Testing Theoretical Game Theory Results on a Large Scale: Prisoner's Dilemma on Facebook

by

Sunny Long

Submitted to the Department of Electrical Engineering and Computer Science
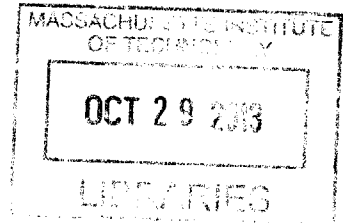in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2013

Author . . . .
Department of Electrical Engineering and Computer Science
July 15, 2013

Certified by . . . .
Constantinos Daskalakis
Associate Professor
Thesis Supervisor

Accepted by . . . .
Professor Albert R. Meyer
Chairman, Masters of Engineering Thesis Committee

Testing Theoretical Game Theory Results on a Large Scale: Prisoner's Dilemma on Facebook

by

Sunny Long

Submitted to the Department of Electrical Engineering and Computer Science
on July 15, 2013, in partial fulfillment of the
requirements for the degree of
Master of Science

Abstract

In my research, I designed and implemented an online game accessable to a large
diverse audience via the Facebook social network to test out game theoretic results
and study social interactions. In this game, we designed scenarios that mirrored clas-
sic game theory situations, most prominantly the Iterated Prisoner's Dilemma game.
The game allows for Facebook users to play each other in the Prisoner's Dilemma
game while competing for a common goal. We then gather data on strategies em-
ployed by players and compare them to classic game theoretic results and try to
explain similarities and differences. We also look at social factors that were promi-
nant in decision-making for players. Analysis includes disecting which factors (e.g.
gender, age, etc.) affect players' decisions and understanding winning strategies. We
discovered that under the social setting of the experiment, factors that came into
play for users' decision-making included gender as well as potential payoffs. On the
other hand, game play history and whether or not two people were friends were not
significant in decision-making. Despite a high overall cooperation rate (over 70%),
the top players predominantly betrayed their opponents while never playing the same
opponent too many times. The game is implemented in Ruby on Rails while all data
analysis is done via Python scripts.

Thesis Supervisor: Constantinos Daskalakis
Title: Associate Professor

Acknowledgments

I would like to express my very deep appreciation to Professor Daskalakis for his guidance, encouragement, and constructive suggestions during the planning and development of this research.

I would also like to extend thanks to Sam Sinai and Brody Childs for their advice and assistance in the implementation and analysis of this project.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Game theory, or more formally, "the study of mathematical models of conflict and cooperation between intelligent rational decision-makers," is a field in economics that seeks to model scenarios ranging from political science to psychology to biology. It is a useful tool to help model and solve situations involving multiple parties making decisions. In this chapter, we would like to introduce two games, the prisoner's dilemma game and the chicken game. Then, we will briefly touch on iterated versions of the prisoner's dilemma game, and finish by addressing how these games relate to situations in real life.

1.1   Prisoner's Dilemma

The prisoner's dilemma is one of the oldest and most influential games studied in game theory, first invented around 1950 by two Rand Corporation scientists. The classic backstory for the prisoner's dilemma goes something as follows: You and your friend have just pulled off the most amazing jewelry heist in the history of jewelry heists. You flee successfully to your hideout and decide that laying low for the next week is the best policy before deciding on what to do with all the stolen jewelry. Days one through six pass without a hitch, but on the morning of the seventh, both you and your friend are brought into the police station for questioning. Unfortunately, you are put in a separate interrogation room from your friend. Your interrogator sits

down opposite you, the lights seem to dim, and he leans in. He says: I have one question for you was your friend involved in robbing the jewelry store?

Of course this is a sticky situation you have the choice to betray your friend, or keep silent. And likewise, you are certain that your friend is in the same situation in the other room. What should you do?

In order to consider what to do, we need to explore the consequences for each possible action. As presented in the story, each player has two choices, either to **betray** the other (rat them out), or to **cooperate** with the other (by staying silent). Thus, there are four possible outcomes:

- (**P1**: Cooperate, **P2**: Cooperate)

- (**P1**: Cooperate, **P2**: Betray)

- (**P1**: Betray, **P2**: Cooperate)

- (**P1**: Betray, **P2**: Betray)

In the classic prisoner's dilemma game, the consequences were as follows:

If one player betrays and the other stayes silent, the betrayer would walk free (cut a deal with the police) while the silent party would have to serve three years. If both players cooperated (stayed silent), then both players would have to serve one year each. If both players ratted the other out (both betrayed), then each party would serve two years. A convenient way to represent these consequences is in a payoff matrix (or payoff diagram) shown in Table 1.1.

1.1.1   Solving the prisoner's dilemma game

The cononical prisoner's dilemma game is what is known as a "one-shot" game, which means that each player submits their strategy simultaneously, and without knowledge of the other's actions - for one game only. Thus, each player does not have information on their opponent's history of play. The only information they have access to is the payoff matrix.

16

	Cooperate	Betray
Cooperate	(-1, -1)	(-3, 0)
Betray	(0, -3)	(-2, -2)

Table 1.1: Payoff diagram representing the cononical prisoner's dilemma game. The first column represents player one's choices while the top row represents player two's choices. For each tuple of numbers, the first number is player one's payoff while the second number is player two's payoff. For instance, if player one plays *Betray* and player two plays *Cooperate*, we wind up at the square with payoff (0, -3). This means that player one gains 0 whereas player two loses 3 (i.e. player one doesn't go to jail and player two serves three years).

From table 1.1, we can logically deduce the ideal strategy for each player. As the game is symmetric (payoffs only depend on the particular strategies played and not on the players playing them), we will just consider player one's ideal strategy. We can see that he should play *Betray* regardless of what player two plays. He will always receive a higher payoff when he plays *Betray* since:

$$0 > -1 (\text{Betray, Cooperate}) \text{ vs. (Cooperate, Cooperate)}$$

$$-2 > -3 (\text{Betray, Betray}) \text{ vs. (Cooperate, Betray)}$$

Likewise, for player two, he will have the same considerations and end up playing *Betray*. Thus, the outcome will be (Betray, Betray) whereby both players will serve two years in prison. This is known as a Nash equilibrium (NE). A Nash equilibrium is defined to be both players choosing strategies such that no player will want to change their strategy (to be better off) while holding the other player's strategy constant. In the prisoner's dilemma, we see that given one player plays *Betray*, the other player will not want to change his strategy from *Betray* to *Cooperate*, thus a NE. On the other hand, if one player is playing *Betray* while the other is playing *Cooperate*, this would not be a Nash equilibrium since the player playing *Cooperate* stands to gain a better payoff by switching to *Betray*.

This is the only pure strategy Nash equilibrium for the prisoner's dilemma game. A pure strategy determines the move a player will make for any situation he could

	Swerve	Straight
Swerve	(0, 0)	(-1, 1)
Straight	(1, -1)	(-5, -5)

Table 1.2: Payoff diagram representing the chicken game.

face. This is in contrast to a mixed strategy which is an assignment of probabilities to each pure strategy.

In [20], using Brouwer's fixed point theorem, Nash showed that for every finite game, there exists a mixed strategy Nash equilibrium. For the prisoner's dilemma, we can think of the pure strategy NE that we found, (Betray, Betray), simply as a degenerate mixed strategy, where each player places a zero probability on playing *Cooperate*.

1.2 The Game of Chicken

Another well-studied game is known as the game of chicken (sometimes referred to as the "hawk-dove" game in biology). In this game, imagine two players are driving towards each other at high speeds on a one-lane street. Each player is faced with two choices: (Swerve, Straight). The two cars get nearer and nearer, and just before impact, each player makes his choice.

As table 1.2 shows, the worst outcome for both players is if they both pursue the *Straight* option, resulting in a car crash. However, if one player stays straight while the other swerves out of the way, the player that stayed straight will be considered the winner, while the player that swerved will be known as the "chicken," or loser.

This game is oftentimes used to model brinkmanshop in politics and international relations. One such example is the Cuban Missile Crisis where both the United States and Russia were pushed to the brink of war, and had both stayed their original courses (*Straight*), then the result could have easily been a war. However, both nations decided to step down and dismantle their armaments *Swerve*), avoiding a catastrophe.

1.2.1 Solving the chicken game

Unlike the prisoner's dilemma, there does not exist a symmetric pure strategy NE. Instead, there are three Nash equilbiria.

- (Swerve, Straight)

- (Straight, Swerve)

- Mixed strategy between *Swerve* and *Straight*

Again, for the pure strategies, we can verify that holding the opponent's player constant, a player will not want to deviate. As for the mixed strategy, we can solve for it by first choosing the mixed strategy for player one in order to make player 2 indifferent between his strategies. Then we do the same for player two to make player one indifferent between his strategies.

