# Predicting Cooperation with Welfare Maximization and Trembles

David K. Levine[1]

**Abstract**

The indefinitely repeated prisoner's dilemma game provides a paradigm for long-term cooperation in the face of short-term incentives for free riding. However, the extent to which players cooperate in the laboratory depends on the parameters of the game. To understand this, I take a simple, direct approach motivated by theory. I hypothesize that players tremble and that they find the best perfect public equilibrium of the induced game of imperfect public information. I calibrate the probability of trembling, using no data from repeated game experiments. This model makes sharp predictions and does a good job both qualitatively and quantitatively in explaining welfare and cooperation rates.

[1]Department of Economics, RHUL

*Acknowledgements:* First version: March 25, 2024. I would like to thank Alessandra Casella, Gary Charness, Vince Crawford, Michalis Drouvelis, Rohan Dutta, Martin Dufwenberg, Ernst Fehr, John Mair, Andrea Mattozzi, Agnieszka Mensfelt, Salvatore Modica, Tom Palfrey, Klaus Ritzberger, Larry Samuelson, Arthur Schram, Francesco Squintani, Catharina Tilmans, Myrna Wooders, participants at the 2025 SAET Conference in Ischia, the experimental workshop at the University Cote d'Azur, at Montpellier University, and the European Economic Review Experimental Summer School in Lisbon. The Leverhulme Trust provided financial support for which I am grateful. My knowledge of the theory and experiments discussed here comes from decades of work with Drew Fudenberg. A particular debt is owed to Pedro Dal Bó and Guillaume Fréchette not only for their intellectual contributions to this work, but for providing me with their original experimental data and the data from their own meta-study. I am similarly grateful to Drew Fudenberg and Gustav Karreskog Rehbinder for making available the data from their meta-study.

## 1. Introduction

This is a meta-study of the indefinitely repeated prisoners' dilemma game. My hypothesis is that the long-term level of cooperation can be well explained as the welfare-maximizing equilibrium in the game in which there is a probability that players tremble uniformly in each period. I refer to this trembling probability as the level of noise. While this paper is about the indefinitely repeated prisoners' dilemma game, the theory is a general one that applies to all games.

There are several ways to pose a question about how well the theory works. First, for a particular treatment, is there a level of noise that well explains the extent of cooperation? Second, while it is unreasonable to suppose that very different populations have the same level of noise, is it the case that there is, for similar populations, in this case, Western college students, a single level of noise that well explains the extent of cooperation across different treatments? An even more stringent question is whether an externally valid level of noise, observed in other types of games entirely, well explains the extent of cooperation in the indefinitely repeated prisoners' dilemma game. There is: the trembling probability of one-third, shown in Levine (2025) to work well in a variety of games, including coordination games, ultimatum bargaining, public goods with punishment and gift exchange, well explains the extent of cooperation in the indefinitely repeated prisoners' dilemma.

The fact that a trembling probability of one-third works well, leads to two new questions. First, is there a wide range of noise levels that explain the level of cooperation, of which one-third is just one of many? There is not: the the data best fits a probability of trembling between 0.32 and 0.34. That leads to the final question: is there another level of noise that does a better job of explaining the extent of cooperation than one-third? There is not: while there are other noise levels between 0.32 and 0.34 that work well, there is little appreciable difference between them.

The infinitely repeated prisoner's dilemma game provides a paradigm for long-term cooperation in the face of short-term incentives for free riding. Such an environment is induced in the laboratory by ending the game with a roll of the dice. In a repeated indefinite horizon prisoner's dilemma game, players are then rematched with new opponents to play another indefinite horizon prisoner's dilemma game. This enables us to see in the laboratory the long-term play of experienced players - the type of play that is most relevant in most applications outside the laboratory. Theoretically, when the discount factor is sufficiently high, on account of the folk theorem of Fudenberg

and Maskin (1986), any degree of cooperation is possible in equilibrium. However, game theorists have long believed, or at least hoped, that players will find their way to efficient equilibria. There are three theories of behavior that motivate this belief and hope, because they point in the direction of players coordinating on welfare-maximizing equilibria. The first is Schelling (1960)'s theory of focal points, the second is that players reason as rule-utilitarians rather than act-utilitarians. The third is that evolutionary forces favor welfare-maximizing equilibria.

Schelling (1960)'s theory of focal points proposes that players in a coordination game coordinate because they agree that a certain equilibrium is focal. The simplest class of games with multiple equilibria are pure coordination games. In these game, the welfare-maximizing equilibrium is thought to be focal. Since the welfare-maximizing equilibrium is generically unique,[2] this can be focal in other games, such as the indefinitely repeated prisoner's dilemma.

A second theory that leads to players coordinating on the welfare-maximizing equilibrium is that they are social minded and reason as rule utilitarians. Rule utilitarianism is an idea dating back to Mill (1861), and is described by Garner and Rosen (1967) as "the rightness or wrongness of a particular action is a function of the correctness of the rule of which it is an instance." This is in contrast to act utilitarianism, a more commonly used concept in economics. A rule utilitarian player follows the rule that answers the question: "what rule that everyone would be willing to follow would be best for the common good?" The idea that this type of reasoning is an individualistic way to solve coordination problems was proposed in Harsanyi (1977). It is closely related to Roemer (2010)'s notion of Kantian players. In recent years, the theory has been widely adapted in voting theory, where it is sometimes called the theory of ethical voters. It was first studied theoretically by Feddersen and Sandroni (2006) and empirically by Coate and Conlin (2004). Subsequently, the model has been used by other authors, such as Herrera, Morelli and Nunnari (2016), Levine and Mattozzi (2020), and Levine, Mattozzi and Modica (2022).

There are also a variety of theoretical results that point to the evolution of agents who find welfare-maximizing solutions to coordination problems. Juang and Sabourian (2021) and Levine (2026) give conditions under which individual evolution leads to this conclusion. Work by Ely (2002), Levine and Modica (2013)

[2]See Levine (2026).

and Levine and Modica (2016) shows how competition between groups also leads to individuals who are able to coordinate.

Outside the laboratory, two Nobel prizes have been awarded for documenting the ability of people to coordinate on good equilibria: Coase (1960) and Ostrom (1990). I would add to this also the work of Townsend (1994) on insurance in rural societies. Indeed: the main premise of contract theory is that people are pretty good at figuring out good solutions to social dilemmas. Outside the laboratory, of course, more time is available, and people can talk, but the problems tend to be more complex.

While there is ample precedent for the idea that players coordinate on the welfare-maximizing equilibrium, I should be clear that this is not strongly established behavioral law, and should be viewed as a parsimonious benchmark. Indeed, the current state of belief is that welfare maximization does not do a good job in predicting play in experimental games. There are two reasons for this. The first is theoretical, based on either evolutionary theory, or lack of common knowledge. The second is empirical based on laboratory findings.

From a theoretical point of view there has been pushback, arguing that payoff dominance, which in the symmetric case is the same as welfare maximization, is less important than risk dominance. Harsanyi and Selten (1988) used the idea of risk dominance as part of their equilibrium selection criterion: in a $2 \times 2$ symmetric game, a pure strategy equilibrium is risk-dominant if it is a best response when half the players are playing that pure strategy. Work on evolutionary game theory, particularly Kandori, Mailath and Rob (1993) and Young (1993), showed that under certain conditions, risk-dominant equilibria were the unique stochastically stable outcome of an evolutionary process. An alternative theoretical approach is the global games approach of Carlsson and Van Damme (1993), based on failure of common knowledge, which leads to the same conclusions for $2 \times 2$ games: it is risk dominance, not payoff dominance, that matters.

I have now discussed two evolutionary literatures: one favors payoff dominance, the other risk dominance. In light of the results of Bergin and Lipman (1996) showing that evolutionary dynamics is not terribly robust; perhaps this is not surprising, so the question should be an empirical one. However, in the experimental arena, the minimum game experiment of Van Huyck, Battalio and Beil (1990) has been thought to provide strong evidence that players are not able to coordinate when risk dominance is an issue. Subsequently, in the indefinitely repeated prisoners' dilemma,

Dal Bó and Fréchette (2011) argued that risk dominance is a better explanation of long-run behavior than payoff dominance. These indefinitely repeated prisoner's dilemma games are the specific topic of this paper. I am going to argue that in these games and more broadly, welfare maximization as an equilibrium selection criterion has been abandoned too quickly.

Let me start with the Van Huyck, Battalio and Beil (1990) minimum game. This is a public goods contribution game with seven levels of contribution, in which only the minimum contribution matters. Payoffs are such that all players agreeing on any level of contribution is a Nash equilibrium, with the welfare-maximizing equilibrium being the maximum contribution by all players. When played with fifteen or more experienced players, it is observed that, rather than coordinating on the welfare-optimal highest level of effort, instead the minimum level of effort observed is the least feasible level of effort. My observation is this: with 15 or more players, in a game where only the minimum level of contribution to a public good matters, then if players tremble, it is highly likely that one trembling player will ruin things for everybody.

Specifically, (see Levine (2025)) if 1/3 of the players tremble uniformly then, in fact, the only equilibrium is for all non-trembling players to provide the least effort. Hence, in the relevant normal form; the one with trembling, they do coordinate on the welfare-optimal equilibrium, as there is only one equilibrium. Now 1/3 of the players trembling uniformly may seem like a lot: 29% of the time they should be playing more than the least level of effort. In fact, in the Van Huyck, Battalio and Beil (1990) data in the final of ten periods, and despite the fact that in the previous six periods the minimum level of effort was the least possible; 29% are providing more than the least level and, indeed, 10% are providing the highest level.

This is my theory: players are good at coordinating, but they also tremble at a substantial rate, and apparent failures of coordination are not due to risk dominance or global games, but rather due to the fact that the environment is noisy. This paper specifically demonstrates that this noise theory works well for indefinitely repeated prisoners' dilemma games. At the end of the paper, I also revisit the meta-study of Dal Bó, Fréchette and Kim (2021) on stag hunt games, again where there is apparent coordination failure, and show that the same probability of trembling, one-third, together with welfare maximization, does a good job in predicting the location and size of the bifurcation between the predominantly Hare and Stag outcomes.

The idea of using noise in understanding experimental play is not new. The noise model here is related to the extensive form quantal response model of McKelvey and Palfrey (1998). However, the level of noise in that model is endogenous, so that the theory of perfect public equilibrium cannot be directly applied. In addition, there is no single value of the noise parameter in quantal response that applies across a broad range of experiments and so can be applied to the games studied here. There is also some recent work done on trembling in the repeated prisoner's dilemma, for example in Graser and van Veelen (2024), albeit, studying strategic issues for small error rates, and not examining experimental data. Regardless, the new element here is in combining a noisy environment with welfare-maximizing coordination.

There is one point I want to mention in this context. In the literature, there is discussion of efficiency, or payoff dominance. In reality, there are two types of efficiency: *ex ante*, before player roles are assigned, and *ex post,* after they are assigned. The former corresponds to efficiency behind Rawls (1971) veil of ignorance and implies welfare maximization. In the experimental setting, where player roles are assigned randomly; this is, perhaps, more compelling than in the world at large. Regardless: all the game studied here are symmetric games, so that the *ex ante* vs. *ex post* distinction is irrelevant.

In order to explain how the degree of long-term cooperation depends on the specific payoff parameters of the game researchers, such as Dal Bó and Fréchette (2011), Blonski, Ockenfels and Spagnolo (2011), Blonski and Spagnolo (2015) and Fudenberg and Rehbinder (2024), have employed a variety of models motivated by individual learning. They have developed theories specifically designed to explain cooperation in the indefinitely repeated prisoners' dilemma. I examine these theories. The simple theories, risk-dominance and global games, do a very poor job quantitatively in explaining the data. Theories specifically designed to explain cooperation in the indefinitely repeated prisoners' dilemma, and estimated from the data with from two to six parameters, do a good job, but are not appreciably different in fit from the simple and general theory - one that applies to all games - of welfare maximization with a one-third chance of trembling. I defer further discussion of these theories to section 5, where I explicitly compare the results of this theory to these earlier theories of cooperation in repeated games. There I show that performance of noise theory, with a noise level that is externally valid, is similar to theories that are estimated from the repeated game data.

## 2. The Model

Two players $i \in \{1,2\}$ repeatedly play a prisoner's dilemma stage game. They take actions $a^i \in \{C,D\}$, $C$ for cooperate and $D$ for defect. Payoffs are symmetric and given by $u^i(a)$ where $a$ denotes the profile. Following the literature, payoffs are normalized so that $u^i(CC) = 1$ and $u^i(DD) = 0$. After each stage game, the match continues with probability $0 < \delta < 1$. Welfare in each period is given by $\omega(a) = (u^1(a) + u^2(a))/2$.

Players tremble uniformly, iid across players and time, with probability $1 > \phi > 0$. This makes it necessary to distinguish between intentions and actions. Players choose intentions, so the game is one of imperfect public information, where the public signals are the outcomes. Following Fudenberg, Levine and Maskin (1994), the solution concept is perfect public equilibrium. The intention profile of the two players induces a probability distribution over outcomes $y \in \{CC, CD, DC, DD\}$. Let $\alpha^i$ denote the probability that $i$ intends to cooperate, and denote by $\pi(y|\alpha)$ the induced probability distribution over outcomes.

I do not attempt to characterize the full set of perfect public equilibria. Rather, as a behavioral restriction on the strategy class, I restrict strategies to have one period of memory. Because of the nature of matching in the laboratory, in equilibrium, they must be symmetric. They are given by $\sigma^i(y)$, a probability vector over intentions $a^i \in A$, conditional on the previous period outcome $y \in \{CC, CD, DC, DD\}$. This is used together with an initial strategy $\sigma_0^i$, also a probability vector over $a^i \in A$. This class of strategies is commonly used in studying experimental repeated games, for example, by Fudenberg and Rehbinder (2024).

For any initial $y$ and one-period memory strategies $\sigma^i$ average expected utility is denoted by $V(y|\sigma)$.

### *2.1. Equilibrium*

The memory-one profile $\sigma$ is *incentive compatible* if for all $i \in 1,2$, $d^i \in \{C,D\}$, $y \in \{CC, CD, DC, DD\}$

$$V^i(y|\sigma) \geq (1-\delta)U^i(d^i, \sigma^{-i}(y)) + \delta \sum_{a^{-i} \in A} \sum_{y' \in A^2} \sigma^{-i}(a^{-i}|y)\pi(y'|d^i, a^{-i})V^i(y'|\sigma).$$

A pair consisting of an initial profile and memory-one profile $(\sigma_0, \sigma)$ is *incentive compatible* if $\sigma$ is incentive compatible and for all $i \in \{1,2\}$,$d^i \in \{C,D\}$

$$\overline{V}^i(\sigma_0,\sigma) \equiv (1-\delta)U^i(\sigma_0) + \delta \sum_{a\in A^2}\sum_{y'\in A^2} \sigma_0^i(a^i)\sigma_0^{-i}(a^{-i})\pi(y'|a)V^i(y'|\sigma) \geq$$

$$(1-\delta)U^i(d^i,\sigma_0^{-i}) + \delta \sum_{a^{-i}\in A}\sum_{y'\in A^2} \sigma_0^{-i}(a^{-i})\pi(y'|d^i,a^{-i})V^i(y'|\sigma).$$

Let $\Sigma$ denote the set of all symmetric incentive compatible pairs $(\sigma_0,\sigma)$.

