *I* claims that were actually matched with player *II*s who did not opt out. In contrast, rows 9a and 9b report the maximum expected profit that a player *II* could achieve when playing against the entire population of player *I* claims made in bargaining sessions with the designated opt-out payment and treatment. The numbers reported in parentheses are the player *II* claims that achieve this expected profit.[28] The maximum possible expected profit obtained by not opting out is less than the opt-out payment for those opt-out payments which exceed $5.00.[29]

These experimental results reflect the tension between optimization and efficiency. In their quest for a hard bargain, player *I*s push player *II*s toward the fifty-fifty outcome. If the outside option for player *II* is enough smaller than $5.00 (i.e., outside options $.90 and $2.50), then the system settles on the fifty-fifty outcome. For higher outside options, hard bargaining on the part of player *I* pushes player *II* to a claim very close to her outside option, with player *I* claiming the rest. This is the deal-me-out outcome, and we expect the system to settle there in a perfect world. However, the experimental world is not perfect. Instead, hard bargaining sometimes leads to disagreements, and this causes the deal-me-out outcome to give player *II*s lower mean payoffs than their outside options. As a result, player *II*s often opt out and the gains from trade go unrealized.

6 Conclusion

We believe that the phenomena studied in this paper are widespread and that efficiency will therefore sometimes fail to be achieved even when the agents involved are as close to being rational as real people are ever likely to get. Contract theorists have recently devoted considerable attention to the problem of avoiding the expropriation of rents meant to compensate parties

[28]In a similar fashion, rows 10a and 10b report the maximum expected profit and the optimal claim for a player *I* who is matched randomly with one from the designated population of claims made by player *II*s who did not opt out.

[29]Because row 9 involves a larger sample of player *I*s, it is not contradictory that the maximum profit in row 9b (the down treatment) falls short of the mean profit in row 8b for outside option 6.40.

for sunk costs. Our results indicate that attention must also be devoted to providing appropriate compensation for opportunity costs that are not sunk at the time contracting occurs.

The traditional line taken by such authors as Coase [17] or Williamson [47] is that efficiency is guaranteed if the parties to the deal have a costless opportunity to negotiate a binding contract before the costs are sunk. More generally, they argue that new property rights and new forms of contracting will emerge to deal with the inefficiencies that can result from a variety of frictions that the literature bundles together under the catch-all notion of a transaction cost. Our paper can be reconciled with this literature by classifying the learning frictions that we study as yet another form of transaction cost whose existence calls for the appearance of new institutions. An obvious possibility is the replacement of primitive bargaining institutions like those built into the Outside Option Game by more sophisticated schemes, but it is necessary to recall that we first turned our attention to the opting-out phenomenon because of its appearance in Rubinstein bargaining models (Binmore *at al* [6, 13]). A more hopeful development might be the increased use of arbitration agencies or bargaining consultancies in those cases where the problem cannot otherwise be internalized. But in the absence of such institutional crutches, it seems wise to soft pedal the claim that all gains from trade will necessarily be exploited in a sufficiently rational society.

7 Appendix: Instructions to Subjects

Bargaining Experiment

In this experiment, you will bargain via the computing equipment in front of you with people seated at other machines in the room. You will participate in a large number of very short bargaining sessions. Whether you are player *I* or player *II* in these sessions is determined randomly. Sometimes you will be player *I* and sometimes player *II*. After each session, you will be randomly paired with a new bargaining partner.

In each bargaining session, you and your counterpart for that session will have the opportunity to split a "cake" between you. You will each simultaneously make a claim. If the two claims sum to no more than the value of the cake, then each of you will receive their claim plus half the surplus after the claims have been met. If the two claims sum to more than the value of the cake, each of you will get *nothing* at all.

Only one thing complicates this very simple scenario. Before each bargaining session begins, player *II only* is offered the opportunity of opting out. If player *II* opts out, he or she gets a payment that may vary from session to session. But, in each session, both players will know what player *II*'s opting out payment is for that session. Player *I* gets *nothing* if player *II* opts out.

The cake is always *nominally* worth $10 but you will be paid the amounts you succeed in securing only for *two* of the bargaining sessions. These will be chosen at random from the final ten sessions in which you participate. The preceding two sets of ten sessions are for practice. In each of these two sets of ten practice sessions, you will also be paid for two sessions chosen at random, but you will not be paid at the full rate. In the first set of ten practice sessions you will be paid at the rate of one dime for each nominal dollar. In the second set of ten practice sessions, you will be paid at the rate of one quarter for each nominal dollar. Only in the third set of ten sessions will you be bargaining for real and getting paid at the full rate for the two sessions the computer chooses at random.