Let p be the probability that player one swerves, $1-p$ be the probability of staying straight. Let's also let q equal the probability that player two serves, and $1-q$ the probability he stays straight. We can now do the following:

$$
\begin{aligned}
U_{P1}(p;q) &= pqu_{p1}(Swerve, Swerve) + p(1-q)u_{p1}(Swerve, Straight)+ \\
&\quad (1-p)qu_{p1}(Straight, Swerve) + (1-p)(1-q)u_{p1}(Straight, Straight) \\
&= p[qu_{p1}(Swerve, Swerve) + (1-q)u_{p1}(Swerve, Straight)]+ \\
&\quad (1-p)[qu_{p1}(Straight, Swerve) + (1-q)u_{p1}(Straight, Straight)] \\
&= p[0q + -1(1-q)] + (1-p)[q + -5(1-q)] \\
&= p(q-1) + (1-p)[-5+6q]
\end{aligned}
$$

We see now that the term $(q-1)$ is the payoff if player one swerves, while $(6q-5)$ is the payoff from staying straight. $U_{p1}(p;q)$ is strictly increasing in p if $(q-1) > (6q-5)$ (i.e. $5/4 > q$); it is strictly decreasing with p if $(q-1) < (6q-5)$ (i.e. $5/4 < q$), and is constant if $(q-1) = (6q-5)$. Thus, player one's best response would be:

- $p = 1$ if $q < 4/5$

- $p = 0$ if $q > 4/5$

- $p \in [0, 1]$ if $q = 4/5$

Similarly, for player 2, since this is a symmetric game, so we see similar results.

We find that the mixed strategy nash equilibrium will be $(4/5, 4/5)$ where each player swerves $4/5$ of the time. We can check that this is indeed a mixed strategy nash equilbrium by showing that by keeping one player's strategy constant (i.e. $4/5$ of the time *Swerve* and $1/5$ of the time *Straight*) the other player has no incentive to change his strategy.

1.3   Iterated Games

To this point, we have only considered "one-shot" games where each pair of players will only play the game once. However, it is a natural extension to think of an iterated verson of each game where each game is played multiple times. We will focus mainly on the prisoner's dilemma game henceforth as that is the main game that we will be analyzing.

1.3.1   Iterated Prisoner's Dilemma (IPD)

Consider playing the prisoner's dilemma game multiple times. We saw previously that the Nash equilibrium for a one-shot game was for both players to play *Betray*. However, once we consider repeated versions of the game, things change. Let us assume that this is an infinitely repeated game, that is, we will play an infinite number of times (or for reality's sake, a large number of repetitions) [1]. Let us also introduce the notion of a "history," which will represent a history of plays that a

---

[1]We distinguish between and infinitely played game versus a finite (but large number of games). In the finite version, we can use backwards induction to show that the same strategy *Betray* will still be the best option for each player: consider the last iteration of the game - this is merely a one-shot game scenario again, and both players will tend to the proven Nash equilibrium. Now, we consider the game right before this last one, again, we can see that both players will play *Betray*. This logic extends forward and we see that the one-shot Nash equilibrium will be played at every stage for a finitely repeated prisoner's dilemma game. The infinitely repeated game has no such "last game," thus backwards induction cannot be used.

	Split	Steal
Split	$(\frac{S}{2}, \frac{S}{2})$	$(0, S)$
Steal	$(S, 0)$	$(0, 0)$

Table 1.3: Payoff diagram for the popular TV show *Golden Balls*.

player has done since the first game. Now, if we look at the IPD where players have access to histories, players may no longer choose to play *Betray* at every time step. As we will see in Chapter 2, Cooperation may start to emerge.

1.4 Games in Real Life

As mentioned before, game theory is used to model many real life situations in order to better calculate the best strategy to use. In politics, when a crucial vote is presented, game theory is an excellent model for voting and creating bills.

In international politics between countries, we saw that the late 20th century was a rough time between the US and Russia. These encounters and decisions are modeled using game theory.

Outside of such large-scale phenomena, one popular manifestation of the prisoner's dilemma is in game shows including *Golden Balls* and *Friend or Foe*.

In *Golden Balls*, two contestants make it to the final round where a sum of money is up for grabs (call it S). In this round, each contestant is presented with a pair of golden balls. Inside each ball is either the word "Split" or the word "Steal," and the players are allowed to know which ball is which for their own pair, but not for the opponent's. Then, each player will be allotted some time to consider their choice, and after this time is up, each will reveal his choice of "Split" or "Steal."

We can see the payoff diagram from 1.3. If both players choose "Split," then the total prize money, S, will be split evenly between the two [2]. If one chooses to "Steal" while the other "Splits," then the player that "Steals" will take home all the money. Finally, if both people try to "Steal," neither of them get any money.

---

[2]There is usually a very large variation in S, but most of the time it is at least \$20,000 and has been as high as \$175,000

We can map this game into the one-shot prisoner's dilemma game where "Split" corresponds to *Cooperate* and "Steal" corresponds to *Betray*. However, there is a slight twist now - there is no longer a unique Nash equilibrium. Any option other than both players choosing "Split" is now a Nash equilibrium. This is a weak form of our classic prisoner's dilemma where *Betray* was always the dominant strategy. The final twist to this game show is the time allotted to the players before they make their decisions. In this time, players are allowed to communicate with each other, and potentially reach a consensus on actions. Most times during the show, this period of time was spent on assuring the opponent that a player would choose "Split" and asking that they do the same. One famous conversation went as follows:

Contestant A: I will not lie to you, I plan on choosing "Steal." I would like for you to choose "Split." After the game ends, I will divide my winnings and give you half. I promise.

Contestant B: In that case, why don't we just both choose "Split"?

Contestant A: Because I am telling you, I will choose "Steal." I will not choose "Split." I encourage you to choose "Split" so I can give you half the money afterwards.

This back and forth went on for a while before the contestants were asked to show their choices. If we assume that Contestant A is telling the truth, then the clear choice for Contestant B is to choose "Split" since choosing "Steal" will most certainly end in him receiving 0. The show ended in both players choosing "Split."

Thus, despite (Split, Split) not being a nash Equilibrium, it was the produced outcome. In fact, in [28], the authors study this show and find that there are "unusually high rates of cooperation when the luck of the game reduces the stakes to 'merely' a few hundred Pounds." Thus, this brings in a new dimension to our analysis - the size of the stakes. In line with this finding, another interesting episode included two contestants vying for a sum of $150,000£$ (one of the largest sums in the history of the show). During the allotted time for discussion, both contestants vowed to choose "Split." Unfortunately, as the time came to reveal their choices, one contestant had veered away from her promise and chosen "Steal." She justified her choice with the simple statement, "I've been burned before." This alludes to the

importance of histories and previous play which we will expand on in coming sections.

Chapter 2

Game Theory & Iterated Prisoner's Dilemma

We have shown in Chapter 1 that the one-shot prisoner's dilemma has a unique pure strategy nash equilibrium. We then hinted at potentially different outcomes for iterative game play. In this chapter, we will explore some of the important findings regarding IPD, strategies for the IPD, and how well each strategy works.

2.1 The Iterated Prisoner's Dilemma − Revisited

Before we launch into the IPD, let us redefine the one-shot prisoner's dilemma a bit more rigorously. We mentioned before about the payoffs of the game, but in table 2.1 we show the general payoff scheme.

The following inequalities must hold from [23]:

	Cooperate	Betray
Cooperate	(λ_1, λ_1)	(λ_2, λ_3)
Betray	(λ_3, λ_2)	(λ_4, λ_4)

Table 2.1: Payoff diagram representing a general prisoner's dilemma game.

$$2\lambda_1 > \lambda_2 + \lambda_3$$

$$\lambda_3 > \lambda_1 > \lambda_4 > \lambda_2$$

The first rule is in place to prevent any tacit agreement for mutual cooperation. This is explained by [5] as a constraint such

> "that the players cannot get out of their dilemma by taking turns exploiting each other. This assumption means that an even chance of exploitation and being exploited is not as good an outcome for a player as mutual cooperation."

The second constraint exists such that *Betray* strictly dominates *Cooperate* as established in [17]

We see that our original nash Equilibrium still holds, and serves as a compelling option for players to choose. However, joint defection yields λ_4 to each player, whereas had they both cooperated, $\lambda_1 > \lambda_4$ could have been achieved. Thus, if the game is a single-shot game, *Betray* is the natural choice, but if the game is iterated over many trials, with access to histories, it is no longer so clear. Learning may take place as reputation comes into play, and players make seek to cooperate.

Thus, the IPD is a more complicated game, harder to predict the outcomes for, and the strategy may no longer be to solely *Betray*.

2.2 The Many Strategies of the IPD

In contrast to the one-shot prisoner's dilemma game, there are now numerous strategies that can be employed to play the IPD. Here, we will consider some of the more famous and influential strategies.

2.2.1 AllD

This is our well documented strategy of always defecting (betraying). One of the major upsides for playing AllD is the fact that you "can't lose." Indeed, by playing this strategy, you will never be hit with your opponent betraying while you cooperate. Furthermore, by playing this strategy, you can take advantage of very altruistic people that play *Cooperate* often. In fact, if you are matched up against someone that plays AllC, you will have the best strategy. The downside to playing AllD is that your opponent will very quickly lose trust in you. The "trust" variable, defined in [24], is the probability of responding cooperatively on the trial after both players have responded with defection. Since the AllD strategy only profits off of the opponent playing *Cooperate*, as the trust variable declines, the profitability of AllD declines as well.

2.2.2 Tit-for-Tat (TFT)

The Tit-for-Tat strategy is described as follows:

- In the first round, always play *Cooperate*

- For every round $t > 1$, play what the opponent played at time $t - 1$.