Maximizing welfare subject to the incentive constraints defines

$$\hat{\omega} = \max_{(\sigma_0,\sigma)\in\Sigma} \overline{V}^i(\sigma_0,\sigma).$$

*Noise theory* predicts that this should describe long-run average welfare in laboratory play.

The empirical long-run average that the theory is designed to describe is, for a particular treatment, defined as

$$\overline{\omega} \equiv (1/T)\sum_{t=1}^{T}(u^1(a_t)+u^2(a_t))/2,$$

where $t = 1,\ldots T$ indexes the individual games in the treatment and $a_t$ denotes the outcome. Note that this accounts for discounting, as discounting determines how long matches last.

*2.2. Details*

Recall that $\alpha^i$ is the probability that $i$ intends to cooperate. The probability that $i$ actually cooperates is then given by $\tilde{\alpha}^i = (1-\phi/2)\alpha^i + (\phi/2)(1-\alpha^i)$. The intention profile of the two players induces a probability distribution over outcomes $y \in \{CC, CD, DC, DD\}$ by

$$\pi(y|\alpha) = \begin{cases} \tilde{\alpha}^1\tilde{\alpha}^2 & y = CC \\ \tilde{\alpha}^1(1-\tilde{\alpha}^2) & y = CD \\ (1-\tilde{\alpha}^1)\tilde{\alpha}^2 & y = DC \\ (1-\tilde{\alpha}^1)(1-\tilde{\alpha}^2) & y = DD \end{cases}.$$

Utility in the one-shot game as a function of intentions is accordingly given by $U^i(\alpha) = \sum_{y \in A^2} \pi(y|\alpha) u^i(y)$.

For any initial $y$ and one-period memory strategies $\sigma^i$ average expected utility is given recursively by

$$V^i(y|\sigma) = (1-\delta)U^i(\sigma^i(y)) + \delta \sum_{a \in A^2} \sum_{y' \in A^2} \sigma^i(a^i|y)\sigma^{-i}(a^{-i}|y)\pi(y'|a)V^i(y'|\sigma^i)$$

Define the $4 \times 4$ matrix for $y, y' \in A^2$ by

$$H(y, y'|\sigma) = \sum_{a \in A^2} \sigma^i(a^i|y)\sigma^{-i}(a^{-i}|y)\pi(y'|a)$$

then $V^i(\sigma) = (I - \delta H(\sigma))^{-1}(1-\delta)U^i$, where $U^i$ is the 4-vector of utilities $U^i(\sigma(y))$.

*2.3. Computation*

Maximization took place searching over a grid of $K+1$ points for each component $\sigma^i(y)$. Each $\sigma^i(\cdot)$, making use of $V^i(\sigma(\cdot)) = (I - \delta H(\sigma(\cdot)))^{-1}(1-\delta)U^i$, induces a $2 \times 2$ game played in the initial period. The algorithm computes the symmetric equilibria of this game, in which indifference is relaxed to mean "payoffs within $\epsilon$." If, for all four values of $y$, $\sigma^i(y)$ and $\sigma^{-i}(y)$ are equilibria, then $\sigma(\cdot)$ is accepted as an approximate equilibrium. The best first period solution, $\sigma_0$, and welfare are then determined as the symmetric welfare-maximizing equilibrium of the $2 \times 2$ game.

The value of epsilon used for the base runs was $\epsilon = .002$. This is to account for the possibility of round-off error, and is chosen to be slightly below what may reasonably be imagined to be the threshold of human perception.

The initial run used $K = 20$. As $\sigma^i(CC)$ was always in $\{0, 1\}$, this restriction was imposed for the primary run with $K = 80$. Recall that in a Harsanyi purification, individuals are restricted to pure strategies and mixing takes place only because different individuals use different pure strategies. Since the number of participants per session was less than 30, using a grid of $K = 80$ is greater precision than can be obtained with a Harsanyi purification.

I discuss the stability of this numerical procedure in Appendix I as well as alternative procedures. There I consider alternative values of $\epsilon \in \{.001, .004, .008\}$ and $K = \{10, 20\}$. The short version is that a more relaxed criterion for equilibrium leads

the algorithm to find a spurious high welfare equilibrium for Bruttel and Kamecke (2012), but otherwise the welfare results are robust. The lack of robustness for Bruttel and Kamecke (2012) is consistent with the results on bifurcation discussed below.

On a mini PC with 8GB RAM and Intel Alder Lake-N95 3.4 GH processor running Ubuntu 22.04, searching grids of $K = 20$ took about thirty seconds to solve for each treatment, while the restricted grid of $K = 80$ took about six minutes.

## 3. The Data

I started with all the data in the meta-studies of Dal Bó and Fréchette (2018) and Fudenberg and Rehbinder (2024). As in Fudenberg and Rehbinder (2024) I excluded treatments with finite endings. As I am interested in long-run play, I selected only experiments in which the sessions lasted more than fifteen matches. I focus on data starting from the tenth match. This window was not chosen based on the data here, it was chosen for consistency with earlier work, and is the same window used in Levine (2025). All data is reported this way unless otherwise indicated. To check the robustness of this tenth match forward window, in section 7 below, using the same data, I report snapshots of the 8th and 15th periods.

Twenty-five treatments from ten papers remained: the ten papers are listed in Table 3.1 below. As indicated in the Table: some of the studies replicated treatments originally done in Dal Bó and Fréchette (2011). In order to fit the tables, I have abbreviated the references to the base (non-replication) studies, with the first three letters of each author's name.

Table 3.1: Studies

| abbreviation | replication | citation |
|---|---|---|
| DalFre | no | Dal Bó and Fréchette (2011) |
| DreEtAl | no | Dreber et al (2008) |
| BruKam | no | Bruttel and Kamecke (2012) |
| SheTarSai | no | Sherstyuk, Tarui and Saijo (2013) |
| KagSch | no | Kagel and Schley (2013) |
| AreCouRan | no | Arechar, Kouchaki and Rand (2018) |
| | yes | Dal Bó and Fréchette (2015) |
| | yes | Dal Bó and Fréchette (2019) |
| | yes | Proto, Rustichini and Sofianos (2019) |
| | yes | Ghidoni and Suetens (2022) |

replication indicates if study replicates treatments in Dal Bó and Fréchette (2011)

In these studies there are nine normalized payoff matrices shown below in Table 3.2.

Table 3.2: Payoff Matrices: player 1 payoffs

| payoffs | $u^1(CC)$ | $u^1(CD)$ | $u^1(DC)$ | $u^1(DD)$ |
|---|---|---|---|---|
| DalFre:A | 1.00 | −1.86 | 3.57 | 0.00 |
| DalFre:B | 1.00 | −0.87 | 1.67 | 0.00 |
| DalFre:C | 1.00 | −0.57 | 1.09 | 0.00 |
| DreEtAl:A | 1.00 | −2.00 | 3.00 | 0.00 |
| DreEtAl:B | 1.00 | −1.00 | 2.00 | 0.00 |
| BruKam | 1.00 | −0.83 | 2.17 | 0.00 |
| SheTarSai | 1.00 | −0.25 | 2.00 | 0.00 |
| KagSch | 1.00 | −0.50 | 2.00 | 0.00 |
| AreCouRan | 1.00 | −0.33 | 1.33 | 0.00 |

C for cooperate, D for defect

These payoff matrices are combined with different discount factors and give rise to twelve distinct treatments shown below in Table 3.3. Treatment tables, including this one, are always ordered first by predicted welfare, given below in Table 4.1, then by empirical welfare. The column CD-DC is the welfare from the off-diagonal, where one player cooperates and one defects. I always aggregate statistics by combining all experiments which are replications of each other.

Table 3.3: Treatments: by payoffs and $\delta$

| payoffs | $\delta$ | CD-DC welfare | welfare $\overline{\omega}$ | cooperation | participants | matches |
|---|---|---|---|---|---|---|
| DalFre:B | 0.50 | 0.41 | 0.15 | 0.18 | 80 | 61 |
| DreEtAl:A | 0.75 | 0.50 | 0.14 | 0.14 | 28 | 21 |
| AreCouRan | 0.13 | 0.50 | 0.19 | 0.19 | 66 | 21 |
| BruKam | 0.80 | 0.67 | 0.35 | 0.32 | 36 | 20 |
| DalFre:A | 0.50 | 0.86 | 0.14 | 0.10 | 196 | 42 |
| DalFre:A | 0.75 | 0.86 | 0.25 | 0.19 | 146 | 33 |
| DalFre:C | 0.50 | 0.26 | 0.39 | 0.46 | 306 | 43 |
| DreEtAl:B | 0.75 | 0.50 | 0.46 | 0.46 | 22 | 27 |
| KagSch | 0.75 | 0.75 | 0.58 | 0.53 | 114 | 30 |
| SheTarSai | 0.75 | 0.88 | 0.69 | 0.59 | 56 | 26 |
| DalFre:B | 0.75 | 0.41 | 0.64 | 0.65 | 68 | 42 |
| DalFre:C | 0.75 | 0.26 | 0.69 | 0.72 | 260 | 44 |

welfare $\overline{\omega}$: average welfare in the data
cooperation: fraction of times cooperation is chosen in the data
CD-DC: average payoff when one cooperates and one defects
matches: average number of matches in the data

### *3.1. Cooperation vs. Welfare*

Welfare is predicted by the theory and has a strong economic meaning, so is a natural measure of "what happened in a treatment." By contrast earlier work has focused on cooperation rates. Of course, the reason that cooperation is interesting in the prisoner's dilemma is precisely because the dilemma is that mutual cooperation welfare dominates the dominant strategy equilibrium. That is, the dilemma is about welfare and cooperation is correlated with welfare.

Cooperation rates do directly measure behavior. Hence, it is sensible to ask how welfare differs from the cooperation rate. First, if players play only $CC$ or $DD$ then the two are the same: because payoffs are normalized, welfare is equal to the cooperation rate. Second, if off-diagonal welfare is one half, as it is in Dreber et al (2008) and Arechar, Kouchaki and Rand (2018) then the two are the same. All of the off-diagonal welfares are reported in Table 3.3 above.

The cooperation rate is an imperfect measure of "what happened in a treatment" because it does not account for correlation. It views a $50-50$ alternation between $CC$ and $DD$ as exactly the same as a $50-50$ alternation between $CD$ and $DC$. From the point of view of behavior, those two alternations are quite different. In a game

like DalFre:C where off-diagonal welfare is 0.26 alternating between one cooperating and one defecting is not terribly good. By contrast, in a game SheTarSai where off-diagonal welfare is 0.88 alternating between one cooperating and one defecting is not so bad.

While, from an economic point of view welfare is meaningful, and so is arguably a better measure, the bottom line is that, as a practical matter, it does not make much difference: the two measures are highly correlated. Using the data in Table 3.3 I regressed cooperation rates on welfare. The results in Table 3.4 below show how highly correlated the two measures are.

Table 3.4: Regression of cooperation on welfare $\overline{\omega}$

| | intercept | slope |
|---|---|---|
| coefficient | $-0.02$ | 1.01 |
| se | 0.03 | 0.07 |

R-squared is 0.95

Since the theory here predicts welfare, and as it has strong economic meaning, I will focus on that. For purposes of robustness, in section 4.5 below, I show that comparing the theoretical and empirical cooperation rates lead to the same picture as welfare comparisons.

## 4. Theory vs. Data

The results of this study are summarized in Table 4.1 below. It shows the twelve treatments, together with the welfare prediction of noise theory $\hat{\omega}$ versus the empirical welfare $\overline{\omega}$. Also shown for comparative purposes are the minimum and maximum welfare achievable with noise. The italics are cases where the theory says that mutual defection is the only equilibrium. There are two substantial anomalies in the data, both highlighted in bold: DalFre:C;$\delta = 0.50$ and DreEtAl:B. The latter is based on little data - only 22 participants - but the former has the most data of any treatment - 306 participants - so is unlikely due to sampling error in measuring $\overline{\omega}$. These anomalies will be discussed in greater detail subsequently.

Table 4.1: Welfare: theory versus data

| payoff matrix | $\delta$ | theory $\hat{\omega}$ | data $\overline{\omega}$ | | min | max | | participants |
|---|---|---|---|---|---|---|---|---|
| *DalFre:B* | 0.50 | 0.14 | 0.15 | | 0.14 | 0.81 | | 80 |
| *DreEtAl:A* | 0.75 | 0.17 | 0.14 | | 0.17 | 0.83 | | 28 |
| *AreCouRan* | 0.13 | 0.17 | 0.19 | | 0.17 | 0.83 | | 66 |
| *BruKam* | 0.80 | 0.21 | 0.35 | | 0.21 | 0.88 | | 36 |
| *DalFre:A* | 0.50 | 0.27 | 0.14 | | 0.27 | 0.93 | | 197 |
| *DalFre:A* | 0.75 | 0.27 | 0.25 | | 0.27 | 0.93 | | 146 |
| DalFre:C | 0.50 | **0.67** | 0.39 | | 0.10 | 0.77 | | 306 |
| DreEtAl:B | 0.75 | **0.67** | 0.46 | | 0.17 | 0.83 | | 22 |
| KagSch | 0.75 | 0.67 | 0.58 | | 0.24 | 0.90 | | 114 |
| SheTarSai | 0.75 | 0.67 | 0.69 | | 0.27 | 0.94 | | 56 |
| DalFre:B | 0.75 | 0.69 | 0.64 | | 0.14 | 0.81 | | 68 |
| DalFre:C | 0.75 | 0.74 | 0.69 | | 0.10 | 0.77 | | 260 |

min/max: welfare of mutual defection/cooperation with noise
italics: theory $\hat{\omega}$ equal to minimum
bold: error $|\hat{\omega} - \overline{\omega}|$ of more than 0.15

The first take-away from the table is that supporting cooperation can be costly. Call the treatments where mutual-defection is not the only equilibrium *cooperative*. From Table 4.4 below, these equilibrium are all supported by a strategy that cooperates in the first period, and whenever the outcome last period was $CC$. Never-the-less for SheTarSai and KagSch the maximum welfare (max) that can be achieved by mutual cooperation with noise is at least 0.23 greater than the best equilibrium predicted by the theory. This means that there is a substantial cost from punishing noisy opponents.

In Appendix III I present information about variation in welfare between sessions of a treatment. It would be concerning if welfare varied between very low and very high, with the differences between treatments depending on the frequency of very high-welfare sessions. This is not what the theory predicts. The take-away from Appendix III is that while there is between-session variation, in the low-welfare treatments, all the sessions have relatively low welfare, and in the high-welfare treatments all the sessions have relatively high welfare.

I turn now to a more detailed assessment of the results in Table 4.1.

*4.1. Qualitative Analysis.*

I first examine how well the theory explains the data from a qualitative point of view. First, in the six italicized treatments where mutual defection is the only equilibrium, in two treatments welfare predictions are off by $0.13 - 0.14$ while in the remaining four cases welfare in the data is within 0.03 of what the theory predicts. The highest empirical welfare in these mutual defection cases is 0.35.