After the bargaining sessions are over, you will be asked to complete a computerized questionnaire. Money prizes will be awarded during the questionnaire for answers to some questions.

When *all* subjects have completed the questionnaire, the computer will display how much money you have earned during the experiment. This will include the amounts you secured during the bargaining, and any prizes you won while completing the questionnaire. It will not include your $2 attendance fee. Please remain in your seat until the supervisor calls your seat number and then bring your seat tag so that you can be paid.

This is not an experiment to find out what kind of person you are. When we see the results, we shall neither know nor care who did what. We are only interested in what happens on average. So please don't feel that some particular sort of behavior is expected of you. However, we do ask that you do not talk to the other subjects or look at their screens. It is important to the experiment that our subjects interact *only* through the computer equipment.

Now press the SPACE BAR on your keyboard. You will see a demonstration that will review the information in these instructions and give you hands-on experience in making claims or opting out. Remember to keep pressing the SPACE BAR to see a new screen. There is no need to hurry. You may have to wait for the other subjects to be ready anyway. If you still have questions after seeing the demonstration, there will be an opportunity to ask the supervisor.

References

[1] Armen Alchian. Uncertainty, evolution, and economic theory. *Journal of Political Economy*, 58:211–221, 1950.

[2] James Andreoni and John H. Miller. Auctions with artificial adaptive agents. Mimeo, University of Wisconsin and Carnegie Mellon University, Madison, Wisconsin and Pittsburg, Pennsylvania, 1993.

[3] Dieter Balkenborg. Tests of the forward induction hypothesis. Mimeo, University of Bonn, 1994.

[4] K. Binmore, J. Swierzsbinski, S. Hsu, and C. Proulx. Focal points and bargaining. *International Journal of Game Theory*, 22:381–409, 1993.

[5] Ken Binmore, John Gale, and Larry Samuelson. Learning to be imperfect: The ultimatum game. *Games and Economic Behavior*, 8:56–90, 1995.

[6] Ken Binmore, Peter Morgan, Avner Shaked, and John Sutton. Do people exploit their bargaining power? An experimental study. *Games and Economic Behavior*, 3:295–322, 1991.

[7] Ken Binmore, Martin J. Osborne, and Ariel Rubinstein. Noncooperative models of bargaining. In Robert J. Aumann and Sergiu Hart, editors, *Handbook of Game Theory*, pages 179–226. North Holland, New York, 1994.

[8] Ken Binmore, Ariel Rubinstein, and Asher Wolinsky. The nash bargaining solution in economic modelling. *RAND Journal of Economics*, 17:176–188, 1986.

[9] Ken Binmore and Larry Samuelson. Muddling through: Noisy equilibrium selection. Ssri working paper 9410, University of Wisconsin, 1993.

[10] Ken Binmore and Larry Samuelson. Evolutionary drift and equilibrium selection. Mimeo, University College London and University of Wisconsin, 1994.

[11] Ken Binmore, Larry Samuelson, and Richard Vaughan. Musical chairs: Modelling noisy evolution. *Games and Economic Behavior*, 1995. Forthcoming.

[12] Ken Binmore, Avner Shaked. and John Sutton. Testing noncooperative bargaining theory: A preliminary study. *American Economic Review*, 75:1178–1180, 1985.

[13] Ken Binmore, Avner Shaked, and John Sutton. An outside option experiment. *Quarterly Journal of Economics*, 104:753–770, 1989.

[14] Gary E. Bolton. A comparative model of bargaining: Theory and evidence. *American Economic Review*, 81:1096–1136, 1991.

[15] T. Börgers and R. Sarin. Learning through reinforcement and replicator dynamics. Mimeo, University College London, 1993.

[16] Antonio Cabrales. Stochastic replicator dynamics. Mimeo, University of California, San Diego, 1993.

[17] Ronald Coase. The problem of social cost. *Journal of Law and Economics*, 2:1–40, 1960.

[18] Russell Cooper, Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross. Alternative institutions for resolving coordination problems: Experimental evidence on forward induction and preplay communication. In James W. Friedman, editor, *Problems of Coordination in Economic Activity*, pages 129–146. Kluwer Academic Publishers, Boston, 1994.