We can see that this strategy derives its name from its nature of reciprocity. The upsides of this strategy are quite apparent - if things are going well (each player is cooperating), then "let the good times roll." However, if the opponent ever defects, we need to punish them by defecting on the next round to send a message. In fact, this theme is prevalent in reality, sometimes going under the name of "Live and Let Live." For instance, during World War I, reciprocity of cooperation was prevalent based on the desire for fresh rations. In his account of army life in [14], Hay writes:

> It would be child's play to shell the road behind the enemy's trenchs, crowded as it must be with ration wagons and water carts, into a bloodstained wilderness...but on the whole there is silence. After all, if you

prevent your enemy from drawing his rations, his remedy is simple: He will prevent you from drawing yours.

This encapusulates the TFT strategy in the IPD as well. The downside to TFT is noise, or the accidental betrayal. Imagine two players both playing TFT - both will be playing *Cooperate* and doing well. Then, if one player accidentally switches to *Betray*, they will forever be locked in an alternating cycle of (Cooperate, Betray) and (Betray, Cooperate). From [5], we know that the payoffs for IPD discourage such outcomes by placing their value below the always cooperating outcomes. Finally, TFT, when faced with AllC, lacks the ability to exploit such play, instead of playing *Betray* every time, it will simply cooperate, which yields a good result, but not the best possible.

2.2.3 Pavlov (Win-stay, Lose-Shift)

As described in [21], the Pavlov strategy seeks to improve on the shortcomings of TFT. The strategy follows the following rule:

- Play *Cooperate* at time t if both players played the same alternative at time $t - 1$

- Else, *Betray*

This strategy gets its name from the fact that if the player receives either λ_1 or λ_3, it is considered a "win" and thus will repeat its previous move. If he receives either λ_2 or λ_4, it is a "loss," and thus will prompt a shift in choice of play.

We see that the Pavlov strategy contains two main advantages over TFT. First, it allows for the exploiting of the AllC strategy to maximize payoff. Second, it is less susceptible to noise. If both players utlize the Pavlov strategy and are cooperating, and one accidentally betrays, the next round both players will betray, but then will return back to (Cooperate, Cooperate) afterwards. This return to joint cooperation will never happen in TFT as seen above. The drawback to Pavlov comes when it

	Cooperate	Betray
Cooperate	(3, 3)	(0, 5)
Betray	(5, 0)	(1, 1)

Table 2.2: Payoff diagram from Robert Axelrod's tournament.

plays an AllD opponent. For every other round, Pavlov will submit *Cooperate* versus AllD's Betray.

2.2.4 Generous Tit-For-Tat (GTFT)

Generous Tit-For-Tat follows the same vein as the original TFT, but tries to improve upon the noise resiliance. This strategy will always be the same as TFT, but after an opponent betrays, there will be a probability $p > 0$ where it will play *Cooperate*. In the original TFT strategy, $p = 0$.

2.3 Seeking the Best Strategy

In 1980, Robert Axelrod held a tournament of various strategies for the IPD game with the payoffs documented in Table 2.2. Game theorests were invited to submit strategies that would then be run by computers against each other to determine the best strategy (as seen from total payoff). The strategies documented above as well as many more were entered. The results of the tournament were groundbreaking. The winner of Axelrod's tournament was the simple reciprocity strategy, Tit-For-Tat (TFT).

After these results were published, Axelrod invited submissions for a second tournament [1]. Having seen the results from the first tournament, many versions of TFT were submitted. However, the best performing entry of the 62 submissions was, once again, Tit-For-Tat. Interestingly enough, despite the possibility for contestants to submit strategies that were designed by others, there was only one submission of

---

[1]The second tournament was identical to the first except that for each round played, there was a probability, $p > 0$ that the game would end after that round. As we will see in the next couple sections, this design mitigates what is known as an "end-game" effect, briefly mentioned before in finite vs. infinite repeated games.

TFT, by its original creator, Anatol Rapoport.

After the second tournament, in [5], Axelrod talks about four factors that determined a winning strategy, namely:

- **Niceness:** Avoidance of unnecessary conflict by cooperating as long as the other player does

- **Provocability:** Getting mad quickly and retaliating in the face of an uncalled-for defection by the other

- **Forgiveness:** After responding to a provocation, don't hold a grudge

- **Clarity:** Act so that the other player can recognize and adapt to your pattern

We see that TFT satisfies these criteria very nicely. It starts by playing *Cooperate*, and plays as long as the opponent cooperates as well - achieving niceness. It responds quickly and retaliates against an opponent that plays *Betray* as it betrays in the subsequent round. It forgives in the sense that once it retaliates, if the opponent playes *Cooperate* again, TFT will revert back to cooperating. Finally, it is faily clear when someone is playing this strategy due to is simplicity. Thus, it is not a huge suprise that TFT faired as well as it did both times.

So, does this mean that TFT is **the best** strategy? Unfortunately, the answer is a resounding **no**.

For instance, in [21], Nowak argues that Pavlov is a more successful strategy than TFT, and backs up his claim with simulation results. This makes some sense given that the Pavlov strategy seeks to patch up some of the shortcomings of TFT, see section 2.2.3.

In [6], the authors demonstrate a strategy known as *Gradual*, which they show to outperform both TFT and Pavlov in a 12-strategy round robin tournament. *Gradual* performs exactly as TFT except when it comes to the opponent playing *Betray*. On the first betrayal at time t, *Gradual* will betray once then cooperate twice. After the second betrayal at time $t' > t$, *Gradual* will betray twice then cooperate twice. After the n^{th} betrayal by the opponent, *Gradual* will betray n times, then cooperate

twice to re-establish cooperative grounds. We can see that this strategy has all the same qualities as TFT except for its simplicity - *Gradual* requires memory since the beginning of the game.

Thus, in both [6, 21], strategies are presented that defeat TFT, even though in Axelrod's original tournament TFT emerges the victor over both of these strategies. This conundrum captures the essence of a strategy's robustness. In [22], the authors define a strategy, σ to be "robust," if in *any* population consisting of copies of two types of strategies, σ itself and some other strategy τ, the strategy σ is never worst than τ. The authors show that *no* strategy exhibits robustness in all population sets of three or more strategy types.

Finally, in [2], even Axelrod himself gives three reasons for why TFT is not *the* best strategy:

- TFT could have been beaten in the second round by any rule which was able to identify and never cooperate with *Random*, while not mistaking other rules for *Random*

- Had only the entries which ranked in the top half been present, TFT would have come in fourth place, after the ones that came in 25th, 16th, and 8th.

- There is no best rule independent of the environment.

Despite not reaching a clear *best* strategy for the IPD, we have discovered a set of rules that govern a successful strategy, [5]. However, even if we follow these rules and design a strategy around them (be it TFT, Pavlov, or Gradual), we are still highly dependent on the opponent's actions for our success. If the opponent plays AllD, then regardless of our strategy, we will never do well. Thus, in order to be successful, the opponent needs to offer some cooperation. This leads us to the next widely studied question of how to achieve cooperation.

31

2.4 Achieving Cooperation

Despite the temptation of betrayal (higher payoff), [7, 16] show the perpetual shift toward a cooperation regime. We now take a look at why cooperation amongst individuals is so high, and what factors affect mutual cooperation.

In [4], Axelrod emphasizes that there are three elements to look at for determining whether cooperation would emerge: Payoff Structure, Future of Gameplay, and Number of Actors.

2.4.1 Payoff Structure

The importance of the payoff structure is emphasized in [19, 27]. We saw from 2.1 and [23] that certain inequalities had to hold for the prisoner's dilemma payoffs. However, in [27], we look at scaling effects on the payoffs. They argue that with a high multiplicative factor, reasons for cooperation increase, whereas a less than one multiplicative factor does not promote cooperation. One possible explanation for behavior linked to payoff scaling is the risk appetite of the competitors. In [8], they look at risk-averse versus risk-loving people and see how they respond to different payoff amounts [2]. Unfortunately, their results do not draw a significant relationship between risk-appetite and behavior in the prisoner's dilemma game. Despite these results, it is not unreasonable to expect a correlation between payoff amounts and cooperation level. We will look at this in the next section pertaining to game design.

2.4.2 Future of Gameplay

Many studies focus on the possibility and certainty of future game play, i.e. the probability that a game will end after any given round. In [4], Axelrod refers to this as "The Shadow of the Future," and presents four factors promoting cooperation:

- Long time horizons

- Regularity of stakes

---

[2]They judge people's risk appetites based on their attitudes over a series of lotteries

32

- Reliability of information regarding others' actions

- Quick feedback about changes in others' actions

The first of these elements is very crucial - if contestants know that they will have future interactions with their opponent, they are more likely to cooperate to maintain good relations. In [19, 10], the importance of the possibility of continued game play is emphasized. Both studies find that as the probability of continuation increases, the probability of cooperation increases as well. In fact, [10] found that "expected iteration decreased retaliation." Thus, not only does the probability of cooperation increase with the probability of a next iteration, but the probability of retaliation also decreases. This may stem from a player assuming that a mistake may have been made by the opponent, and is more willing to forgive.

The regularity of stakes allows for steady game play without abnormal jumps. Consider two stakes that are vastly different - it may be conceivable that a player will cooperate in the lower stakes game. Then, once this trust is established, he may turn around and betray in the higher stakes game - and dependent on the payoff multiple, this may be a viable and profitable strategy.

The last two points highlight the transparency aspect of the IPD in that players should be informed of how the opponent is playing in order for trust to build between them. If one player is always surprised by the opponent's game play, then this does not lend well to building cooperation between them - since he will never know when the opponent may just choose to betray.