By contrast, in the cooperative cases welfare in the data is at least 0.39 and exceeds the welfare of mutual-defection with noise by at least 0.29. That is, the treatments that are predicted to have mutual defection have less cooperation (as measured by welfare) than those that are predicted to be cooperative. The theory does a good job sorting out the cooperative treatments.

The theory also sorts well based on cooperation rates. The italicized treatments all have a predicted cooperation rate of 0.17 due to the assumed 1/6 probability of trembling onto an unintended action. The empirical cooperation rates for five of the treatments range from 0.10 to 0.19 with the sixth being equal to 0.32. By contrast, the empirical cooperation rates in the cooperative treatments are all at least 0.46.

*4.2. Quantitative Analysis*

I turn now to the quantitative fit of the theory. In section 5 below I will compare it to other theories, but in an absolute sense, how big is big? Should a difference between the theory and data of 0.05 as is the case in DalFre:B;$\delta = 0.75$ be regarded as large or small? There is no unique answer to this question, but insight can be gained by looking at past practice.

Below in Table 4.2 I have gathered together data from games where Nash equilibrium uniquely predicts an equilibrium path of mutual defection. These are the finite horizon games from Dal Bó (2005) together with the two treatments in this paper, DalFre:A;$\delta = 0.5$ and AreCouRan. My reading of the literature is that mutual defection is regarded as a good description of what happens in those games. However, as reported in the second column of Table 4.2 absolute welfare errors are small but not zero.

Table 4.2: Normalized Welfare Errors for Nash equilibrium paths

| game | abs error |
|---|---|
| Dal:PD1;$H = 1$ | 0.13 |
| Dal:PD2;$H = 1$ | 0.06 |
| Dal:PD1;$H = 2$ | 0.12 |
| Dal:PD2;$H = 2$ | 0.13 |
| Dal:PD1;$H = 4$ | 0.02 |
| Dal:PD2;$H = 4$ | 0.24 |
| DalFre:A;$\delta = 0.5$ | 0.15 |
| AreCouRan | 0.19 |

finite games: average over final two matches, $H$ is number of supergame periods
DalFre and AreCouRan: average is over the 10th and later matches from Table 4.1

I take this as an indication of what is an acceptable error for a theory. Below in Figure 4.1 below I compare the histogram from Table 4.2 with the histogram from Table 4.1. I want to emphasize that these errors are by no means random variables. Rather, a histogram is a convenient way of summarizing what have been considered acceptable errors historically, and what sorts of errors I have found.

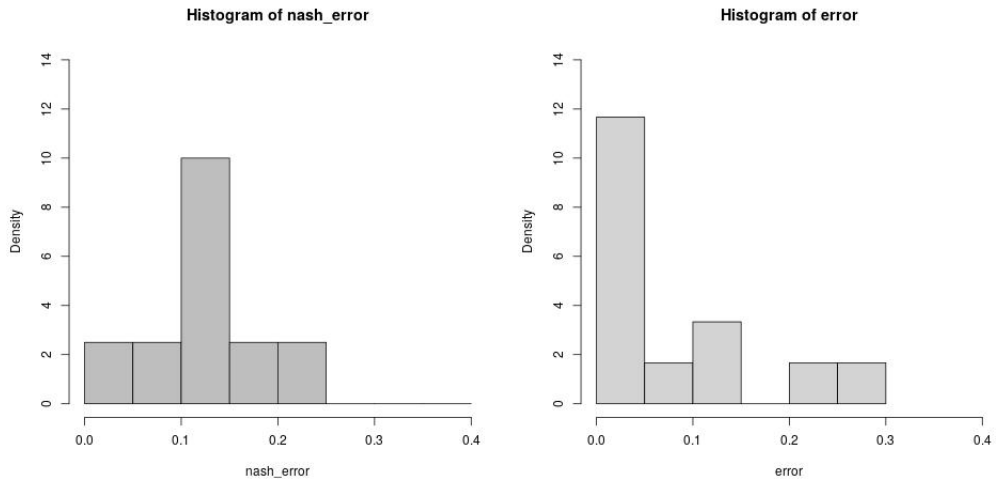


Figure 4.1: Error Histograms
left are the Nash welfare errors from Table 4.2
right are the welfare errors for noise theory from Table 4.1
the vertical axis is the continuous density, so:
since the width of the cells is 0.05 the columns sum to 20

There are two takeaways from this. First, in general, the quantitative errors in

noise theory are small compared to what has been viewed as acceptable for subgame perfect equilibrium. Second, as highlighted in Table 4.1, the two bold-faced observations, DalFre:C;$\delta = 0.50$ and DreEtAl:B, are indeed anomalous, with errors larger than what would ordinarily be considered "acceptable."

Below in Figure 4.2 I show a graphical view of the theory against the data. The horizontal axis is the theoretical prediction of welfare and the vertical axis the empirical welfare.

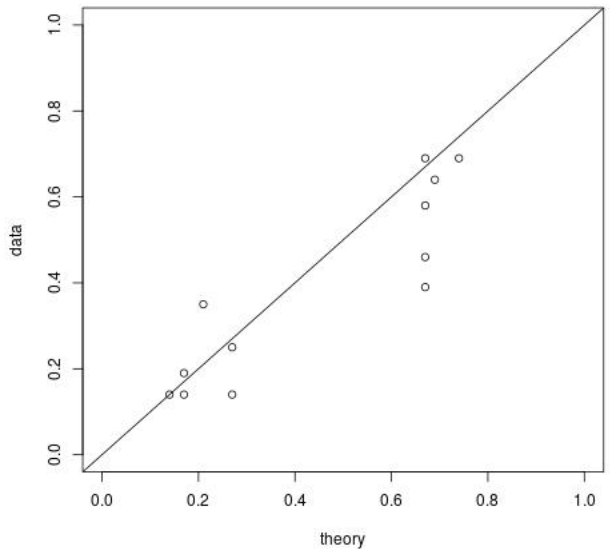


Figure 4.2: Scatter Plot of Welfare: Data $\overline{\omega}$ Against Theory $\hat{\omega}$

If the theory was perfectly correct, then the dots representing observations in Figure 4.2 would lie on the 45° line. With the exception of the two anomalous outliers DalFre:C;$\delta = 0.50$ and DreEtAl:B they are generally close.

*4.3. Sampling Error vs. Specification Error*

The difference between the theory and the data as reported in Table 4.1 has two sources: sampling error and specification error. In the case of Nash equilibrium, the theory is deterministic: there is no sampling error, and so the error is entirely specification error. By contrast noise theory is stochastic due both to trembling and the use of mixed strategies. Standard errors can be computed by Monte Carlo simulation, but there is an easy upper bound that is fairly accurate.

Welfare is between zero and one. The greatest variance possible for any such random variable is 1/4 attained by the binomial with mean equal to 1/2. According to the theory, trembles and mixing are independent between periods, so each match is an independent draw. Hence, for each treatment, the number of observations is half the number of participants times the average number of matches after the ninth. These can be computed from Table 3.3 above. The least number of observations is DreEtAl:A with 168. This gives rise to upper bounds on the standard error of 0.039. Even the largest of these upper bounds, 0.039, is inadequate to explain the anomalous treatments. The smallest of these divergences of data from theory is 0.21 in DreEtAl:B. This is more than five times the uppermost possible standard error.

In the data, individual play is serially correlated. However, the first matches, at least, are independent, so there are at least as many independent observations as the half the number of participants. This gives rise to a larger upper bound on the standard errors, reported below in Table 4.3 below.

Table 4.3: Standard Error Upper Bounds

| payoff matrix | $\delta$ | theory $\hat{\omega}$ | data $\overline{\omega}$ | max se | participants |
|---|---|---|---|---|---|
| *DalFre:B* | 0.50 | 0.14 | 0.15 | 0.07 | 80 |
| *DreEtAl:A* | 0.75 | 0.17 | 0.14 | 0.13 | 28 |
| *AreCouRan* | 0.13 | 0.17 | 0.19 | 0.09 | 66 |
| *BruKam* | 0.80 | 0.21 | 0.35 | 0.12 | 36 |
| *DalFre:A* | 0.50 | 0.27 | 0.14 | **0.05** | 197 |
| *DalFre:A* | 0.75 | 0.27 | 0.25 | 0.06 | 146 |
| DalFre:C | 0.50 | **0.67** | 0.39 | **0.04** | 306 |
| DreEtAl:B | 0.75 | **0.67** | 0.46 | **0.15** | 22 |
| KagSch | 0.75 | 0.67 | 0.58 | **0.07** | 114 |
| SheTarSai | 0.75 | 0.67 | 0.69 | 0.09 | 56 |
| DalFre:B | 0.75 | 0.69 | 0.64 | 0.09 | 68 |
| DalFre:C | 0.75 | 0.74 | 0.69 | 0.04 | 260 |

max se: upper bound on standard error

bold: bigger than 1.645 for non-cooperative equilibria, 1.282 for cooperative equilibria

Reported in bold face are differences between the theory and data that are "significantly different at the 90% level." For the non-cooperative treatments, I used a two-tailed test. For the cooperative treatments, because the theory consistently over-

states welfare,[3] I use a one-tailed test. The results show that, even with using an extreme upper bound for standard errors, I cannot "excuse" the failures of the theory as being due to sampling error.

This is all to say that it is specification error, not sampling error, that is important for the anomalies.

### *4.4. Best Equilibrium Strategies*

Below in Table 4.4 are shown the intended strategies that support the maximum welfare reported in Table 4.1 above. Recall that this is a $36 \times 36$ game. Levine (2026) shows that for any finite game, the solution of the welfare maximization problem is generically unique. This does not "prove" that for any particular problem the solution is unique, but the algorithm found only a single optimal strategy in each case.

Table 4.4: Best Equilibrium Intended Strategies

| payoffs | $\delta$ | initial | CC | CD | DC | DD |
|---|---|---|---|---|---|---|
| *DalFre:B* | 0.50 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| *DreEtAl:A* | 0.75 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| *AreCouRan* | 0.13 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| *BruKam* | 0.80 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| *DalFre:A* | 0.50 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| *DalFre:A* | 0.75 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| DalFre:C | 0.50 | 1.00 | 1.00 | 0.00 | 0.00 | 1.00 |
| DreEtAl_08:B | 0.75 | 1.00 | 1.00 | 0.01 | 1.00 | 0.00 |
| KagSch | 0.75 | 1.00 | 1.00 | 0.00 | 0.53 | 0.01 |
| SheTarSai | 0.75 | 1.00 | 1.00 | 0.00 | 0.28 | 0.01 |
| DalFre:B | 0.75 | 1.00 | 1.00 | 0.30 | 1.00 | 0.08 |
| DalFre:C | 0.75 | 1.00 | 1.00 | 0.91 | 1.00 | 0.26 |

player 1 conditional probabilities of cooperating
CC,CD,DC,DD are last the period outcome

The strategy $(0, 0, 0, 0, 0)$ is mutual-defection and is the only available equilibrium in the italicized treatments. The strategy $(1, 1, 0, 0, 0)$ is the grim-trigger strategy, which is never used. Mixed strategies are rare. All of the cooperative strategies cooperate initially and when the state is $CC$. They have the general form $(1, 1, M, H, L)$

[3] With the exception of SheTarSai, where there is practically no difference.

where $M < H$. This can be described as grim-trigger with forgiveness ($L$) and apology ($H > M$). When one player defects and the other cooperates the defecting player "apologizes" by cooperating with a higher probability than the other.

Each of the memory-one strategies gives rise to a $2 \times 2$ game with payoffs

$$(1-\delta)U^i(a) + \delta \sum_{y \in A} \pi(y|a)V^i(y|\sigma).$$

The best equilibrium memory-one strategies give rise to three types of games. The first are prisoner's dilemma games in which it is dominant to defect. Here the only equilibrium of the supergame is mutual defection. The second are coordination games in which $CC$ and $DD$ are equilibria with a mixed one in between (which, however, is never optimal to use). Here the equilibrium can be chosen in a state-dependent way to provide incentives for cooperation. The third are indifference games where the payoff of a player is independent of the action of the other player, and the player is indifferent between cooperating and defecting. In these latter games, mixing sometimes takes place, and, in an off-diagonal states $y \in \{CD, DC\}$, the mixing need not be symmetric and generally is not.

The issue of strategies is further pursued in Appendix IV, which examines empirical and conditional cooperation rates. The bottom line is that the predictions of noise theory about conditional cooperation are mediocre, but considerably better than alternative theories.

### *4.5. Cooperation Rates*

The uniqueness of the equilibrium strategies makes it possible to compute the cooperation rates for each solution. This is shown in Table 4.5 below.

Table 4.5: Theoretical and Empirical Cooperation Rates

| payoffs | $\delta$ | theory cooperation | cooperation |
|---|---|---|---|
| *DalFre:B* | 0.50 | 0.17 | 0.18 |
| *DreEtAl:A* | 0.75 | 0.17 | 0.14 |
| *AreCouRan* | 0.13 | 0.17 | 0.19 |
| *BruKam* | 0.80 | 0.17 | 0.32 |
| *DalFre:A* | 0.50 | 0.17 | 0.10 |
| *DalFre:A* | 0.75 | 0.17 | 0.19 |
| DalFre:C | 0.50 | **0.74** | 0.46 |
| DreEtAl:B | 0.75 | **0.67** | 0.46 |
| KagSch | 0.75 | 0.59 | 0.53 |
| SheTarSai | 0.75 | 0.56 | 0.59 |
| DalFre:B | 0.75 | 0.73 | 0.65 |
| DalFre:C | 0.75 | 0.81 | 0.72 |

The picture is essentially the same as for welfare: BruKam is marginal, under predicting cooperation as it does welfare, and DalFre:C and DreEtAl:B are anomalous, over predicting cooperation as they do welfare. The fit in the remaining treatments remains good.

## 5. Learning Theoretic Approach

In the last twenty years there has been extensive research into when cooperation emerges in repeated prisoner's dilemma games in the laboratory. There are four main theories: subgame perfect equilibrium, risk dominance, the critical discount factor, and reinforcement learning. With the exception of reinforcement learning, none explicitly account for the presence of trembling, although there is a different literature that examines perfect public equilibrium where the experimenter introduces noise in observation: a good review of the entire literature can be found in Dal Bó and Fréchette (2018).

The first approach to the issue of cooperation is that of subgame perfect equilibrium studied in Dal Bó (2005). He shows that finite games have low cooperation rates, while indefinitely repeated games with cooperative equilibria have high cooperation rates. Subsequent research has borne out the idea that if mutual defection is the only equilibrium, then cooperation rates are low. This can be seen in Table 4.2 above, although, in fact, welfare ranges from 0.02 to 0.24, while the theory predicts 0.00.

Subsequent research by Dal Bó and Fréchette (2011) (the results of which are included here) showed that the reverse need not be true. Even if cooperation is an equilibrium, cooperation rates can be low. This can be seen in Table 4.1 above, while DalFre:B;$\delta = 0.50$, DalFre:A;$\delta = 0.75$ and DreEtAl:A all admit cooperative equilibria (without noise), all have low cooperation rates - welfare is no higher than 0.25 and cooperation rates no more than 0.19.