[19] Vincent P. Crawford. An "evolutionary" interpretation of Van Huyck, Battalio, and Beil's experimental results on coordination. *Games and Economic Behavior*, 3:25–59, 1991.

[20] Vincent P. Crawford. Adaptive dynamics in coordination games. *Econometrica*, 1995. Forthcoming.

[21] Eddie Dekel and Drew Fudenberg. Rational behavior with payoff uncertainty. *Journal of Economic Theory*, 52:243–267, 1990.

[22] Milton Friedman. *Essays in Postive Economics*. University of Chicago Press, Chicago, 1953.

[23] W. Guth, R. Schmittberger, and B. Schwarze. An experimental analysis of ultimatum bargaining. *Journal of Behavior and Organization*, 3:367–388, 1982.

[24] Werner Güth and Reinhard Tietz. Ultimatum bargaining behavior: A survey and comparison of experimental results. *Journal of Economic Psychology*, 11:417–49, 1990.

[25] G. W. Harrison and M. McKee. Experimental evaluation of the Coase theorem. *Journal of Law and Economics*, 28:653–670, 1985.

[26] Morris W. Hirsch and Stephen Smale. *Differential Equations, Dynamical Systems, and Linear Algebra.* Academic Press, New York, 1974.

[27] E. Hoffman and M. Spitzer. The Coase theorem: Some experimental tests. *Journal of Law and Economics*, 25:73–98, 1982.

[28] E. Hoffman and M. Spitzer. Entitlements, rights and fairness: An experimental examination of subjects' concepts of distributive justice. *Journal of Legal Studies*, 14:259–297, 1985.

[29] Elizabeth Hoffman, Kevin McCabe, Keith Shachat, and Vernon Smith. Preferences, property rights and anonymity in bargaining games. Mimeo, University of Arizona, 1992.

[30] Elon Kohlberg and Jean-Francois Mertens. On the strategic stability of equilibria. *Econometrica*, 54:1003–1038, 1986.

[31] Ian M. McDonald and Robert R.Solow. Wage bargaining and employment. *American Economic Review*, 71:896–908, 1981.

[32] John H. Miller and James Andreoni. Can evolutionary dynamics explain free riding in experiments? *Economics Letters*, 36:9–15, 1991.

[33] Barry Nalebuff and Avinash Dixit. *Thinking Strategically.* W. W. Norton and Company, New York, 1991.

[34] John F. Nash. Two-Person cooperative games. *Econometrica*, 21:128–140, 1953.

[35] Georg Nöldeke and Larry Samuelson. An evolutionary analysis of backward and forward induction. *Games and Economic Behavior*, 5:425–454, 1993.

[36] Jack Ochs and Alvin E. Roth. An experimental study of sequential bargaining. *American Economic Review*, 79:355–384, 1989.

[37] Martin J. Osborne and Ariel Rubinstein. *Bargaining and Markets*. Academic Press, New York, 1990.

[38] Chris Proulx. A numerical analysis of the replicator dynamics applied to an outisde option game. Mimeo, University of Michigan, 1995.

[39] Alvin E. Roth. Bargaining experiments. In John Kagel and Alvin E. Roth, editors, *Handbook of Experimental Economics*. Princeton University Press, 1994. Forthcoming.

[40] Alvin E. Roth and Ido Erev. Learning in extensive–form games: Experimental data and simple dynamic models in the intermediate term. Mimeo, University of Pittsburgh and The Technion, 1993.

[41] Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50:97–109, 1982.

[42] Larry Samuelson and Jianbo Zhang. Evolutionary stability in asymmetric games. *Journal of Economic Theory*, 57:363–391, 1992.

[43] Karl H. Schlag. Why imitate, and if so, how? Mimeo, University of Bonn, 1994.

[44] Richard H. Thaler. *The Winner's Curse*. Princeton University Press, Princeton, 1992.

[45] Eric van Damme. Stable equilibria and forward induction. *Journal of Economic Theory*, 48:476–509, 1989.

[46] Jürgen Weibull. Evolutionary game theory. Mimeo, University of Stockholm, 1994.

[47] Oliver Williamson. *The Economic Institutions of Capitalism*. Free Press, New York, 1985.

Institut für Höhere Studien
Institute for Advanced Studies
Department of Economics
Stumpergasse 56
A-1060 Vienna
Austria

Tel.: +43-1-599 91-149
Fax.: +43-1-599 91-163
e-mail: woergoet@ihssv.wsr.ac.at