2.4.3 Number of Actors

As shown in Axelrod's tournaments, reciprocity was a very effective strategy. However, for this to be possible, players need to be able to identify defectors, focus on retaliation, and have long-run incentives to punish defectors. As we increase the number of actors, this becomes more difficult. As [4] mentions - none of the cooperators has the incentive to play the role of policeman and are content to be a free-rider on

the willingness of others to enforce the rules. This difficulty of preventing defection through decentralized retaliaton is known as the "sanctioning problem."

Thus, in a large world of betrayers, how can cooperation come about?

To answer this exact question, in [3], Axelrod looks at the concept of invasion of strategies. He defines a strategy A as invading a strategy B if $V(A|B) > V(B|B)$, that is to say the expected payoff an A gets from playing a B is greater than that from a B playing a B. This ties in closely with the evolutionary approach, namely, the concept of invasion is parallel to the concept of a single mutant individual being able to do better than the population average. He then goes on to show that TFT can successfully invade AllD as long as there is at least a small portion of TFT's interacting with each other. His results suggest that we only need around 5% of TFT to invade a society of AllD.

Chapter 3

Our Game – Designs and Considerations

In this chapter, we explore the design of our IPD game as well as the questions that we seek to answer. The first section will talk about the environment that we have. The second and third section will look at the front-end design of our game and the backend design. The final section will present the questions that we would like to answer, questions that have either been theoretically (through simulation) shown proven as well as some phenomena that come up in the social setting we utilize to test the IPD.

3.1 The 21st Century as a Testing Ground

The 21st Century is a ripe testing ground for theories all around, especially economic theories. The world has never before been so connected, where opinions can travel around so quickly and so far. Take for instance the field of politics, shrouded in secrecy just three decades ago. With the advent of the internet and information sharing, not only are people learning about politics, they are engaging in discussions. In [9, 26, 1], the authors discuss how the popular site Intrade seeks to predict polls and elections. On Intrade, people are allowed to buy and sell contracts whose prices are governed by the probability that a certain party will win the election. Some view

this as a viable proxy for the outcome of elections.

Another widespread network is the Twitter network - connecting millions of people monitoring each other's 140-character status updates, or "tweets." To get a sense of how important these networks and social bindings are to us today, we witnessed the power of a single tweet. In April of 2013, the Twitter account of the Associated Press was hacked, and the following tweet was broadcast through the network:

> Breaking: Two Explosions in the White House and Barrack Obama is injured

To see the ripple effect of such a post, the Dow Jones immediately dropped 100 points, nearly 1% of its value.

Finally, the largest network today, Facebook commands over 1 billion users, connected in a network of friends. It is in this network that we seek to study the empirical results of running the IPD game. The Facebook network serves as an excellent source for testing as users are allowed to play each other, and we can obtain a number of details regarding the user.

3.2 Game Design (rev. 1)

We sought to implement a classic version of the IPD whereby Facebook users would be allowed to play each other and compete against each other. We also included a chicken game as well. In this section, we will look at some of the independent variables we considered when designing the game, and the reasons behind the choices that we made.

3.2.1 Stakes & Scaling

As we mentioned before in 2.4.1, we saw that [19, 27] emphasize the importance of the payoff structure of the IPD game. Axelrod argues in [4] that payoff structures is one of the three elements that would determine whether or not cooperation could come about.

From [11], the authors argue that from their empirical experiments, people only learned to cooperate with $\frac{\lambda_3}{-\lambda_2} > 2.5$, and they would learn faster for $\frac{\lambda_3}{-\lambda_2} > 4$ (and have a higher forgiveness rate). Thus, for our stakes, we decided to implement two scenarios where the ratio of $\frac{\lambda_3}{\lambda_2}$ was 2.5 and two scenarios where the ratio was 4. For the pair that had the ratio of 2.5, the two scenarios differed in that one scenario's payoffs were scaled upward by a multiple of four from the other one. For the pair of scenarios with a ratio of 4, the scaling factor on payoffs was 2.

By having these separate scenarios, we would be able to look at both how payoffs and the $\frac{\lambda_3}{\lambda_2}$ ratio affected decisions.

3.2.2 Friends

The Facebook network is built on the concept of "friends." Each user is connected to his or her friends, and that expands the network outwards. Presumably, one would only be "friends" with someone else on Facebook if they actually knew each other in real life. Thus, we decided that it would be interesting to see the social dynamic of how we interface with our "friends" via the IPD game. In [18], they explore the aspects of human friendship and how that affects cooperation in the prisoner's dilemma. They conclude that humans tend to cooperate more with friends.

Unfortunately, simply being a friend with someone on Facebook does not guarantee that the pair are friends in reality (a simple look through this author's Facebook friends will reveal that he still interfaces with less than a quarter of them). Thus, to serve as a proxy for friendship, we looked at how many mutual friends two people shared. This perhaps better captures the friendship of two people. If two people are friends in real life, there is a high probability that they will share a group of friends (reflected via number of mutual friends).

3.2.3 User Qualities

Similar to the friendship information, Facebook keeps tabs on a number of personal details that we will access and record. These details include:

- gender

- political affiliation

- religious affiliation

- age

We will focus on both age and gender as everyone with a Facebook account is required to have those fields. By including these options, we can explore interactions between users who have different qualities and test for their significance. For instance, we can see whether or not males are more likely to cooperate with fellow males or females.

3.3   Backstory & Frontend

In this section, we will look at the backstory of the game as well as decisions governing the design and implementation of the frontend. The ordering will closely follow the layout of the game itself.

3.3.1   Homepage

As we thought about how to frame our game's backstory, we sought a bit of inspiration from [15], where Kaminski describes the innerworkings of Polish prisons. Thus, as we were going to test the IPD game in a social setting, we figured that we might as well make the backstory one about serving time in prison.

Therefore, we decided to craft the story around the user being sent to prison and trying to get out in as little time as possible. We required users to log in via Facebook before they were allowed to play such that we would have access to their information. Appendix B-1 shows the homepage that we created to our game.

Each player would start with 250 weeks of prison time. Then for each game that they played, dependent on the outcome of the game, they could potentially lose time off of their sentence. The goal was to reach 0 weeks of prison time the quickest by

playing the least number of games (time spent in prison is equivalent to the number of games a user plays).

3.3.2 Stages

Upon signing into their Facebook account, users were then presented with the "Stages" page where they could choose which "stage" they wanted to play on, seen in B-2. As we talked about earlier in 3.2.1, different stakes and payoffs required different stages that users could play on. We created six stages, aptly named for a prison backdrop:

- Stage 1 → Dining Hall

- Stage 2 → Kitchen

- Stage 3 → Cell

- Stage 4 → Yard

- Stage 5 → Hospital

- Stage 6 → Office (Interrogation Room)

Each stage represented an independent game with different payoffs, seen in figs. B-3 to B-8. The dining hall allowed for players to get familiar with the interface and play with an AI without affecting their prison time. The yard was a chicken game and the other four were IPD games with varying payoffs.

The payoff diagrams for each stage (other than the dining hall) can be seen in tables A.1 to A.5. Note that the hospital and the office represent the "high-stakes" IPD games while the kitchen and the cell have lower stakes.

3.3.3 Game Information

There are two additonal informational pages, one to show the user's own statistics for the games that he has played (B-9), and the other to view the leaderboard (B-10). The statistics page allows for users to see how many times they have played

Cooperate versus the times they have played *Betray*. Furthermore, they can see these same statistics for how opponents have played against them (i.e. how many times they have been cooperated with or betrayed).

The leaderboard page lets users see how well they are doing compared to others. Users can rank people by time left or by how quickly other users have gotten out of prison. We implemented this to increase the competitive level amongst users playing our game. When users complete the game (get out of prison), their prison time resets to 250 weeks and the second column shows how many weeks they required to get out of prison (games played). This helps to encourage users to continue playing to try and beat their previous record.

3.4 Game Mechanics

We tried to make sure that the game was as easy and smooth to play as possible. Thus, to play a game a user will:

1. Login on our site

2. Select stage they wish to play on

3. Challenge a friend or select Challenge random user

4. Input their strategy

This will send a challenge (or open game) to their opponent so that when he logs in to our game, he will see that he has an outstanding game. He can then go to the stage where the chalelnege was issued and click on Betray or Cooperate. Immediately, he will be able to see the results of that game while the original challenger will see it the next time he logs on.

To promote people to play, we allow for users to invite their Facebook friends directly from our website. Furthermore, when a user issues a challenge (starts a game with) an opponent, a Facebook notification is sent to the opponent to remind them to play.

3.5 Additional Considerations

One of the items we wrestled with during the design and construction of the game was how to label the choices each player faced when starting or responding to a game. In [29] the authors explore how the labeling of users' choices in the prisoner's dilemma game can influence a player's decision on whether to cooperate or not. They compared situations when the labels were the researchers' conventional vocabulary (such as "cooperate" and "defect") versus non-representational labels (such as "a" and "b"). They found that when users were exposed to the conventional vocabulary used by researchers, the cases of cooperation increased.

When we were designing our game, we wanted to establish a connection between the researcher's interpertation of the option and the presentation to the players. In the end, we decided to label our strategies as we would normally refer to them: Betray and Cooperate.