The anomaly of games in which there are cooperative equilibria, yet cooperation does not emerge, led to learning theoretic considerations. Theories of learning and evolution such as Kandori, Mailath and Rob (1993) and Young (1993) suggest that in a $2 \times 2$ game risk dominance plays a key role in determining which equilibrium will emerge when there are multiple equilibria. To make this operational in an indefinitely repeated game where there are 32 memory-one strategies the literature has focused on just two such strategies: mutual-defection and grim-trigger. Before discussing how this works and the findings, I want to point out that this choice incorporates the basic hypothesis of this paper - that players are trying to achieve a good outcome. Specifically, if the opponent is thought to randomize equally over all 32 memory-one strategies, for all the treatments here, always-defect is a best response. By limiting attention to grim-trigger, the question becomes: if players are trying to find a good solution using grim-trigger, will they be able to learn to do so?

Dal Bó and Fréchette (2011) examine risk dominance in the first six games of Table 3.3. They give mixed reviews. Their assessment is based on whether cooperation rose or fell over time, and they do not have a clear theory of what constitutes high and low cooperation - standard theories being limited to 0.00 and 1.00. Moreover, rising and falling is a difficult criterion because initial match cooperation ranges from 28% to 56%, and doing well in the early matches may mean doing less well in the later matches.

Before looking at the data, I first introduce two key theoretical constructs connected with risk dominance. The first is the basin: the greatest proportion of players who can be playing always defect for grim-trigger to be a best response. However, when grim-trigger is not an equilibrium, this is not well-defined, as grim-trigger will not be a best response, regardless of how few players are playing always defect. By contrast, the the basin of always defect is always well-defined, so I define the *modified basin,* $\beta^M$,of grim-trigger to be one minus the basin of always defect. When grim-trigger is an equilibrium this is the same as the basin and is given by

$$\beta^M = \frac{(1-\delta)(u^1(DC)-1)-\delta}{(1-\delta)(u^1(DC)+u^1(CD)-1)-\delta}.$$

If grim-trigger is not an equilibrium then the modified basin is zero, that is, $\beta^M = 0$.If the modified basin is greater than 0.50 then grim-trigger is risk-dominant; if it is smaller, then mutual defection is risk-dominant.

A related idea is that of the critical discount factor. This was originally introduced by Blonski, Ockenfels and Spagnolo (2011) based on axiomatic ideas, the crucial one being that for any given payoff matrix there is a critical discount factor above which there will be cooperation and below which there will not. They have additional axioms concerning additive separability and so forth, but the bottom line is that they compute

$$\Delta = \delta - \frac{u^1(DC)-u^1(CD)-1}{u^1(DC)-u^1(CD)}$$

where the negative term is the critical discount factor. In other words, there should be cooperation if $\Delta$ is positive and should not be if it is negative. Blonski and Spagnolo (2015) subsequently pointed out that the critical discount factor is the discount factor that makes players indifferent between grim-trigger and always-defect when the population is $50-50$ between the two, so this has a similar flavor to the basin.

Finally, in a recent paper, Fudenberg and Rehbinder (2024) consider a six-parameter reinforcement learning model using $\Delta$ as an explanatory variable. They undertake the ambitious program of predicting play in every round of every match. The details of that model are well described in the paper. For the purposes here I generated artificial data using the parameters from their Table 4. For each treatment in Table 3.3 I conducted a Monte Carlo simulation with 1000 trials. I used the greatest even number less than the mean number of participants per session (14) and the mean number of matches for the treatment from Table 3.3. Averaging welfare over the tenth and subsequent matches generated predictions $\hat{\bar{\omega}}$ of welfare for each treatment.

The values of $\beta^M, \Delta, \hat{\omega}$ and $\hat{\bar{\omega}}$ for the different treatments are below in Table 5.1 below.

Table 5.1: Models and Values

| payoff matrix | $\delta$ | $\hat{\omega}$ | $\beta^M$ | $\Delta$ | $\hat{\hat{\omega}}$ | data $\overline{\omega}$ |
|---|---|---|---|---|---|---|
| *DalFre:B* | 0.50 | 0.14 | 0.28 | −0.11 | 0.22 | 0.15 |
| *DreEtAl:A* | 0.75 | 0.17 | 0.33 | −0.05 | **0.30** | 0.14 |
| *AreCouRan* | 0.13 | 0.17 | **0.00** | **−0.27** | 0.24 | 0.19 |
| *BruKam* | 0.80 | 0.21 | **0.77** | **0.13** | **0.58** | 0.35 |
| *DalFre:A* | 0.50 | 0.27 | 0.00 | −0.32 | 0.12 | 0.14 |
| *DalFre:A* | 0.75 | 0.27 | **0.19** | **−0.07** | 0.31 | 0.25 |
| DalFre:C | 0.50 | **0.67** | **0.61** | **0.10** | 0.47 | 0.39 |
| DreEtAl:B | 0.75 | **0.67** | **0.67** | **0.08** | 0.50 | 0.46 |
| KagSch | 0.75 | 0.67 | **0.80** | **0.15** | 0.60 | 0.58 |
| SheTarSai | 0.75 | 0.67 | **0.89** | **0.19** | 0.67 | 0.69 |
| DalFre:B | 0.75 | 0.69 | **0.73** | **0.14** | 0.57 | 0.64 |
| DalFre:C | 0.75 | 0.74 | **0.84** | **0.35** | 0.79 | 0.69 |

italics: $\hat{\omega}$ equal to minimum with noise
bold: error of more than 0.15

One thing should be clear: the basin and the critical discount factor are non-starters for quantitative predictions: no theory that predicts that welfare is either zero or one is quantitatively useful. In particular, the highest welfare in any session (see Appendix III) is 0.82, and no treatment higher than 0.69. Moreover, in only three treatments does welfare fall below 0.16, and the lowest welfare in any treatment is 0.14.

I turn now to a more detailed analysis how $\beta^M, \Delta$, $\hat{\omega}$ and $\hat{\hat{\omega}}$ fare with respect to the data.

### *5.1. Sorting Cooperation Levels*

As indicated those treatments predicted to have mutual defection by $\hat{\omega}$ have strictly lower welfare than those that are cooperative. How do $\beta^M, \Delta$ do on this score? The $\beta$ theory predicts mutual defection when $\beta^M < 0.5$ and $\Delta$ theory when $\Delta < 0$. These generate the same predictions about mutual defection as $\hat{\omega}$ except that $\hat{\omega}$ predicts mutual defection for BruKam and the other two do not. As welfare in BruKam is higher than any of the other italicized treatments, this means that while $\beta^M, \Delta$ disagree with $\hat{\omega}$ about where the dividing line is, they also correctly yield lower welfare in the treatments they predict to have mutual defection than those they predict to have cooperation.

*5.2. Bifurcation*

As just indicated, all four theories do a good job sorting the data into high and low welfare treatments. It is a characteristic of the static theories - $\hat{\omega}, \beta^M, \Delta$ - that there is a bifurcation. That is, as the parameters are gradually changed from a game in which the only equilibrium is mutual defection, welfare should jump when cooperative equilibria emerge. In other words, the prediction is that a histogram of welfare should be U-shaped, exhibiting low and high welfare, but nothing in between. Below in Figure 5.1 is the histogram of empirical welfare.

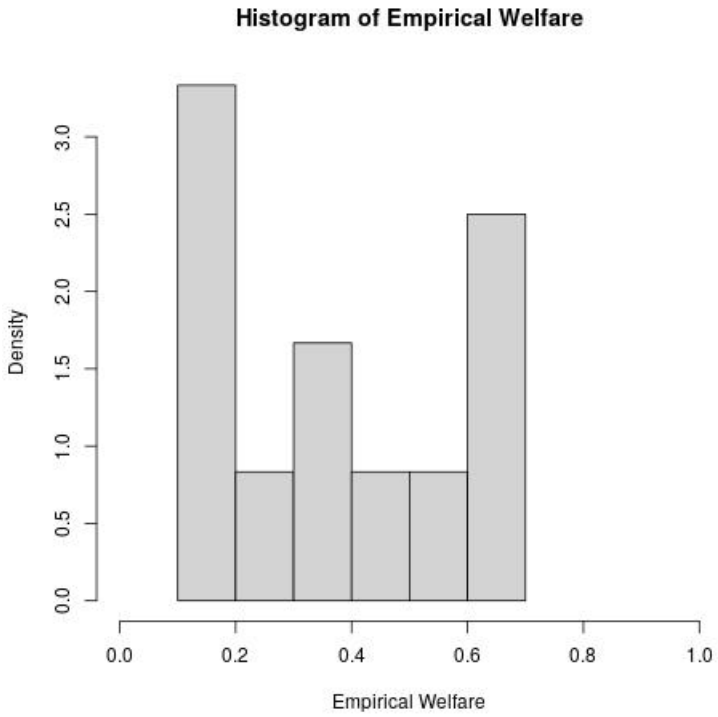


Figure 5.1: Histogram of Empirical Welfare

The first thing to observe is that the histogram is compressed: the low welfare treatments give welfare well greater than zero and the high-welfare treatments are even further away from one. Consequently, the $\beta^M, \Delta$ theories in their purest form cannot contend with the data quantitatively. They predict welfare of 0 when they predict mutual defection, while it ranges from 0.10 to 0.25 in the data. They predict grim-trigger with welfare of 1.00 in the remaining cases, while empirical welfare in those cases ranges from 0.35 to 0.69.

The second observation is that while the histogram is generally U-shaped, there is also a middle-range spike between 0.35 and 0.46. These correspond to the two anomalous treatments in noise theory - DalFre:C;$\delta = 0.5$ and DreEtAl:B, plus BruKam. This poses a challenge for all the theories.

*5.3. Monotonicity*

According to the relevant theories empirical welfare should be increasing in $\hat{\omega}, \beta^M, \Delta, \hat{\hat{\omega}}$. If empirical welfare is increasing in the theoretical constructs, then the theoretical constructs should be increasing in the data. I examine this below in Figure 5.2 plotting the welfare data on the horizontal axis against $\hat{\omega}$ and $\hat{\hat{\omega}}$ and values of $\beta^M$ and $\Delta$ normalized so they will conveniently fit on the graph using the regression coefficients from below.

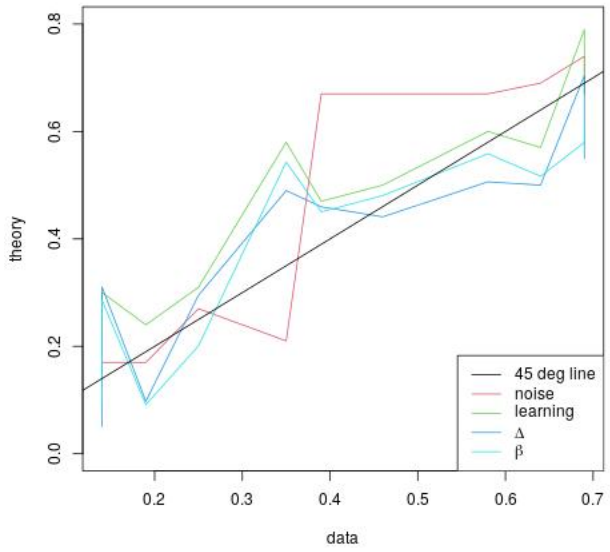


Figure 5.2: Theory Against Data

As can be seen, while all curves are generally upward sloping, none are increasing. Note also that, consistent with discussion of bifurcation above, all of the theories are quite challenged in the middle range between 0.35 and 0.46.

To investigate further, observe that two theories, $\hat{\omega}$ and $\hat{\hat{\omega}}$ make specific predictions of welfare. Unless we take the discredited theory of 0 for mutual-defection and 1 for cooperation, the theories $\beta^M, \Delta$ do not. However, we can regress the aggregated data

for each treatment on the theoretical constructs (in Table 5.3 below) to get reasonable predictions.[4] The results are shown in Table 5.2 below.

Table 5.2: Predictors of Welfare $\overline{\omega}$

| | parameters | MAE | $R^2$ | intercept (se) | slope (se) |
|---|---|---|---|---|---|
| $\hat{\omega}$ | 0 | 0.09 | 0.75 | 0.00 | 1.00 |
| $\beta^M$ | 2 | 0.09 | 0.77 | 0.09 (0.06) | 0.58 (0.10) |
| $\Delta$ | 2 | 0.09 | 0.75 | 0.36 (0.03) | 0.98 (0.18) |
| $\hat{\hat{\omega}}$ | 6 | 0.08 | 0.75 | 0.00 | 1.00 |

parameters: number of parameters estimated from repeated game data
MAE: mean absolute error

The bottom line is that there is little difference in accuracy between the theories. All have similar mean square error[5] ($R^2$) and mean absolute error. Put another way: we might think of $\beta^M, \Delta, \hat{\hat{\omega}}$ as the best we can do looking at the data. Perhaps imperfection is unavoidable because of difference in subject populations, laboratory conditions, locations, and other factors not accounted for by the theories. Noise theory, which does not look at the data at all, does just as well.

*5.4. Predicted Values*

As indicated $\hat{\omega}$ and $\hat{\hat{\omega}}$ make specific predictions about welfare in each treatment and for $\beta^M, \Delta$ we can use the fitted values from the regression in Table 5.2 as predictions. Below in Table 5.3 I report the predictions of the four theories along with the data.

[4] As the relationship might not be linear, I also ran regressions with quadratic terms, but these added little. I also ran a logistic regression that allowed an S-shaped function: again, this made little difference.

[5] $R^2$ is not "adjusted" and is defined as one minus the ratio of the sum of squared differences between the theory and the data to the sum of squared deviations of the data from its mean. For the two regression treatments, this is guaranteed to be non-negative, but in principle, it could be negative for $\hat{\omega}$ and $\hat{\hat{\omega}}$.

Table 5.3: Welfare Predictions vs. Data

| payoff matrix | $\delta$ | $\hat{\omega}$ | $\beta^M$ | $\Delta$ | $\hat{\hat{\omega}}$ | data $\overline{\omega}$ |
|---|---|---|---|---|---|---|
| *DalFre:B* | 0.50 | 0.14 | 0.25 | 0.25 | 0.22 | 0.15 |
| *DreEtAl:A* | 0.75 | 0.17 | 0.29 | **0.31** | **0.30** | 0.14 |
| *AreCouRan* | 0.13 | 0.17 | 0.09 | 0.10 | 0.24 | 0.19 |
| BruKam | 0.80 | 0.21 | **0.54** | 0.49 | **0.58** | 0.35 |
| *DalFre:A* | 0.50 | 0.27 | 0.09 | 0.05 | 0.12 | 0.14 |
| *DalFre:A* | 0.75 | 0.27 | 0.20 | 0.30 | 0.31 | 0.25 |
| DalFre:C | 0.50 | **0.67** | 0.45 | 0.46 | 0.47 | 0.39 |
| DreEtAl:B | 0.75 | **0.67** | 0.48 | 0.44 | 0.50 | 0.46 |
| KagSch | 0.75 | 0.67 | 0.56 | 0.51 | 0.60 | 0.58 |
| SheTarSai | 0.75 | 0.67 | 0.61 | 0.55 | 0.67 | 0.69 |
| DalFre:B | 0.75 | 0.69 | 0.52 | 0.50 | 0.57 | 0.64 |
| DalFre:C | 0.75 | 0.74 | 0.58 | 0.70 | 0.79 | 0.69 |

italics: $\hat{\omega}$ equal to minimum with noise
bold: error of more than 0.15

The takeaway is this. Noise theory and $\hat{\hat{\omega}}$ has the most outliers, with two major anomalies, while each of the other theories has only a single anomaly. For the five low-welfare treatments, with the exception of DalFre:A, noise theory does better than any of the other theories. Indeed: what 5.3 highlights is that noise theory does very well most of the time, but fails badly twice - for DalFre:C;$\delta = 0.50$ and DreEtAl:B. Since learning theory does better on these, it may be that some hybrid theory, combining noise and learning as complements, would be effective.