Another item that we discussed thoroughly was whether or not to show player histories during games. This is definitely a consideration especially since it may affect a user's decision. In [13], the authors conclude that with more memory (i.e. a longer history), players were more likely to cooperate since opponents would be able to see their past behavior. For our game, we decided on a hybrid approach to try to confirm this result. We would randomly allow roughly half of the games access to user histories [1]. For these games, both users would be able to see up to the past five actions of their opponent on this level.

Along with histories, players would be able to view how much time their opponent had left in prison. This option allows us to look at actions by and against users with little time left. In [25], Rapoport explores start and end effects for the IPD game. He concludes that when the probability the existance of another round diminishes, so does the probability of cooperation. Thus, for players with little prison time

---

[1]We accomplished this by seeing if the parity of each user's user id (sequentially generated on order of signup) matched. If they did, then each user would be allowed access to histories. We chose this strategy over just randomly allowing for history access since that approach would be potentially exploitable - constant refreshing of the page would eventually yield a user access to their opponent's history
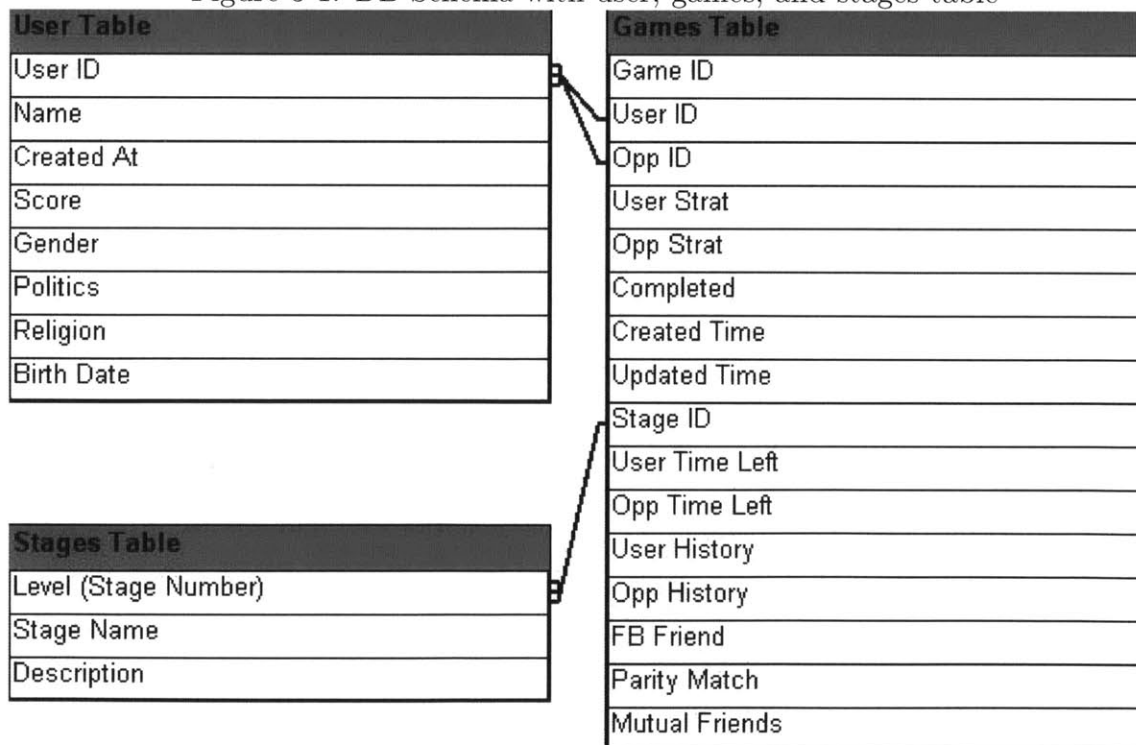
remaining, it is possible to see how this results plays out. Furthermore, we can also see how players respond to an opponent with little time left.

The final consideration was the aspect of communication between players. The classic IPD game did not allow for communication between the two players. In [12], the authors study the impact of allowing players to send messages and promises to each other during the IPD. For our game design, we chose to not allow for user communication within the game itself - though there is not much that we can do to stop users from communicating outside of the game environment.

3.6 Backend Design

To keep track of the all of the data and game play, we designed a three table modular database as seen in 3-1.

Figure 3-1: DB Schema with user, games, and stages table

User Table	Games Table
User ID	Game ID
Name	User ID
Created At	Opp ID
Score	User Strat
Gender	Opp Strat
Politics	Completed
Religion	Created Time
Birth Date	Updated Time
	Stage ID
	User Time Left
	Opp Time Left
Stages Table	User History
Level (Stage Number)	Opp History
Stage Name	FB Friend
Description	Parity Match
	Mutual Friends

The User Table allows us to obtain information on each user as they sign up for

our game via Facebook. For each new user, a new line will be generated in this table with all the required fields (if the user provided them on Facebook).

For each level created, we update the Stages Table. Level corresponds to a number (1-6) that represents the stage.

The Games Table keeps track of every game that is created and played. Below is a breakdown of some of the important fields:

- User ID - Identifies which user is playing (this user is the one being challenged)

- Opp ID - Identifies the user that is the challenger (creator of the game)

- User Strat - Strategy played by the user

- Opp Strat - Strategy played by the challenger

- Completed - Whether or not the game has been completed (i.e. played by both parties)

- Created Time - Time when the game was created

- Updated Time - Time when the game was played by the opponent (user that was challenged)

- User/Opp History - Percentage of last 5 games that *Betray* was played

- FB Friend - Whether or not the User and Opp are friends on Facebook

- Parity Match - Whether or not to show player histories (see 3.5)

- Mutual Friends - Number of mutual friends on Facebook

With these three tables, we can piece together the game play history for different players and look at their strategies across time, players, etc.

43

3.7 Language and Implementation

We used Ruby on Rails to create the application and currently utilize PostgreSQL as our backend system. Our application is hosted on Heroku and is accessible via: prisoners-dilemma.herokuapp.com

3.8 Game Design (rev. 2)

After launching our application, we received over 1000 user signups and over 3000 completed games. Despite these promising results, we still noticed that there was an overall waning interest in continuing to play the game as well as a number of users that signed up but never played a single game. Thus, we decided to improve the game for a second launch, with changes geared towards more intuitive game play for new users as well as incentives to keep users playing.

3.8.1 Jumpstarting New Players

With over 400 signups that never played a game, we decided that the largest barrier to entry was getting the first game started for the user. Essentially, these 400 users would sign in, try and start a game unsuccessfully, and then give up. To solve this, we implemented two new features - game recommendations and an AI masquerading as a regular player.

Game Recommendations

Once a user signs in and goes to a stage, there will be three recommended players to play against. As seen in fig. B-11, a user will be provided two friends and one random user to challenge - he simply needs to press Betray or Cooperate, and the game begins. This is a vast improvement from before where a user needed to figure out how to start a game - now all it takes is the click of a button.

AI

We have also implemented a bot that plays *TFT* on Stage 2 and *Gradual* on Stage 3. We have it such that if a user does not have any outstanding games waiting for

him to finish, the bot will start one against him - thus providing a game against new users once they log in.

We hope that these two new features will help to draw in new users and promote game play.

3.8.2 Incentives to Continue Playing

In our first release, we had a number of users that played around 30 games, got released from prison, and then never played again. However, as we would like users to play as many games as possible, we implemented two new changes to try to keep users in the game longer. First, we decided to increase the jail time to 350 weeks. This would naturally promote more game play for users seeking to get out prison. Second, we decided to offer a (monetary) prize for the top four users (with the quickest prison release times). With more than just bragging rights on the line, users will play more games and play more carefully.

Beyond just Facebook users, we also added in the capability to distribute our game to Amazon Mechanical Turk (MTurk) players. These players would be able to go through a 350-week cycle, and upon finishing would receive a small monetary compensation. Utilizing MTurk allows for us to have more users play consistently. This will hopefully result in more games generated on a consistent basis and keep users engaged beyond their first prison cycle.

3.8.3 Other Changes

Beyond these changes, we once again brought up the topic of showing user histories. Our original approach of randomizing when users saw their opponents' histories seemed, in retrospect, too opaque for users. We often received questions of why some user histories were shown while others were not. Furthermore, it was not clear to users whether or not their history would be shown against a certain opponent in the future if it wasn't shown currently (i.e. will my current game play affect my future reputation against the same opponent). Due to these confusions (see the Analysis

section below for why this was detrimental to our results), we decided to allow users to always see their opponent's cooperation rates over the past few games (against all users). This hones in on the aspect of reputation and makes it a key factor in decision-making for players.

A final change we made for Rev. 2 was allowing game request receivers to see their oppoenent's recent game history against them. This may result in a shift towards more cooperators starting games - we shall see once we gather the data.

Chapter 4

Questions and Hypotheses

With the plethora of data that the Facebook network provides, there are many interesting questions that we can ask. In this chapter, we look at questions and hypotheses that we can form. The next chapter will seek to analyze the data we have collected and discuss the veracity of our hypotheses.

4.1   Some Basic Questions and Hypotheses

Overall Play

The first question that comes to mind is how people played? Specifically, did people betray or cooperate more in a general sense? When thinking about this, we consider that if this game were memoryless, then as the authors in [13, 11] wrote, *AllD* would be a dominant strategy. However, since we display histories for roughly 50% of the games played along with the ability for users to remember their opponents' names (from Facebook), our IPD is not memoryless. In [3, 16], the authors argue that cooperation exists and grows over time.