## 6. How Much Noise?

I come now to the question of the choice of $\phi$. So far, this has been taken to be 1/3, the value from a study of a completely different class of games. How closely is this pinned down by the data, and is there another choice of $\phi$ which yields substantially better results? I turn now to the question of how welfare varies with $\phi$.

The big picture is this. The game under study is a finite game with 32 strategies. Levine (1986) shows that the problem of maximizing welfare in such a game has a generically unique solution that varies continuously with parameters except on a lower-dimensional bifurcation set, where strategies and welfare can change abruptly. A key first step in examining how welfare depends on noise, is to identify where this

bifurcation set is. Following that, I report welfare for different values of $\phi$ and ask what it means.

### *6.1. Bifurcation*

Bifurcations are reported as ranges in the third column of Table 6.1 below. The top of the range has mutual-defection as the only equilibrium. The bottom of the range has a best equilibrium that is cooperative. Hence, the bifurcation takes place in the range shown. For those treatments with bifurcation points nearer $\phi = 1/3$ the ranges were computed with a higher degree of precision. As expected, DalFre:A and AreCouRan have only mutual defection as an equilibrium regardless of noise.

Table 6.1: Bifurcation Ranges and Equilibrium Structure

| matrix | $\delta$ | $\phi$ bifurcation | equilibrium welfare | | | $\overline{\omega}$ | equilibria |
|---|---|---|---|---|---|---|---|
| | | | min | mid | max | | |
| *DalFre:B* | 0.50 | $(0.1, 0.2)$ | | 0.14 | | 0.15 | 1 |
| *DreEtAl:A* | 0.75 | $(0.1, 0.2)$ | | 0.17 | | 0.14 | 1 |
| *AreCouRan* | 0.13 | none | | 0.17 | | 0.19 | 1 |
| *BruKam* | 0.80 | $(0.332, 0.333)$ | | 0.21 | | 0.35 | 1 |
| *DalFre:A* | 0.50 | none | | 0.27 | | 0.14 | 1 |
| *DalFre:A* | 0.75 | $(0.0, 0.1)$ | | 0.27 | | 0.25 | 1 |
| DalFre:C | 0.50 | $(0.5, 0.6)$ | 0.10 | 0.46 | 0.67 | 0.39 | 6 |
| DreEtAl:B | 0.75 | $(0.339, 0.340)$ | 0.17 | $[0.33, 0.67]$ | 0.67 | 0.46 | 2 |
| KagSch | 0.75 | $(0.338, 0.339)$ | 0.24 | $[0.33, 0.67]$ | 0.67 | 0.58 | 3 |
| SheTarSai | 0.75 | $(0.338, 0.339)$ | 0.27 | $[0.33, 0.67]$ | 0.67 | 0.69 | 2 |
| DalFre:B | 0.75 | $(0.35, 0.4)$ | 0.14 | $[0.29, 0.68]$ | 0.68 | 0.64 | 7 |
| DalFre:C | 0.75 | $(0.6, 0.7)$ | 0.10 | $[0.18, 0.73]$ | 0.73 | 0.69 | 52 |

equilibrium columns explained in section 6.1.1
$\phi$ bifurcation: range in which the bifurcation is located
for $\phi$ at the top mutual-defection is the only equilibrium

Focusing on the first three columns, most of the treatments are quite robust, with their bifurcation points lying well away from $\phi = 1/3$. Four of the treatments, however, DreEtAl:B, BruKam, SheTarSai and KagSch, have bifurcations close to $\phi = 1/3$.

Getting the classification wrong is extremely costly in both mean absolute and mean square error. For example, as noise is increased from $\phi = 1/3$, the SheTarSai and KagSch treatments are the first to bifurcate. This takes place in the range

$\phi \in (0.338, 0.339)$ for both. Predicted welfare drops from 0.67 in both cases to 0.27 for SheTarSai and 0.24 for KagSch, while the empirical welfares are 0.69 and 0.58 respectively. In other words, any estimation procedure will need to get the classifications of KagSch and SheTarSai as right as possible.

*6.1.1. Structure of Equilibria*

How many equilibria are there, and what are they like? The final five columns of table 6.1 address this issue in welfare space. The question is of some importance: if, for example, all equilibria yielded about the same level of welfare then saying "the best equilibrium" would be meaningless. However, this cannot be the case, since always defect (with trembling) is always the "worst" equilibrium, and when "cooperative" equilibria are available, the "worst" equilibrium yields substantially less welfare than the "cooperative" equilibria.

Recall that an equilibrium strategy has two parts, the state dependent memory one part $\sigma^i(y)$ and the initial cooperation rate $\sigma_0^i$. The former are computed on a grid, so are necessarily finite in number. The number that are found is reported in the column "equilibria." The former, $\sigma_0^i$, are either finite (at most three) in number or, in the case of belief-free equilibria, a continuum. In belief-free equilibria, in the first-period game, players are indifferent between cooperating and defecting regardless of what the other player is doing in that first period (see also Appendix IV).

To probe this further, in order to allow for mixing in the state $CC$, I used a coarse grid of $K = 20$. As an indication of the distribution of equilibrium welfare, I have reported the minimum, the maximum and a "mid" number. The mid is the median if there are no belief-free equilibria. When there are belief free equilibria, the welfare intervals corresponding to different state-dependent equilibrium strategies $\sigma^i(\cdot)$ overlap, so I merge these into a single interval, and report the bottom and top of that interval. For comparison, I also report the empirical welfare $\overline{\omega}$.

To understand how this works, bifurcation is crucial. For high $\phi$ there is a unique equilibrium, which is "always defect." This is the case for the first six rows in table 6.1. Then there is a bifurcation in which new equilibria emerge. The three treatments DreEtAl:B, KagSch, and SheTarSai all have noise $\phi = 1/3$ just below the bifurcation point. In this case cooperative equilibria emerge, but relatively few state-dependent parts $\sigma^i(\cdot)$.

As noise is further lowered it becomes substantially less than the bifurcation point,

and a wider variety of state-dependent parts arise. This is the case for the remaining three treatments, DalFre:C;$\delta = 0.50$, DalFre:B;$\delta = 0.75$, and DalFre:C;$\delta = 0.75$.

The wide range of equilibrium welfares that emerge once cooperative equilibria become available is consistent with how the folk theorem works. What the computations show is that the folk theorem largely works even with the limitation to one-period memory strategies and in the absence of a public randomizing device. It does not work completely, because there is a gap between the worst equilibrium and the second worst equilibrium, even when there are belief-free equilibria.

The theory of "best" equilibrium should not be taken to mean literally that players succeed in finding the very best: that would be a miracle. Rather, it should be taken as a benchmark: when there are many equilibria close to the best, as there are once cooperative equilibria emerge, it should be taken as an prediction that players find something "close" to the best.

### *6.2. Varying the Noise*

Below in Table 6.2 I report welfare for different noise values, along with the overall mean absolute error of the theory from the data.

Table 6.2: Welfare for different values of $\phi$

| | | $\phi$ | | | | | | |
|---|---|---|---|---|---|---|---|---|
| payoff matrix | $\delta$ | 0.26 | 0.32 | 1/3 | 0.337 | 0.34 | 0.40 | data $\overline{\omega}$ |
| *DalFre:B* | 0.50 | 0.11 | 0.13 | 0.14 | 0.14 | 0.14 | 0.17 | 0.15 |
| *DreEtAl:A* | 0.75 | 0.13 | 0.16 | 0.17 | 0.17 | 0.17 | 0.20 | 0.14 |
| *AreCouRan* | 0.13 | 0.13 | 0.16 | 0.17 | 0.17 | 0.17 | 0.20 | 0.19 |
| *BruKam* | 0.80 | **0.75** | **0.66** | 0.21 | 0.22 | 0.22 | 0.25 | 0.35 |
| *DalFre:A* | 0.50 | 0.21 | 0.26 | 0.27 | 0.27 | 0.27 | 0.31 | 0.14 |
| *DalFre:A* | 0.75 | 0.21 | 0.26 | 0.27 | 0.27 | 0.27 | 0.31 | 0.25 |
| DalFre:C | 0.50 | **0.80** | **0.76** | **0.67** | **0.67** | **0.67** | **0.63** | 0.39 |
| DreEtAl:B | 0.75 | **0.74** | **0.68** | **0.67** | **0.66** | **0.17** | **0.20** | 0.46 |
| KagSch | 0.75 | **0.77** | 0.69 | 0.67 | 0.67 | **0.24** | 0.43 | 0.58 |
| SheTarSai | 0.75 | 0.77 | 0.69 | 0.67 | 0.67 | **0.28** | **0.32** | 0.69 |
| DalFre:B | 0.75 | 0.76 | 0.70 | 0.69 | 0.69 | 0.68 | **0.17** | 0.64 |
| DalFre:C | 0.75 | 0.80 | 0.75 | 0.74 | 0.74 | 0.74 | 0.69 | 0.69 |
| increase in MAE over $\phi = 1/3$ | | 0.06 | 0.02 | 0.00 | 0.00 | 0.06 | 0.07 | |

bold: error more than 0.15
MAE: mean absolute error

As can be seen from the final row, values of $\phi$ that lie outside the range $\phi \in (0.32, 0.34)$ yield higher mean absolute error than $\phi = 1/3$.

Larger values of $\phi$ than one-third lead to bifurcations in DreEtAl:B, Kag_Sch and SheTarSai, creating two new anomalies, and increasing the size of the DreEtAl:B anomaly from 0.20 to 0.26.

Smaller values of $\phi$ have a less dramatic effect, creating one new anomaly with BruKam, but also increasing the size of the DalFre:C;$\delta = 0.50$ anomaly substantially, from 0.28 to 0.36.

I would particularly like to emphasize the mean absolute error at $\phi = 0.26$. This is 0.04 larger even than $\phi = 0.32$, and is 0.15, meaning nearly an anomaly on average, and, in fact, leads to a fourth anomaly. Many people reading this paper, or Levine (2025), for the first time, express surprises at a noise level so high as $\phi = 1/3$. Hence, I want to emphasize that $\phi = 0.26$, which some might still regard as large, does a considerably worse job explaining the data than $\phi = 1/3$.

Is there a better solution than $\phi = 1/3$? The answer is, if so, it is a minuscule improvement. Compare $\phi = 1/3$rd to $\phi = 0.337$. As can be seen, this leads to either no change or a trivial change in predicted welfare for each treatment. The conclusion that emerges from this is that there cannot be a substantially lower mean absolute error than for $\phi = 1/3$: without a bifurcation, the continuous change in welfare over the short range $\delta \in (0.32, 0.34)$ is simply not large enough to get a much better fit.

## 7. Data Robustness

The base window used for the analysis is the tenth period on. This is prespecified, and is the same as used in Levine (2025) for a different class of games. As a robustness check, and to see if the results are dependent on this choice of window, Table 7.1 gives snapshots of welfare computed only in the eight and fifteenth period, and it gives the twenty-period forward window, along with the theory, the value of $\overline{\omega}$ computed from the base ten-period forward window, and the number of matches. It provides comparisons of the mean absolute error difference between the alternative data window and the base window, and the mean absolute error between the theory and data computed in the alternative window for three values of $\phi$, the base value of 1/3 and the two surrounding values of 0.32 and 0.40. The latter is chosen specifically to see if greater noise leads to a better fit is obtained for the "early" snapshot at period eight.

Table 7.1: Data Windows

| payoff matrix | $\delta$ | theory $\hat{\omega}$ | data $\overline{\omega}$ | data 8 | data 15 | data 20+ | matches |
|---|---|---|---|---|---|---|---|
| *DalFre:B* | 0.50 | 0.14 | 0.15 | 0.25 | 0.20 | 0.14 | 61 |
| *DreEtA:A* | 0.75 | 0.17 | 0.14 | 0.32* | 0.19 | 0.11 | 21 |
| *AreCouRan* | 0.13 | 0.17 | 0.19 | 0.21 | 0.18 | 0.18 | 21 |
| *BruKam* | 0.80 | 0.21 | 0.35 | **0.43** | 0.29 | 0.36 | 20 |
| *DalFre:A* | 0.50 | 0.27 | 0.14 | 0.14 | 0.14 | **0.11** | 42 |
| *DalFre:A* | 0.75 | 0.27 | 0.25 | 0.21 | 0.25 | 0.26 | 33 |
| DalFre:C | 0.50 | 0.67 | **0.39** | **0.35** | **0.37** | **0.40** | 43 |
| DreEtAl | 0.75 | 0.67 | **0.46** | **0.50** | 0.70* | **0.47** | 27 |
| KagSch | 0.75 | 0.67 | 0.58 | **0.38*** | 0.58 | 0.57 | 30 |
| SheTarSai | 0.75 | 0.67 | 0.69 | 0.63 | 0.79 | 0.71 | 26 |
| DalFre:B | 0.75 | 0.69 | 0.64 | **0.44*** | 0.66 | 0.64 | 42 |
| DalFre:C | 0.75 | 0.74 | 0.69 | **0.57** | **0.53*** | 0.73 | 44 |
| MAE data vs data $\overline{\omega}$ | | | 0.00 | 0.09 | 0.06 | 0.02 | |
| MAE $\phi = 0.32$ | | | 0.11 | 0.18 | 0.13 | 0.11 | |
| MAE $\phi = 1/3$ | | | 0.09 | 0.16 | 0.09 | 0.09 | |
| MAE $\phi = 0.40$ | | | 0.16 | 0.17 | 0.20 | 0.17 | |

data 8: $\overline{\omega}$ 8th period only; data 15: $\overline{\omega}$ 15th period only
data 20+: $\overline{\omega}$ average 20th and later periods
matches: average number of matches in the data
bold: discrepancy between theory and data exceeds 0.15
*: difference with base window exceeds 0.15

As can be seen, the twenty-plus window is virtually identical to the ten-plus window, the snapshot at period fifteen is similar, and the snapshot at period eight is the least similar to the base window. Note that regardless of the window, $\phi = 1/3$ continues to give strictly smaller mean absolute error than either of the alternatives, $\phi = 0.32$, $\phi = 0.40$. The finding for $\phi = 0.40$ is interesting, because it indicates that the poor fit of the theory to the snapshot at period eight cannot be "explained" as "best equilibrium with more noise."

More broadly, I would generally expect that over time greater "equilibriumness" would be observed. In the non-cooperative treatments this would mean decreasing welfare over time, and in the cooperative treatments, increasing welfare. This is generally, although not always, the case.