We hypothesized that more people would play *Cooperate* given both the existence of histories as well as human forgiveness. People may come to realize that to achieve the long run success, cooperation may be necessary.

We can then look at the best overall strategy - i.e. what did the best-performing player do in his gameplay?

Factors Affecting Decisions

From Facebook, we are able to harvest a number of attributes for each player, and it is possible to consider which of these factors are statistically significant in swaying a player's decision.

History - We believe that the higher the betrayal rate of an opponent, the more likely a user will betray that person.

Stage (Payoffs) - We believe that the higher the payoffs, the more likely betrayal will occur. We refer to the "raw" payoff and not the ratio of the payoffs. This would seem probable because a betrayal (versus an opponent's cooperation) would yield a significant payoff.

Friends - We believe that being friends (and having more mutual friends) would tend to promote cooperation as well as forgiveness.

Gender - We think that playing against a female would promote cooperation whereas playing against a male may lead to more betrayals. The generalization of females being "kinder" might suggest this hypothesis.

Age - We think that perhaps older people would tend to cooperate more, as they are more trusting.

4.2 Some More Advanced Questions

Beyond the simple questions, we can also consider some more advanced approaches to understanding how people played.

Pair plays

For instance, for each pair of people, over extended iterations, what strategy does each player play, and is it dependent on the aforementioned factors?

Between pairs of players, we can also look at the string of plays each does - are there longs strings of C's or D's?

Chapter 5

Results and Analysis

In this chapter, we will present the data received from our first release (rev. 1). We will then seek to answer the questions posed above and draw some conclusions.

5.1 Results

The following is a snapshot of the statistics on users and games played during our first release:

Games

We ran our application on Facebook for around three weeks. In that time, we had 12,000 games started, 6,000 finished games, and 3,500 finished games played between real players (2,500 were against a tutorial bot).

Users

In this time, we had nearly 800 users sign up. 26% were female and 74% male. The average age of a user was just under 30 years old with a standard deviation of 12 years.

5.2 Analysis

5.2.1 Overall Analysis

In the previous section, we raised the question of how overall game play would look. Our data on game play for completed games is shown in 5-1.

Figure 5-1: How users play over time
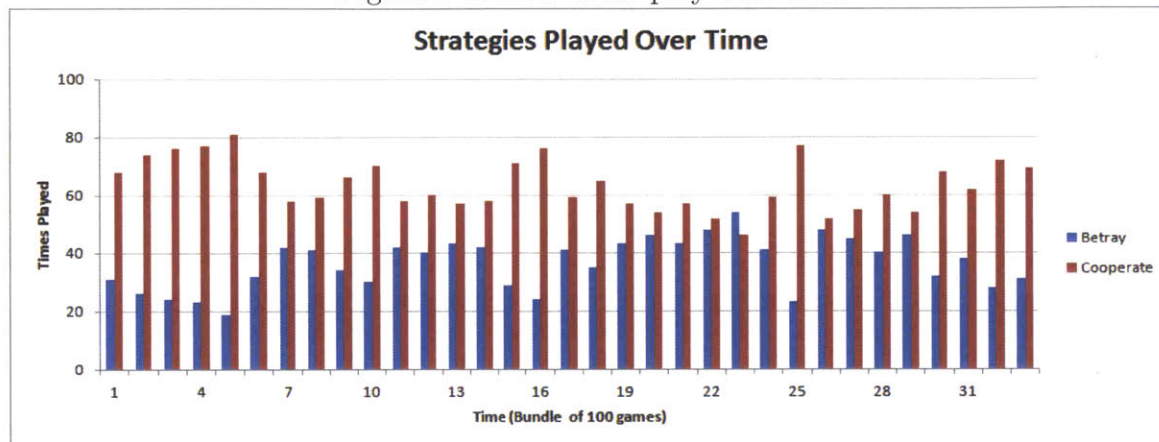
The graph bundles games into 100-game segments and orders them according to time. We can see that there is an overwhelming number of *Cooperates* as compared to *Betrays.* In fact, there is only one set of 100 where the betray option ever overtakes cooperate, and a slim margin at that. From this graph, we can also see that in the beginning (first 600 completed games), cooperation levels were very high. Thus, when people started playing, a "start effect" as discussed in [25].

A more conclusive calculation yielded (for all IPD games):

$$Pr(\text{Opp Betrays}|\text{Opp's First Game}) = 0.38$$
$$Pr(\text{User Betrays}|\text{User's First Game}) = 0.36$$

These results align nicely with our graphical representation - we see that there is a majority of users that cooperate on their first move (both starting and answering games).

5.2.2 Factors Affecting Decisions

In the previous section, we listed a number of factors that may affect how players make decisions. Here, we delve into our dataset and see which factors were significant, and how they affected the outcome.

First, we created two regressions where the dependent variables are the opponent's strategy (game starter) and the user's strategy (responder). We took all of the factors mentioned above and set them as independent variables to see how significant they were.

- User Strategy → *user_strat_num* where 1 = Betray, 0 = Cooperate

- Opponent Strategy → *opp_strat_num* where 1 = Betray, 0 = Cooperate

- User Gender → *user_gender*

- User Age → *user_age*

- Opponent Gender → *opp_gender*

- Opponent Age → *opp_age*

- Facebook Friends or Not → *fb_friend_num* where 1 = True, 0 = False

- Number of Mutual Friends → *mutual_friends*

- Maximum Reward of a Stage → *max_reward*

- Cost Benefit Ratio (see section 3.2.1) → *cdratio*

The regressions run (in STATA):

1. regress user_strat_num user_gender user_age opp_gender opp_age fb_friend_num mutual_friends max_reward cdratio, robust

2. regress opp_strat_num user_gender user_age opp_gender opp_age fb_friend_num mutual_friends max_reward cdratio, robust

Table 5.1: User Strategy Regression

	user_strat_num	t-statistic
user_gender	-0.0398	(-1.85)
user_age	-0.00263**	(-3.11)
opp_gender	0.0668*	(2.47)
opp_age	0.000868	(0.55)
fb_friend_num	-0.104	(-1.93)
mutual_friends	-0.000306	(-0.70)
max_reward	0.00647***	(3.72)
cdratio	-0.0295*	(-2.11)
_cons	0.524***	(7.22)
N	2213	
t statistics in parentheses		
=*$p < 0.05$	**$p < 0.01$	***$p < 0.001$

Table 5.2: Opponent Strategy Regression

	opp_strat_num	t-statistic
user_gender	-0.113***	(-6.52)
user_age	-0.000587	(-0.78)
opp_gender	0.251***	(9.51)
opp_age	0.00220	(1.51)
fb_friend_num	-0.00478	(-0.09)
mutual_friends	0.000632	(1.36)
max_reward	0.00712***	(4.99)
cdratio	-0.0459***	(-4.10)
_cons	0.243***	(3.97)
N	2213	
t statistics in parentheses		
=*$p < 0.05$	**$p < 0.01$	***$p < 0.001$

History

The two regressions results are seen in 5.1 and 5.2, respectively.

We use max_reward and cdratio as proxies for stages instead of using the level number (since larger doesn't necessarily mean a higher payoff).

It seems that different factors come into play depending on whether or not a user is starting a game (opp_strat) or finishing a game (user_strat). However, shared between the two are gender considerations, stage considerations, and a lack of Facebook friend considerations.

5.2.3   Gender

From both 5.1 and 5.2, we see that gender plays a significant role. In particular for the game starter (Opp game play):

$$Pr(\text{Betray}|Opp = Female) > Pr(\text{Betray}|Opp = Male)$$
$$Pr(\text{Betray}|User = Female) < Pr(\text{Betray}|User = Male)$$

And for the game finisher (User game play):

$$Pr(\text{Betray}|Opp = Female) > Pr(\text{Betray}|Opp = Male)$$
$$Pr(\text{Betray}|User = Female) < Pr(\text{Betray}|User = Male)$$

In words, this means:

- When responding to games, players choose to betray more when the opponent (game starter) is female rather than male

- When responding to games, males chose to betray more than females

- A female starting a game will more likely choose to betray than a male starting a game

- Players betray less when starting games against females than against males

We look more closely in our data, and see the following:

Females starting games chose to betray 32.6% of the time, whereas males only chose to betray 19.3% of the time. In responding to games, males betrayed 44.1% of the time versus only 29.2% of the time for female respondents.