Focusing on the $\phi = 1/3$ case, for the period fifteen snapshot, the results are not terribly different from the period ten forward analysis. The DalFre:C;$\delta = 0.50$ treatment remain robustly anomalous. The DreEtAl treatment is no longer anomalous

due to the peculiar fact that while welfare rises from 0.50 in period eight to 0.70 in period fifteen, it then falls to $0.46 - 0.47$ for the ten period forward window. On the other hand, the DalFre:C;$\delta = 0.75$ treatment becomes anomalous as welfare actually falls slightly from 0.57 to 0.53 from period eight to fifteen, but then rises to 0.69, relatively close to what is predicted by the theory. Overall, the mean absolute error remains is the same to two significant digits.

For the twenty-plus window, welfare in the DalFre:A;$\delta = 0.50$ treatment falls enough that it becomes anomalous, but the mean absolute error also remains unchanged to two significant digits.

The bottom line here is that which "long-term" is chosen, ten plus, fifteen, or twenty plus, makes little difference.

## 8. Application to Stag Hunt

A good theory ought to be able to predict behavior in more than one class of games. In previous work, in Levine (2025), I studied non-repeated games using the model described here augmented to allow for a form of altruism that is of little importance in the repeated prisoner's dilemma setting.[6] The same level of noise used here works well in that setting. Here I stick to pure noise theory and examine two-player two-action stag hunt games. This analysis is drawn from the meta-study of Dal Bó, Fréchette and Kim (2021).

### *8.1. The Game*

In stag hunt is a two-player symmetric two-action game in which the actions can be written as $S$ for stag and $H$ for hare. Below in Table 8.1 is a normalized payoff matrix for stag hunt, where $\Lambda, \lambda > 0$.

Note that for $\Lambda > 1$ this is usually called the assurance game, and one treatment, in Feltovich, Iwasaki and Oda (2012), does in fact have $\Lambda = 2$.

[6]The key point is that, according to the calibration in Levine (2025), half the players are only willing to sacrifice about half a cent to improve efficiency, and the other half nothing. When only the non-cooperative equilibria are available; this very slightly raises welfare, but Levine (2025) shows in the Appendix on the one-shot prisoner's dilemma game that this is very small relative to the increase in welfare due to noise. On the other hand, in the cooperative equilibria, any attempt by altruistic players to be kind and not punish will undermine the equilibrium. More details can be found in Levine (2024). However, the theory here rests on the assumption that players are pro-social in the sense of attempting to coordinate on a good equilibrium. Recent experiments by Kolle, Quercia and Tripodi (2025) suggest that if this fails, less good equilibria may occur.

Table 8.1: Normalized Stag Hunt Payoffs

| | $S$ | $H$ |
|---|---|---|
| $S$ | $1,1$ | $-\lambda, 1-\Lambda$ |
| $H$ | $1-\Lambda, -\lambda$ | 0,0 |

The basin of stag for this game, in the ordinary sense, can easily be computed as

$$\beta = \frac{\Lambda}{\Lambda + \lambda}.$$

Noise theory applied to this game is extremely simple. If $\beta < \phi/2 = 1/6$ then the probability of stag is $\phi/2 = 1/6$. If $\beta > \phi/2 = 1/6$ then the probability of stag is $1 - \phi/2 = 5/6$.

### *8.2. The Data*

Dal Bó, Fréchette and Kim (2021) conduct a meta-study of 8 experiments with 18 treatments of two-player two-action stag hunt games, which have strangers treatments and last at least eight periods. To this I added the two treatments of Dal Bó, Fréchette and Kim (2021) that are comparable with the others, their "one-game" treatments.[7]

### *8.3. Theory versus Data*

Figure 8.1 below is recreates Figure 3 from Dal Bó, Fréchette and Kim (2021). It shows the frequency of stag in the eighth period, plotted against the basin $\beta$. I have added to this figure the two "one-game" treatments from Dal Bó, Fréchette and Kim (2021) and I have plotted the prediction of noise theory as the solid line and that of risk dominance by the dashed line.

[7] The studies used by Dal Bó, Fréchette and Kim (2021) are Battalio, Samuelson and Van Huyck (2001), Clark, Kay and Sefton (2001), Cooper et al (1992), Dal Bó, Fréchette and Kim (2021), Dubois, Willinger and Van Nguyen (2012), Duffy and Feltovich (2002), Feltovich, Iwasaki and Oda (2012), Schmidt et al (2003), and Straub (1995).

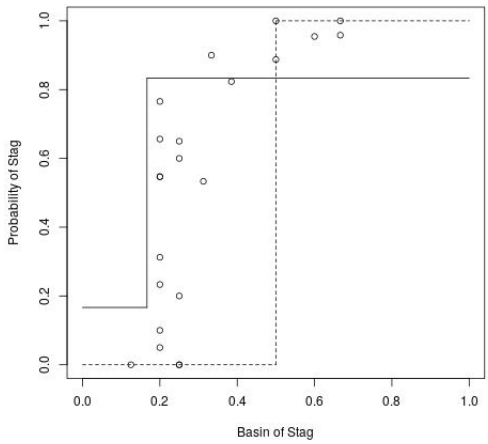


Figure 8.1: Two Player Stag Hunt: Period 8
solid line: noise theory
dashed line: risk dominance

Noise theory shows a sharp bifurcation at 1/6. The data shows a less sharp bifurcation between about 0.2 and 0.3, however, there is only one observation with $\beta < 0.2$ so we cannot be sure what happens in that range. Nevertheless the same simple theory that does a good job in indefinitely repeated prisoner's dilemma games seems to deliver in stage hunt games as well. In particular, while near the bifurcation point the data is all over the map, for the range above 0.3 the theory predicts about 83% stag while the average in the data is about 90%. This should be contrasted with risk dominance: in the indefinitely repeated prisoner's dilemma, high cooperation rates were seen only with large basins on the order of 2/3. In stag hunt a small basin of 1/3 seems to be enough to achieve a good outcome. By contrast the simple noise theory seems to be pretty close in both cases.

Why does the indefinitely repeated prisoners' dilemma require large basins for cooperation, while stag hunt requires small basins? Noise theory provides an answer. In stag hunt there is a single effect of noise: it makes hare more attractive because hare is more frequently played when players intend to play stag. Similarly, in the indefinitely repeated prisoners' dilemma, noise makes defection more attractive as it is more frequently played when players intend to cooperate. However, in the indefinitely repeated prisoner's dilemma, there is a second effect. More noise means more forgiveness: defection is not always punished. This further reduces the attractiveness of cooperating, and so payoffs must more strongly favor cooperation (larger basin)

than in stag hunt.

## 9. Conclusion

I have examined a simple theory of noisy play in which there is a 1/6 probability of a participant choosing an unintended action. This number is not estimated from repeated game data, but was used successfully in an earlier study of non-repeated games. I hypothesize that in the resulting game of imperfect public information welfare is the highest achievable in perfect public equilibrium with memory-one strategies.

With this theory, I examine predicted and actual welfare starting from the tenth match in twelve different treatments - all of those from Fudenberg and Rehbinder (2024) that last more than fifteen matches. I find that in ten out of the twelve treatments the theory is accurate to within 0.15, and in seven out of the twelve within 0.05. In the two other treatments, the theory does well qualitatively, predicting substantial cooperation, but fares poorly quantitatively, with participants achieving, in one case, 0.34 less than is predicted by the theory.

I compare the theory to existing theories that are primarily motivated by learning. These theories are estimated from data on repeated prisoner's dilemma games, but do no better than noise theory.

## References


Abreu, D., D. Pearce and E. Stacchetti (1986): "Optimal cartel equilibria with imperfect monitoring," *Journal of Economic Theory* 39: 251-269.

Arechar, A. A., M. Kouchaki, and D. G. Rand (2018): "Examining spillovers between long and short repeated prisoner's dilemma games played in the laboratory," *Games*, 9(1), 5.

Battalio, R., L. Samuelson and J. Van Huyck (2001): "Optimization Incentives and Coordination Failure in Laboratory Stag Hunt Games," *Econometrica* 69: 749-764.

Bergin, J., and B. Lipman (1996): "Evolution with state-dependent mutations," *Econometrica*: 943-956.

Blonski, M., P. Ockenfels and G. Spagnolo (2011): "Equilibrium Selection in the Repeated Prisoner's Dilemma: Axiomatic Approach and Experimental Evidence," *American Economic Journal: Microeconomics* 3: 164–92.

Blonski, M. and G. Spagnolo (2015): ""Prisoners' other dilemma,"" *International Journal of Game Theory* 44: 61-81.

Breitmoser, Y. (2015): "Cooperation, but no reciprocity: Individual strategies in the repeated prisoner's dilemma," *American Economic Review* 105: 2882-2910.

Bruttel, L. and U. Kamecke (2012): "Infinity in the lab. How do people play repeated games?" *Theory and Decision* 72: 205-219.

Carlsson, H. and E. Van Damme (1993): "Global games and equilibrium selection," *Econometrica*: 989-1018.

Clark, K., S. Kay and M. Sefton (2001): "When are nash equilibria self-enforcing? an experimental analysis," *International Journal of Game Theory* 29: 495–515.

Case, R. H. (1960): "The Problem of Social Cost," *Journal of Law and Economics* 3: 1-44.

Coate, S. and M. Conlin (2004): "A Group Rule-Utilitarian Approach to Voter Turnout: Theory and Evidence," *American Economic Review* 94: 1476-1504.

Cooper, R. W., D. V. DeJong, R. Forsyth and T. W. Ross (1992): "Communication in coordination games," *Quarterly Journal of Economics* 107: 739–771.

Dal Bó, P. Dal (2005): "Cooperation under the shadow of the future: experimental evidence from infinitely repeated games," *American Economic Review* 95: 1591-1604.

Dal Bó, P. and G. R. Fréchette, G. R. (2011): "The evolution of cooperation in infinitely repeated games: Experimental evidence," *American Economic Review* 101: 411-429.

Dal Bó, P. and G. R. Fréchette (2015): "Strategy Choice in the Infinitely Repeated Prisoner's Dilemma," unpublished

Dal Bó, P., and G. R. Fréchette (2019): "Strategy choice in the infinitely repeated prisoner's dilemma," *American Economic Review* 109: 3929-3952.

Dreber, A., D. G. Rand, D. Fudenberg and M. A. Nowak(2008): "Winners don't punish," *Nature* 452: 348-351.

Dal Bó, P., and G. R. Fréchette, G. R. (2018): "On the determinants of cooperation in infinitely repeated games: A survey," *Journal of Economic Literature* 56: 60-114.

Dal Bó, P., G. R. Fréchette and J. Kim (2021): "The determinants of efficient behavior in coordination games," *Games and Economic Behavior* 130: 352-368.

Dubois, D., M. Willinger and P. Van Nguyen (2012): "Optimization incentive and relative riskiness in experimental stag-hunt games," *International Journal of Game Theory* 4: 369–380.

Duffy, J. and N. Feltovich (2002): "Do actions speak louder than words? an experimental comparison of observation and cheap talk," *Games and Economic Behavior* 39: 1–27.

Ely, J. C. (2002): "Local conventions," *BE Journal of Theoretical Economics*: 1-32.

Fedderson, T. and A. Sandroni (2006): "A Theory of Participation in Elections," *American Economic Review* 96: 1271–1282.

Feltovich, N., A. Iwasaki and S. H. Oda (2012): "Payoff levels, loss avoidance, and equilibrium selection in games with multiple equilibria: an experimental study," *Economic Inquiry* 50: 932–952.

Fudenberg, D., D. Levine and E. Maskin (1994): "The Folk Theorem With Imperfect Public Information," *Econometrica* 62: 997-1039.

Fudenberg, D., and E. Maskin (1986), "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econome*trica 54: 533.

Fudenberg, D., and G. Karreskog Rehbinder (2024): "Predicting Cooperation with Learning Models," *American Economic Journal: Microeconomics* 16: 1-32.

Garner, Richard T. and Bernard Rose (1967): *Moral Philosophy: A Systematic Introduction to Normative Ethics and Meta-ethics* New York: Macmillan.

Ghidoni, R. and S. Suetens(2022): "The effect of sequentiality on cooperation in repeated games," *American Economic Journal: Microeconomics* 14: 58-77.

Graser, C., and M. van Veelen, M. (2024): "Repeated prisoner's dilemmas with errors: How much subgame-perfection, how much forgiveness, and how much cooperation?" University of Amsterdam

Harsanyi, J. C. (1977): "Rule Utilitarianism and Decision Theory," *Social Ethics*: 25-53.

Harsanyi, J. C. and R. Selten (1988): *A general theory of equilibrium selection in games*, MIT Press Books.

Herrera, H., M. Morelli, M. and S. Nunnari (2016): "Turnout Across Democracies," *American Journal of Political Science* 60: 607-624.

Juang, W.-T. and H. Sabourian(2021): "Rules and Mutation - A Theory of How Efficiency and Rawlsian Egalitarianism/Symmetry May Emerge," Cambridge Working Papers in Economics, CWPE: 2101 CINET: 2101

Kagel, J. H. and D. R. Schley (2013): "How economic rewards affect cooperation reconsidered," *Economics Letters* 121: 124-127.

Kandori, Michihiro, George Mailath, and Rafael Rob (1993): "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica* 61: 29-56.

Kolle, F., S. Quercia, and E. Tripodi (2025): "Social preferences under the shadow of the future," *Experimental Economics*, 28: 22–42.

Levine, D. K. (1986): "Modeling altruism and spitefulness in experiments," *Review of Economic Dynamics* 1: 593-622.

Levine, D. K. (2024): "Behavioral Mechanism Design in the Repeated Prisoner's Dilemma," Levine's Working Paper Archive

Levine, D. K. (2025): "Behavioral Mechanism Design as a Benchmark for Experimental Studies," mimeo RHUL.

Levine, D. K. (2026): "The Largesse Design Problem," RHUL.

Levine, D. K., and S. Modica (2013): "Anti-malthus: Conflict and the evolution of societies," *Research in Economics*: 289-306.

Levine, D. K., and S. Modica (2016): "Dynamics in stochastic evolutionary models," *Theoretical Economics*: 89-131.

Levine, D. K. and A. Mattozzi (2020): "Voter Turnout with Peer Punishment," *American Economic Review* 110: 3298-3314.

Levine, D. K., A. Mattozzi and S. Modica (2022): *Social Mechanisms and Political Economy: When Lobbyists Succeed, Pollsters Fail and Populists Win*, mimeo RHUL.

Levine, D. K. (2026): "The Evolution of Resilience," *Journal of Economic Theory*, forthcoming.

Mailath, G. J., and L. Samuelson, L. (2006), *Repeated Games and Reputations: Long-run Relationships*, Oxford university press.

McKelvey, R. D. and T. R. Palfrey (1998): "Quantal response equilibria for extensive form games," *Experimental Economics* 1: 9-41.

Mill, J. S. (1861): *Utilitarianism.*

Nikiforakis, N. and H. T. Normann (2008): "A Comparative Statics Analysis of Punishment in Public-good Experiments," *Experimental Economics* 11: 358-369.

Ostrom, Elinor (1990): *Governing the commons: The evolution of institutions for collective action*, Cambridge university press.