From 5-2, we see that females tend to cooperate more (77%) versus males (71.5%):

Figure 5-2: Male vs. Female Cooperation Statistics
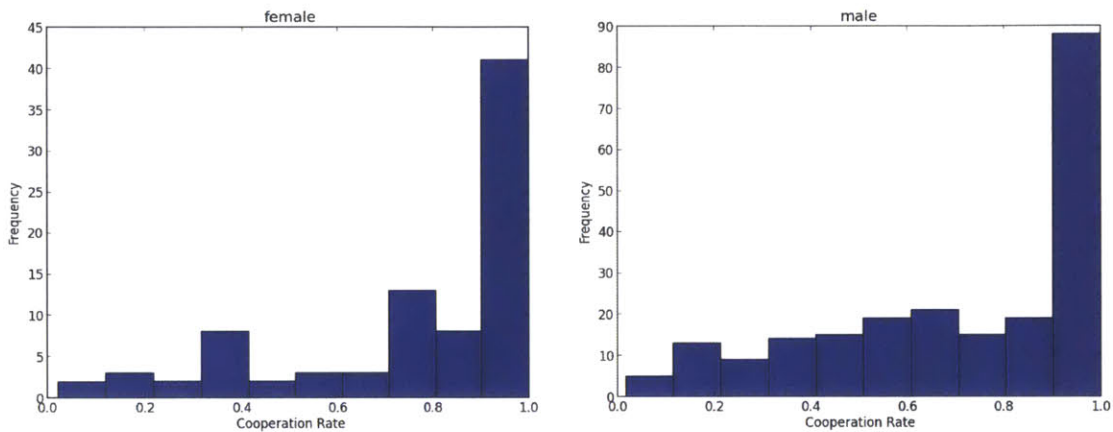
If we disegard who start's the game, Table 5.3 shows overall gender vs. gender cooperation statistics:

Table 5.3: Cross gender cooperation dynamics

Player Gender	Opponent Gender	Player Cooperation Rate
Male	Male	0.676
Male	Female	0.78
Female	Male	0.677
Female	Female	0.79

We see that when playing a female, people cooperate more than if they play against a male. The highest cooperation rates occured when females played females, while the lowest occured when two males played against each other. The two middle rows suggest that males tend to cooperate more with females while females betray males much more.

5.2.4 Stage & Payoffs

We saw from our regressions that the stage (proxied via max payoff and *cdratio*) was statistically significant in affecting a player's decision. The experiments from [11] showed that a higher $\frac{\lambda_3}{-\lambda_2}$ ratio promoted more cooperation. Since it is not a memoryless game, this would make sense. However, for stages with higher payoffs, it would seem that the temptation of betrayal would be much higher (as the potential gain would be greater).

Our data yields Table 5.4:

Table 5.4: Cooperation rates by stage

	Cooperation Rate	Games Played	*cdratio*
Stage 2	0.725	763	2.5
Stage 3	0.746	574	4
Stage 4	0.798	630	-
Stage 5	0.576	816	2.5
Stage 6	0.673	766	4

Our data lines up very nicely to the results presented in [11]. Between the two lowest (max payoff) stages, 2 and 3, stage 3 had a higher cooperation rate, correlating with a higher *cdratio*. Likewise, for the two highest payoff stages, 5 and 6, stage 6 has a higher cooperation rate, while maintaining a higher *cdratio*. The highest cooperation rate came on stage 4, the chicken game, which has the highest *cdratio*. Thus, it makes sense that in this non-memoryless setting, people are less willing to be betray on that stage.

Another observation we see here is the effect of max payoff on betrayal rates. Comparing stages with the same *cdratio*, we see that the stage with a higher max payoff (i.e. higher temptation for betrayal) has a lower rate of cooperation.

5.2.5 Facebook Friend Status

One of the major benefits of utilizing the Facebook network as our proliferation method is the ability to gather information about users and their opponents - one

aspect being whether or not two users are friends. We used two metrics to determine friendship between players:

1. Whether or not the two players were friends on Facebook

2. The number of mutual friends between two players on Facebook

We believed that the second option was actually a better proxy for how close two people were since the criteria for "friending" someone on Facebook is not terribly demanding.

However, as tables 5.1 and 5.2 show, these factors were not significant in a user's game play. If we only look at mutual friends between players, we can then set different thresholds (number of mutual friends) to see cooperation rates when the pair of players share more mutual friends than the threshold versus when they share less than the threshold.

Table 5.5: Effect of different thresholds of mutual friends on cooperation rates

	Cooperation Rates		
Threshold	Above Threshold	Below Threshold	Difference
1	0.716	0.668	0.048
3	0.703	0.67	0.033
5	0.71	0.686	0.024
10	0.712	0.687	0.025
20	0.7	0.693	0.007
50	0.704	0.693	0.011
100	0.655	0.698	-0.043

From Table 5.5, we can confirm the statistical insignificance of friendship on cooperation rates. For the most part - having mutual friends leads to a higher cooperation rate - however, it is not true that having more mutual friends would promote higher cooperation rates.

5.2.6   Histories

The original topic of whether or not to show histories for players proved to be a difficult decision. In the end, for rev. 1, we decided to "randomly" show histories

56

between players - at least hoping it woud look random to the users. We would then be able to see how viewing history would affect players' decisions.

We did a stage by stage analysis of how showing histories (and not the history itself) would affect user choice.

Table 5.6: Stage by stage cooperation rates with vs. without history shown

	Without History		With History	
	Cooperation Rate	Games Played	Cooperation Rate	Games Played
Stage 2	**0.754**	374	0.698	389
Stage 3	0.73	283	**0.763**	291
Stage 4	0.775	327	**0.823**	303
Stage 5	0.569	403	**0.582**	413
Stage 6	0.649	396	**0.7**	370

Table 5.6 shows that there is a slight edge towards cooperation when histories are shown (except for stage 2).

If we restrict our data set to only when histories are shown, we can determine the significance of the history itself (i.e. if a user has betrayed 20% or 80% in their last five games).

Table 5.7: Regression on user strategy when histories are shown

	user_strat_num	t-statistic
user_gender	0.00239	(0.08)
user_age	-0.00226*	(-2.29)
opp_gender	0.0493	(1.41)
opp_age	-0.000151	(-0.09)
fb_friend_num	-0.129	(-1.59)
mutual_friends	-0.000210	(-0.28)
max_reward	0.00438	(1.82)
cdratio	0.00731	(0.40)
user_history	**0.00523*****	**(16.07)**
opp_history	**0.000672**	**(1.62)**
_cons	0.303**	(3.13)
N	1109	
t statistics in parentheses		
=*$p < 0.05$	**$p < 0.01$	***$p < 0.001$

57

Table 5.8: Regression on opponent strategy when histories are shown

	opp_strat_num	t-statistic
user_gender	-0.00837	(-0.43)
user_age	-0.000373	(-0.68)
opp_gender	0.120***	(3.66)
opp_age	-0.000448	(-0.52)
fb_friend_num	-0.0870	(-1.14)
mutual_friends	0.000644	(0.92)
max_reward	-0.00224	(-1.37)
cdratio	-0.00762	(-0.61)
user_history	**0.000697***	**(2.56)**
opp_history	**0.00802***	**(27.88)**
_cons	0.121	(1.91)
N	1109	
t statistics in parentheses		
$=^*p < 0.05$	$^{**}p < 0.01$	$^{***}p < 0.001$

From tables 5.7 and 5.8, we can see that when histories are shown, users (respondants to games) are heavily influenced by their own past actions, while opponents (game starters) are influenced by their own past actions. The tables suggest that players are not swayed much by their counterparty's histories, and hold steady to their own patterns.

Shortcomings

Despite these statistically significant results, they are not 100% straightforward. There necessarily exists some omitted variable bias in these results. For one, users were not informed on when histories were shown (hoping that they would assume this system was random). Thus, a user could not draw the conclusion that if a game against a player revealed that player's history, then the next game against the same player would also show histories. Furthermore, players were not informed that history-showing was symmetric, i.e. if I can see your history, you can see mine. Therefore, when players decide whether or not to betray, given that a history is shown, the "reputation" effect comes into play. Because of the lack of information to players regarding histories, different players proabably reached different conclusions - and thus reacted differently to the "reputation" effect.

5.2.7 Other Ineffective Factors

The Facebook network provided us a number of other fields to record as well, including religion, education, etc. However, these fields on Facebook are not required, and thus we did not have this information on a majority of people. The other field that we did know for everyone was the age field. As shown in tables 5.1 and 5.2, a player's age was not significant in determining game play.

5.3 Pairs

We can look at repeated interactions between pairs of people. Due to the fact that we have multiple levels and players don't play at the same time, the chronological order of players' decisions is hard to track. However, we can still look at overall cooperate and defect rates for extended interactions. We restrict the data to pairs of players that have played at least 20 games between them. [1]

Table 5.9: Cooperation rates between pairs of players with at least 20 games played against each other

Player 1 ID	Player 2 ID	P1 Coop Rate	P2 Coop Rate	Games Played
231	68	1.000	0.963	27
2	22	0.731	0.654	26
231	2	0.896	0.483	29
228	223	0.955	0.818	22
2	37	1.000	1.000	24
2	11	1.000	0.913	23

Table 5.9 is very telling - especially in the IPD context. We see that there is a majority that perfer to Cooperate when they play against the same opponent for a large number of games. The pair [231, 2] is the only exception. The reason for this discrepancy lies in the fact that 231 was an *AllC* bot. Thus, as soon as player 2 caught on, the best strategy would be *AllD* which is what caused the dip in cooperation rate.

---

[1]There was one pair that had 93 games played between them. They had colluded such that for one round, one player would always *Cooperate* while the other always Betrayed, and then the next round they would switch roles. They are not counted in this data

5.4 Winners

We utilized the Facebook network for many reasons - one of them being that friends are competitive with friends. Due to this competitive edge, people playing our game would most likely want to beat their friends (hence the existance of a leaderboard). In this section, we will briefly look at the top players and how they played. Again, similar to the analysis of pair play in 5.3, we are limited to looking at sketches of overall play.