Proto, E., A. Rustichini and A. Sofianos (2019): "Intelligence, personality, and gains from cooperation in repeated interactions," *Journal of Political Economy* 127: 1351-1390.

Rawls, J. (1971): *A Theory of Justice*, Taylor and Francis.

Roemer, J. E. (2010): "Kantian equilibrium," *Scandinavian Journal of Economics* 112: 1-24.

Schelling, T. C. (1960): *The Strategy of Conflict*, Harvard University Press.

Schmidt, D., R. Shupp, J. M. Walker and E. Ostrom (2003): "Playing safe in coordination games: the roles of risk dominance, payoff dominance, and history of play," *Games and Economic Behavior* 42: 281–299.

Sherstyuk, K., N. Tarui and T. Saijo (2013): Payment schemes in infinite-horizon experimental games, Experimental Economics 16: 125-153.

Straub, P. G. (1995): "Risk dominance and coordination failures in static games," *Quarterly Review of Economics and Finance* 35: 339–363.

Townsend, R. M. (1994): "Risk and insurance in village India," *Econometrica*, 539-591.

Van Huyck, J. B., R. C. Battalio and R.O. Beil (1990): "Tacit coordination games, strategic uncertainty, and coordination failure," *American Economic Review* 80: 234-248.

Young, P. (1993): "The Evolution of Conventions," *Econometrica* 61: 57-83

**Appendix I: Computation of the Equilibria**

Recall that the algorithm, described in Section 2.3, searched over a $K+1$ grid using an indifference tolerance of $\epsilon$. The base run used $K = 80$ and $\epsilon = .002$.

An alternative to computing equilibria in the $2 \times 2$ game with an $\epsilon$ indifference tolerance would be to find grid points at which a deviation leads to no greater utility gain than $\epsilon$. Although this algorithm was tried and yields similar results, it has an important glitch. When the only equilibrium is where all players intend to defect, approximate equilibrium allows them to cooperate with a low probability, as the loss from doing so is low. The algorithm proved all too capable of finding, depending on the tolerance, a solution one or two grid points away from the actual pure strategy equilibrium. Although this is perhaps more a matter of aesthetics - it does not really change the analysis - I switched to the more robust computational method described in Section 2.3.

To examine the numerical stability of the algorithm, I first report results for different grid sizes, below, in Table 9.1.

Table 9.1: Robustness of Welfare to Grid Size

| payoff matrix | $\delta$ | $K = \mathbf{80}$ | $K = 20$ | $K = 10$ | data $\overline{\omega}$ |
|---|---|---|---|---|---|
| *DalFre:B* | 0.50 | 0.14 | 0.14 | 0.14 | 0.15 |
| *DreEtAl:A* | 0.75 | 0.17 | 0.17 | 0.17 | 0.14 |
| *AreCouRan* | 0.13 | 0.17 | 0.17 | 0.17 | 0.19 |
| *BruKam* | 0.80 | 0.21 | 0.21 | 0.21 | 0.35 |
| *DalFre:A* | 0.50 | 0.27 | 0.27 | 0.27 | 0.14 |
| *DalFre:A* | 0.75 | 0.27 | 0.27 | 0.27 | 0.25 |
| DalFre:C | 0.50 | 0.67 | 0.67 | 0.67 | 0.39 |
| DreEtAl_08:B | 0.75 | 0.67 | 0.67 | 0.67 | 0.46 |
| KagSch | 0.75 | 0.67 | 0.67 | 0.67 | 0.58 |
| SheTarSai | 0.75 | 0.67 | 0.67 | **0.27** | 0.69 |
| DalFre:B | 0.75 | 0.68 | 0.68 | **0.50** | 0.64 |
| DalFre:C | 0.75 | 0.73 | 0.73 | 0.73 | 0.69 |

bold: greater than 0.04 from base

For any given value of $\epsilon > 0$ a fine enough grid will find a grid point near every mixed equilibrium. As can be seen, with $K = 10$, the algorithm has trouble finding mixed equilibria on the grid. However, from $K = 20$, even quadrupling the grid size to 80 does not change welfare for any treatment. As the BruKam treatment is near a

bifurcation point, and also sensitive to the choice of $\epsilon$, and to verify that the "solution" found for large $\epsilon$ is spurious, I also checked BruKam with a grid size of 160. This also resulted in a welfare of 0.21.

Table 9.2 reports results for the base value of $\epsilon = .002$, for quadruple, double, and half that value.

Table 9.2: Robustness of Welfare to Tolerance

| payoff matrix | $\delta$ | $\epsilon = 0.001$ | $\epsilon = \mathbf{0.002}$ | $\epsilon = 0.004$ | $\epsilon = 0.008$ | data $\overline{\omega}$ |
|---|---|---|---|---|---|---|
| *DalFre:B* | 0.50 | 0.14 | 0.14 | 0.14 | 0.14 | 0.15 |
| *DreEtAl:A* | 0.75 | 0.17 | 0.17 | 0.17 | 0.17 | 0.14 |
| *AreCouRan* | 0.13 | 0.17 | 0.17 | 0.17 | 0.17 | 0.19 |
| *BruKam* | 0.80 | 0.21 | 0.21 | **0.65** | **0.67** | 0.35 |
| *DalFre:A* | 0.50 | 0.27 | 0.27 | 0.27 | 0.27 | 0.14 |
| *DalFre:A* | 0.75 | 0.27 | 0.27 | 0.27 | 0.27 | 0.25 |
| DalFre:C | 0.50 | 0.67 | 0.67 | 0.67 | 0.67 | 0.39 |
| DreEtAl:B | 0.75 | 0.67 | 0.67 | 0.67 | 0.68 | 0.46 |
| KagSch | 0.75 | 0.67 | 0.67 | 0.68 | 0.70 | 0.58 |
| SheTarSai | 0.75 | 0.67 | 0.67 | 0.68 | 0.71 | 0.69 |
| DalFre:B | 0.75 | 0.69 | 0.69 | 0.69 | 0.70 | 0.64 |
| DalFre:C | 0.75 | 0.74 | 0.74 | 0.75 | 0.75 | 0.69 |

bold: greater than 0.04 from base

Choosing $\epsilon$ is not as easy as choosing $K$. While, for given $\epsilon$, larger $K$ is unambiguously better, for given $K$, too small an $\epsilon$ will miss legitimate equilibria, while too large an $\epsilon$ will find spurious equilibria. As indicated, $\epsilon = .002$ is stable in the sense that refining the grid from $K = 20$ to $K = 80$ does not change welfare. Increasing $\epsilon$ to $\epsilon = .004$ or larger increases welfare, allowing for spurious solutions. Reducing $\epsilon$ by a factor of two does not change welfare, indicating that there are no further spurious solutions to be eliminated.

As an additional check on BruKam, I verified that $\epsilon = .004$, taking $\phi = .338$, just below bifurcation point for other treatments, results in welfare of 0.21.

**Appendix II: Strongly Symmetric Equilibrium**

As a further check on the numeric stability of the code, I computed a different kind of equilibrium where exact calculations are possible. These are strongly symmetric equilibria, as described by Abreu, Pearce and Stacchetti (1986). These differ from one-period memory equilibrium in three ways.

- Regardless of the history, players must always play identical strategies. This limits the set of equilibrium, so reduces the maximum welfare.
- Players are allowed access to a public randomization device, although they did not have access to such a device in the actual experiment. This increases the maximum welfare.
- There is no limitation on memory. This increases the maximum welfare.

I indicate below how to compute the best pure strategy strongly symmetric equilibrium. Unlike memory-one equilibrium, mixing is not needed to create indifference: the public randomization device is used to coordinate punishments. However, mixing can potentially improve welfare if the gain to defecting when the opponent is cooperating is greater than when the opponent is defecting. Mixing, in this case, reduces the gain to defecting when both players are cooperating, so requires less punishment, potentially increasing welfare.

It is relatively easy to compute an upper bound on how much welfare is possible with mixing. Let $v$ denote the best pure strategy welfare and $\sigma_C$ the probability of cooperating at the best equilibrium. Since welfare when $\sigma_C$ is used is at most $(\sigma_C)^2$, it must be that $(\sigma_C)^2 \geq v$. Let $g$ be the gain to defecting when the opponent is cooperating and $\ell$ be the gain to defecting when the opponent is defecting. If $\ell < g$ then the game in which the gain to defecting when the opponent is cooperating is given by $g' = g\sqrt{v} + \ell(1 - \sqrt{v})$ has at least as much welfare as the original game, and the best welfare is attainable in this game with pure strategies.

The calculations of welfare for strongly symmetric equilibrium are shown below in Table 9.3.

Table 9.3: Best Strongly Symmetric Equilibrium

| payoff matrix | $\delta$ | theory $\hat{\omega}$ | strong | upper | data $\overline{\omega}$ |
|---|---|---|---|---|---|
| *DalFre:B* | 0.50 | 0.14 | 0.14 | 0.14 | 0.15 |
| *DreEtAl:A* | 0.75 | 0.17 | 0.17 | 0.17 | 0.14 |
| *AreCouRan* | 0.13 | 0.17 | 0.17 | 0.17 | 0.19 |
| *BruKam* | 0.80 | 0.21 | **0.62** | **0.65** | 0.35 |
| *DalFre:A* | 0.50 | 0.27 | 0.27 | 0.27 | 0.14 |
| *DalFre:A* | 0.75 | 0.27 | 0.27 | 0.27 | 0.25 |
| DalFre:C | 0.50 | **0.67** | **0.74** | **0.74** | 0.39 |
| DreEtAl:B | 0.75 | **0.67** | 0.60 | 0.60 | 0.46 |
| KagSch | 0.75 | 0.67 | 0.63 | 0.68 | 0.58 |
| SheTarSai | 0.75 | 0.67 | 0.65 | 0.73 | 0.69 |
| DalFre:B | 0.75 | 0.69 | 0.70 | 0.70 | 0.64 |
| DalFre:C | 0.75 | 0.74 | 0.74 | 0.74 | 0.69 |

strong: welfare at best pure strongly symmetric equilibrium
upper: upper bound on welfare at best mixed strongly symmetric equilibrium
italics: theory $\hat{\omega}$ equal to minimum
bold: error of more than 0.15

The highlights are these. There is a good strongly symmetric equilibrium in BruKam. As indicated, this is the one treatment that is not numerically robust. Otherwise, as expected, all the treatments predicted to have no cooperation yield the same prediction and the same welfare.

Among the cooperative treatments, DalFre:C;$\delta = 0.5$ has an even better strongly symmetric equilibrium than one period memory equilibrium, fitting the data even less well. The treatment DreEtAl:B has no strongly symmetric equilibrium with welfare as high as the best memory one equilibrium, and consequently fits the data better. Note that empirically in DreEtAl:B players in fact do not use symmetric strategies, cooperating 0.17 for $CD$ and 0.44 for $DC$.

The remaining four treatments have very similar strongly symmetric and memory one welfare.

*Solving the LP Problem for Best Pure Strategy Strongly Symmetric Equilibrium*

It is convenient to define the probability of a tremble $\pi = \phi/2$. Strongly symmetric equilibrium makes use of two signals: $DD$ and $DC/CD$. Denote by $P_{DD}$ the average present value punishment for the first signal and $P_{DC}$ that for the second signal. If the actual play is $CC$ the probability of the $DD$ signal is $p_{DD} = \pi^2$ and the probability of

the $DC/CD$ signal is $p_{DC} = 2\pi(1-\pi)$. Similarly if actual play is $DC$, the probability of the $DD$ signal is $q_{DD} = (1-\pi)\pi$ and the $DC/CD$ signal is $q_{DC} = (1-\pi)^2 + \pi^2$. Denote the optimal welfare by $v$.

Define the relative gain to cooperation

$$R \equiv \frac{U^1(CC)}{U^1(DC) - U^1(CC)},$$

the abbreviations, $u = U^1(CC), n = U^1(DD)$ and $g = U^1(DC) - U^1(CC)$.

Then define the interest rate

$$\rho \equiv \frac{1-\delta}{\delta}$$

and the noisy interest rates

$$\eta_{DD} \equiv \frac{\rho + p_{DD}}{q_{DD} - p_{DD}}, \eta_A \equiv \frac{\rho + p_{DD} + p_{DC}}{(q_{DD} + q_{DC}) - (p_{DD} + p_{DC})}.$$

**Proposition 9.1.** *If $\eta_A \leq R < \eta_{DD}$ and $q_{DC} - ((1+\eta_{DD})/\eta_{DD})p_{DC} \leq 0$ the solution is the static Nash equilibrium $P_{DC} = P_{DD} = 0$ and $v =$. Otherwise it is given below*

Table 9.4: Best Strongly Symmetric Equilibrium

| | $P_{DC}$ | $v$ | $P_{DD}$ |
|---|---|---|---|
| $R < \eta_A$ | 0 | $n$ | 0 |
| $\eta_A \leq R < \eta_{DD}$ | $\rho\left(\frac{1-\eta_{DD}R}{q_{DC}-p_{DC}-\eta_{DD}p_{DD}}\right)g$ | $u - \left(\frac{p_{DD}(u-n)+p_{DC}P_{DC}}{\rho+p_{DD}}\right)$ | $v - n$ |
| $R \geq \eta_{DD}$ | 0 | $u - \left(\frac{p_{DD}}{q_{DD}-p_{DD}}\right)g$ | $\left(\frac{\rho}{q_{DD}-p_{DD}}\right)g$ |

*Note that the columns are ordered so that the solutions in each column depend only on the results from the previous columns.*

*Proof.* Average present value $v$ in the best strongly symmetric equilibrium is then characterized by solving the LP problem of choosing $v, P_{DD}, P_{DC}$ to maximize $v$ subject to the two incentive constraints

$$v = (1-\delta)u + \delta((1 - p_{DD} - p_{DC})v + p_{DD}(v - P_{DD}) + p_{DC}(v - P_{DC}))$$

$$v = (1-\delta)(u+g) + \delta((1 - q_{DD} - q_{DC})v + q_{DD}(v - P_{DD}) + q_{DC}(v - P_{DC}))$$

and the feasibility constraints $v - P_{DD}, v - P_{DC} \geq n$ and $P_{DD}, P_{DC} \geq 0$.

Rewrite the first incentive constraint as

$$\rho(u - v) = p_{DD}P_{DD} + p_{DC}P_{DC}.$$

Similarly the second constraint is

$$\rho(u + g - v) = q_{DD}P_{DD} + q_{DC}P_{DC}.$$

Alternatively the two can be combined to give a second equation entirely in terms of the punishments

$$\rho g = (q_{DD} - p_{DD})P_{DD} + (q_{DC} - p_{DC})P_{DC}.$$

Maximizing $v$ is the same as minimizing the aggregate on-path punishment cost $p_{DD}P_{DD} + p_{DC}P_{DC}$ subject to the punishment constraint. Equivalently $Q_{DD}, Q_{DC}$ can be chosen to minimize

$$\frac{p_{DD}}{q_{DD} - p_{DD}}Q_{DD} + \frac{p_{DC}}{q_{DC} - p_{DC}}Q_{DC}$$

subject to $\rho g = Q_{DD} + Q_{DC}$. Hence the maximum $Q$ (and $P$) should be allocated to the smallest $p_K/(q_K - p_K)$.

Computing the ratios for the two punishments gives

$$\frac{p_{DD}}{q_{DD} - p_{DD}} = \frac{1}{1 - 2\pi}\pi$$

$$\frac{p_{DC}}{q_{DC} - p_{DC}} = \frac{1}{1 - 2\pi}\pi\frac{2 - 2\pi}{1 - 2\pi}$$

so that $DD$ is more efficient.

**Only $DD$ Signal Used**

If only $DD$ is punished and the $v - P_{DD} \geq n$ constraint does not bind, then the two equality constraints can be solved to find

$$P_{DD} = \rho\frac{1}{q_{DD} - p_{DD}}g$$

$$v = U^1(CC) - \frac{p_{DD}}{q_{DD} - p_{DD}} g.$$

Hence

$$v - P_{DD} = U^1(CC) - \frac{p_{DD}}{q_{DD} - p_{DD}} g - P$$

$$= U^1(CC) - \frac{\rho + p_{DD}}{q_{DD} - p_{DD}} g \geq n$$

or

$$\frac{\rho + p_{DD}}{q_{DD} - p_{DD}} \leq R. \tag{9.1}$$

**Both Signals Used**

If inequality 9.1 fails then both signals must be used and $P_{DD} = v - n$. The two equalities are now

$$\rho(u - v) = p_{DD}(v - n) + p_{DC} P_{DC}$$

$$\rho g = (q_{DD} - p_{DD})(u - n)(q_{DC} - p_{DC}) P_{DC}.$$

Notice that increasing $P_{DC}$ has two effects: it increases the punishment received when $DC/CD$ occur, but it also reduces $v - n$ which is the punishment received when $DD$ occurs. If $q_{DC} - p_{DC} - \frac{q_{DD} - p_{DD}}{\rho + p_{DD}} p_{DC} \leq 0$ the latter effect weakly dominates the first so the only solution is that static Nash. Otherwise

$$P_{DC} = \rho g \frac{1 - \frac{q_{DD} - p_{DD}}{\rho + p_{DD}} R}{q_{DC} - p_{DC} - \frac{q_{DD} - p_{DD}}{\rho + p_{DD}} p_{DC}}$$

$$v = \frac{\rho u + p_{DD} n}{\rho + p_{DD}} - \frac{p_{DC} P_{DC}}{\rho + p_{DD}}.$$

It remains to check that $v - P_{DC} \geq n$. This is

$$\frac{1}{\rho + p_{DD}} \left(\rho(u - n) + p_{DC} P_{DC}\right) \geq P_{DC}$$

which can be written as

$$R \geq \frac{\rho + p_{DD} + p_{DC}}{(q_{DD} + q_{DC}) - (p_{DD} + p_{DC})}.$$

If this fails then the only strong symmetric equilibrium is the static Nash. □

### Appendix III: Session Data

It would be concerning if welfare varied between very low and very high, with the differences between treatments depending on the frequency of very high welfare sessions. This is not what the theory predicts. Below in Table 9.5 I give the minimum and maximum welfare across sessions for each treatment.

Table 9.5: Session Data

| payoff matrix | $\delta$ | theory $\hat{\omega}$ | data $\overline{\omega}$ | min $\overline{\omega}$ | max $\overline{\omega}$ | sessions |
|---|---|---|---|---|---|---|
| *DalFre:B* | 0.50 | 0.14 | 0.15 | 0.12 | 0.20 | 6 |
| *DreEtAl:A* | 0.75 | 0.17 | 0.14 | one session | | 1 |
| *AreCouRan* | 0.13 | 0.17 | 0.19 | 0.13 | 0.23 | 7 |
| *BruKam* | 0.80 | 0.21 | 0.35 | 0.29 | **0.39** | 3 |
| *DalFre:A* | 0.50 | 0.27 | 0.14 | **0.05** | 0.21 | 12 |
| *DalFre:A* | 0.75 | 0.27 | 0.25 | 0.16 | 0.32 | 9 |
| DalFre:C | 0.50 | **0.67** | 0.39 | **0.23** | **0.49** | 17 |
| DreEtAl:B | 0.75 | **0.67** | 0.46 | one session | | 1 |
| KagSch | 0.75 | 0.67 | 0.58 | **0.50** | 0.67 | 6 |
| SheTarSai | 0.75 | 0.67 | 0.69 | 0.57 | 0.82 | 4 |
| DalFre:B | 0.75 | 0.69 | 0.64 | **0.53** | 0.80 | 6 |
| DalFre:C | 0.75 | 0.74 | 0.69 | **0.54** | 0.81 | 17 |

bold: theory differs by more than 0.15

The bottom line here is that while there is between session variation, in the low welfare treatments all the sessions have relatively low welfare, and in the high welfare treatments all the sessions have relatively high welfare.

### Appendix IV: Strategies

How does the theory do in predicting play conditional on the initial and four subsequent states, $y \in \{CC, CD, DC, DD\}$? I present the results in two tables below, Table 9.6 covering the low welfare cases and Table 9.7 covering the high welfare cases. Following that, I consider an alternative theory for predicting conditional cooperation, the semi-grim trigger equilibria of Breitmoser (2015).

Table 9.6: Conditional Probabilities of Player 1 Cooperating

| payoffs | $\delta$ | initial | CC | CD | DC | DD | CC frac |
|---|---|---|---|---|---|---|---|
| *DalFre:B* | 0.50 | 0.17 | **0.17** | 0.17 | 0.17 | 0.17 | |
| | data | 0.19 | 0.88 | 0.32 | 0.31 | 0.06 | 0.02 |
| *DreEtAl:A* | 0.75 | 0.17 | **0.17** | 0.17 | 0.17 | **0.17** | |
| | data | 0.27 | 0.74 | 0.12 | 0.26 | 0.00 | 0.04 |
| *AreCouRan* | 0.13 | 0.17 | **0.17** | 0.17 | 0.17 | 0.17 | |
| | data | 0.19 | 0.75 | 0.19 | 0.13 | 0.10 | 0.00 |
| *BruKam* | 0.80 | **0.17** | **0.17** | 0.17 | 0.17 | 0.17 | |
| | data | 0.45 | 0.95 | 0.23 | 0.27 | 0.05 | (0.17) |
| *DalFre:A* | 0.50 | 0.17 | **0.17** | 0.17 | 0.17 | 0.17 | |
| | data | 0.09 | 0.96 | 0.29 | 0.33 | 0.02 | 0.03 |
| *DalFre:A* | 0.75 | 0.17 | **0.17** | 0.17 | **0.17** | 0.17 | |
| | data | 0.23 | 0.96 | 0.30 | 0.36 | 0.03 | 0.08 |

CC,CD,DC,DD are last the period outcome
bold: error of more than 0.15
CC frac: fraction of games in which CC is observed

For the low-welfare case there are 9 anomalies. Of these, 6 are due to the fact that empirically there is a very high cooperation rate following CC, so, excluding *CC*, there are 3. By contrast, other theories predict always defect, with the exception of BruKam, where the basin and critical discount factor from Section 5 predict grim-trigger. These theories generate 20 anomalies, and 15, excluding CC. Overall, noise theory does much better.

The fact that stands out is the high level of cooperation following CC in the data, a fact for which, except for BruKam, no theory offers an explanation. There are two *ad hoc* explanations for this high level of cooperation. The first is the observation that, except for BruKam, the actual frequency with which CC occurs, reported in the final column of Table 9.6, is quite low, so it could be argued that participants are poorly informed or do not care. In this case, however, it is not clear why they cooperate rather than being all over the map.

My best guess is players who tremble onto cooperating in the initial period actually tremble onto grim-trigger. While that can explain the data reasonably well, and is a sensible way to "test the waters" from an individual perspective, it falls far short of being a coherent theory.

Next, in Table 9.7, I examine the high welfare cases.

Table 9.7: Conditional Probabilities of Player 1 Cooperating

| payoffs | $\delta$ | initial | CC | CD | DC | DD |
|---|---|---|---|---|---|---|
| DalFre:C | 0.50 | **0.83** | 0.83 | 0.17 | 0.17 | **0.83** |
| | data | 0.52 | 0.97 | 0.31 | 0.21 | 0.04 |
| DreEtAl:B | 0.75 | **0.83** | 0.83 | 0.18 | **0.83** | 0.17 |
| | data | 0.62 | 0.94 | 0.17 | 0.44 | 0.04 |
| KagSch | 0.75 | 0.83 | 0.83 | 0.17 | **0.52** | 0.18 |
| | data | 0.69 | 0.96 | 0.27 | 0.32 | 0.03 |
| SheTarSai | 0.75 | **0.83** | 0.83 | **0.17** | 0.35 | 0.18 |
| | data | 0.63 | 0.96 | 0.48 | 0.42 | 0.08 |
| DalFre:B | 0.75 | 0.83 | 0.83 | 0.36 | **0.83** | **0.22** |
| | data | 0.71 | 0.98 | 0.37 | 0.39 | 0.06 |
| DalFre:C | 0.75 | 0.83 | **0.83** | **0.77** | **0.83** | **0.36** |
| | data | 0.81 | 0.99 | 0.27 | 0.27 | 0.06 |

CC,CD,DC,DD are last the period outcome
bold: error of more than 0.15

The theory does less well with the high-welfare treatments. The theory predicts higher forgiveness, and greater first-period cooperation, than is observed in the data. There are, in total, 13 anomalies. What can be said in favor of the theory is that grim-trigger is even worse, generating 18 anomalies.

A one-period memory theory without noise, designed to predict strategies, is Breitmoser (2015)'s theory of semi-grim equilibrium. After conducting a meta-study of indefinitely repeated prisoners' dilemma games, including the Dal Bó (2005) treatments discussed here, Breitmoser (2015) observes that the probability of cooperating following CD is similar to that following DC. This leads him to propose equilibrium in semi-grim strategies. These are memory-one mixed strategies that are constrained to have the same probability of cooperating following CD as DC. There are either two types of these equilibria or none at all. In case there are none, Breitmoser (2015) predicts always defect. For the low-welfare games this is always the case except for BruKam. For the remaining five low welfare treatments, excluding CC, always defect generates 12 anomalies. By contrast noise theory generates only 2 anomalies so does much better.

For the high welfare treatments and BruKam there are two different types of equilibria, depending on the characteristics of the first-period game. In the indifference game, also called belief free, in the first-period game players are indifferent between

cooperating and defecting regardless of what the other player is doing in that first period. Notice that noise theory often generates such strategies. There are a continuum of such equilibria, indexed by a parameter $r$, and for each of the seven games, Table 9.8 below shows the equilibrium that minimizes the mean square error to the data as reported in the table. Because there is a closed-form solution for the equilibria, and the theoretical conditional probabilities are linear in $r$, there is a unique value of $r$ that minimizes mean square error, and there is a closed-form solution for it. In no case does a semi-grim equilibrium yield any prediction about period one, since players are completely indifferent, regardless of beliefs about opponent's first period strategy.

Table 9.8: Conditional Probabilities of Player 1 Cooperating: Semi-Grim Indifference

| payoffs | $\delta$ | initial | CC | CD | DC | DD |
|---|---|---|---|---|---|---|
| *BruKam* | 0.80 | | 0.83 | 0.36 | 0.35 | 0.00 |
| | data | 0.45 | 0.95 | 0.23 | 0.27 | 0.05 |
| DalFre:C | 0.50 | | **0.79** | **0.68** | **0.68** | 0.00 |
| | data | 0.52 | 0.97 | 0.31 | 0.21 | 0.04 |
| DreEtAl:B | 0.75 | | 0.89 | **0.44** | 0.44 | 0.00 |
| | data | 0.62 | 0.94 | 0.17 | 0.44 | 0.04 |
| KagSch | 0.75 | | 0.86 | 0.33 | 0.33 | 0.06 |
| | data | 0.69 | 0.96 | 0.27 | 0.32 | 0.03 |
| SheTarSai | 0.75 | | 0.97 | 0.38 | 0.38 | **0.23** |
| | data | 0.63 | 0.96 | 0.48 | 0.42 | 0.08 |
| DalFre:B | 0.75 | | **0.82** | 0.47 | 0.47 | 0.01 |
| | data | 0.71 | 0.98 | 0.37 | 0.39 | 0.06 |
| DalFre:C | 0.75 | | **0.57** | **0.50** | **0.50** | 0.04 |
| | data | 0.81 | 0.99 | 0.27 | 0.27 | 0.06 |

first row: semi-grim prediction
CC,CD,DC,DD are last the period outcome
bold: error of more than 0.15

For these seven treatments, excluding the initial period, since semi-grim equilibrium makes no prediction about that, semi-grim, for which 7 parameters are estimated, generates 9 anomalies. In comparison, noise theory generates 11, so does worse, but not terribly so.

The second type of semi-grim equilibrium is when first period play is a coordination game. Breitmoser (2015) refers to this as Eff LimLog SG. There are three

symmetric equilibria: cooperating, defecting, and mixed. For each treatment there is a unique equilibrium of this type, and the prediction for initial period play is that it must be either cooperate, defect, or play the unique mixed equilibrium. In each of the seven treatments below, in Table 9.9, I have chosen the initial period equilibrium that best fits the data.

Table 9.9: Conditional Probabilities of Player 1 Cooperating: Semi-Grim Coordination

| payoffs | $\delta$ | initial | CC | CD | DC | DD |
|---|---|---|---|---|---|---|
| *BruKam* | 0.80 | **0.17** | 1.00 | 0.17 | 0.17 | 0.00 |
| | data | 0.45 | 0.95 | 0.23 | 0.27 | 0.05 |
| DalFre:C | 0.50 | **0.33** | 1.00 | 0.33 | 0.33 | 0.00 |
| | data | 0.52 | 0.97 | 0.31 | 0.21 | 0.04 |
| DreEtAl:B | 0.75 | **1.00** | 1.00 | 0.27 | **0.27** | 0.00 |
| | data | 0.62 | 0.94 | 0.17 | 0.44 | 0.04 |
| KagSch | 0.75 | **1.00** | 1.00 | 0.13 | **0.13** | 0.00 |
| | data | 0.69 | 0.96 | 0.27 | 0.32 | 0.03 |
| SheTarSai | 0.75 | **1.00** | 1.00 | **0.07** | **0.07** | 0.00 |
| | data | 0.63 | 0.96 | 0.48 | 0.42 | 0.08 |
| DalFre:B | 0.75 | **1.00** | 1.00 | **0.21** | **0.21** | 0.00 |
| | data | 0.71 | 0.98 | 0.37 | 0.39 | 0.06 |
| DalFre:C | 0.75 | **1.00** | 1.00 | 0.14 | 0.14 | 0.00 |
| | data | 0.81 | 0.99 | 0.27 | 0.27 | 0.06 |

first row: semi-grim prediction
CC,CD,DC,DD are last the period outcome
bold: error of more than 0.15

The coordination game does better than semi-grim indifference for conditional probabilities, with 6 anomalies, versus 9, but the predictions for initial period play are all anomalies. Hence the total anomalies for the seven treatments generated by semi-grim coordination are 13 versus 15 for noise theory. Again, noise theory does worse, but not terribly so.

The overall conclusion is that noise theory does not do terribly well, but does considerably better than alternatives.