Table 5.10: Statistics of Top Players

Player ID	Prison Time	Coop Rate	# of Unique Opps.	2nd Longest Pair Streak
108	24	0.000	19	3
166	24	0.240	13	2
74	25	0.351	24	5
113	26	0.000	19	3
110	26	0.459	43	15
224	27	0.372	58	2
41	28	0.000	34	2

When we see Table 5.10, we note a few ingredients to winning shared amongst the top players.

5.4.1 Betray, betray, betray

There is not a signle player in the top 7 that has a cooperation rate higher than 50% whereas the overall cooperation rate for players was high up in the 70s. We see that taking advantage of others cooperating is key in winning and getting ahead.

5.4.2 Play one, move on

We tracked the number of unique opponents as well as the second longest streak of game playes amongst a pair of players. We note that the number of unique opponents is quite high (compared to the time in prison) for all of these top players. [2]

---

[2]Many of these players finished the game (got out of prison) multiple times, thus their unique opponent numbers might be a bit high.

Furthermore, we tracked the second longest streak of games between pairs of players. For instance, for the top player, 108, we looked at all of the games that he played, and took note of the second largest number of games played between him and the same opponent. We looked at second largest due to the fact that the actual largest number was due to playing the *AllC* bot - and exploiting it repeatedly.

From these statistics, we see that the top players maintain a strategy of playing different opponents, and never playing too many games against the same opponent.

5.4.3 How to Win

We can now start to form a strategy that these top players utilize:

- Play different players

- Betray them

- Never play the same player too many times - move on to new opponents

This strategy works nicely in our environment as we saw the high overall cooperation rates, as well as start effects of cooperation.

Chapter 6

Conclusion

Here, we will briefly talk about the design and implementation of our Facebook game. Then, we will summarize our results and analysis, and comment on some of the shortcomings of these results - and suggestions for future improvements.

6.1 Design and Implementation

We created a Facebook game to run the Iterated Prisoner's Dilemma in a social setting. The Facebook network provided us with a number of independent variables (gender, age, etc.) to allow us to study the social effects of playing the IPD. Furthermore, being based on Facebook introduced a competitve aspect between players that hopefully promoted "smarter" play by users.

We implemented the game in Ruby on Rails, hosted on Heroku. The analysis of the data was done via Python.

6.2 Results and Analysis

Our first version (rev. 1) garnered over 800 user sign ups and 3,000 completed games. Our analysis showed an overwhelming percentage (over 70 percent) of cooperation from players. Furthermore, we saw that gender and payoff (stage) played significant roles when users played. Other factors such as age and whether or not two users were

friends did not influence players' decisions.

When we analyzed the top players of our game, we noticed a pattern of game play:

- Play different players (don't play the same player too much)

- Betray

This was extremely effective against a predominantly cooperative user base. Furthermore, by varying opponents, these top players could betray without having to deal with the repercussions that may have resulted due to access to histories and reputation effects.

6.3   Shortcomings and Improvements

We conducted our experiment on the Facebook network, and thus were not able to use a "controlled setting." Therefore, we were not able to study the chronology of decision making of players (and how it evolved).

Despite this shortcoming, we were still able to analyze long-term pair relationships as well as winning strategies. Furthermore, the social setting afforded us new variables that would not be available in a controlled setting - yielding great results and analysis from a social perspective as well.

Appendix A

Tables

	Cooperate	Betray
Cooperate	(3, 3)	(-2, 5)
Betray	(5, -2)	(0, 0)

Table A.1: Payoff diagram for the Kitchen.

	Cooperate	Betray
Cooperate	(6, 6)	(-2, 8)
Betray	(8, -2)	(0, 0)

Table A.2: Payoff diagramfor the Cell.

	Cooperate	Betray
Cooperate	(6, 6)	(0, 12)
Betray	(12, 0)	(-12, -12)

Table A.3: Payoff diagram for the Yard.

	Cooperate	Betray
Cooperate	(12, 12)	(−8, 20)
Betray	(20, -8)	(0, 0)

Table A.4: Payoff diagram for the Hospital.

	Cooperate	Betray
Cooperate	(12, 12)	(-4, 16)
Betray	(16, -4)	(0, 0)

Table A.5: Payoff diagram for the Office.

Appendix B

Figures

Figure B-1: The landing page for our game
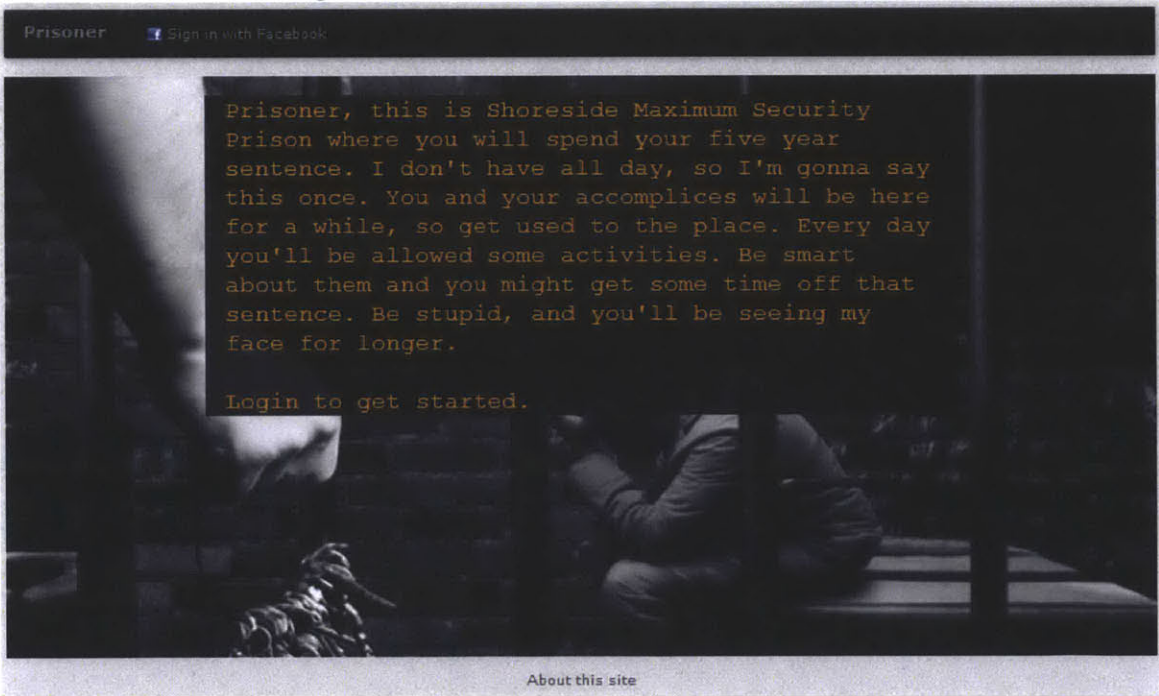
Figure B-2: The stage selection page for users
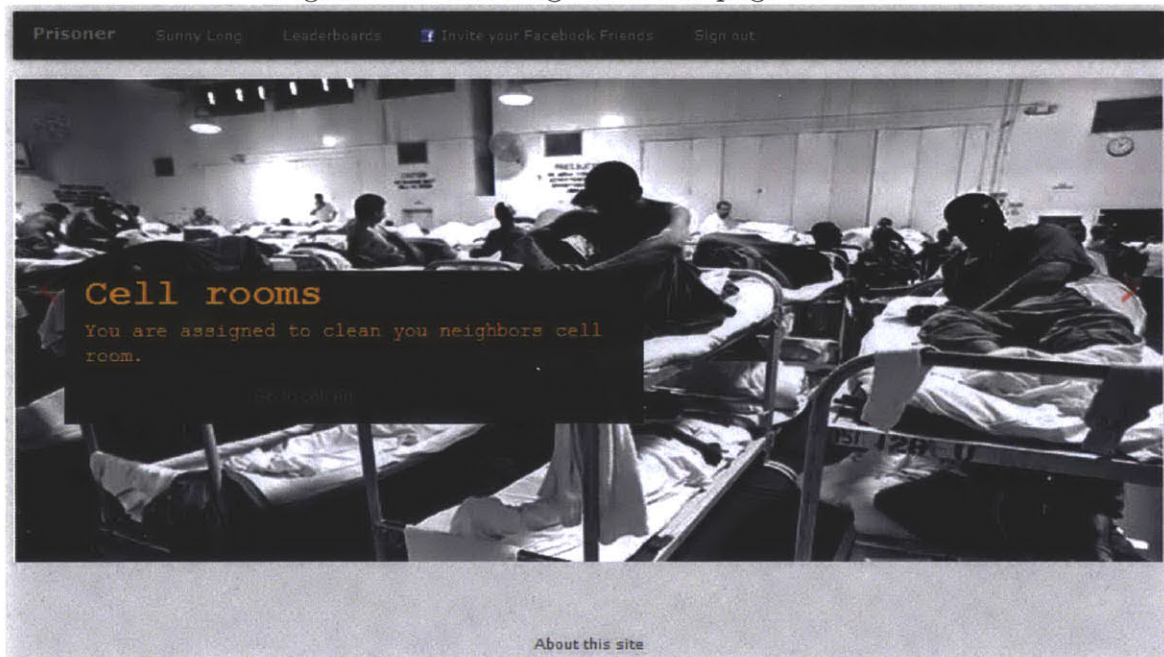
Figure B-3: Stage 1 - The Dining Hall

Figure B-4: Stage 2 - The Kitchen
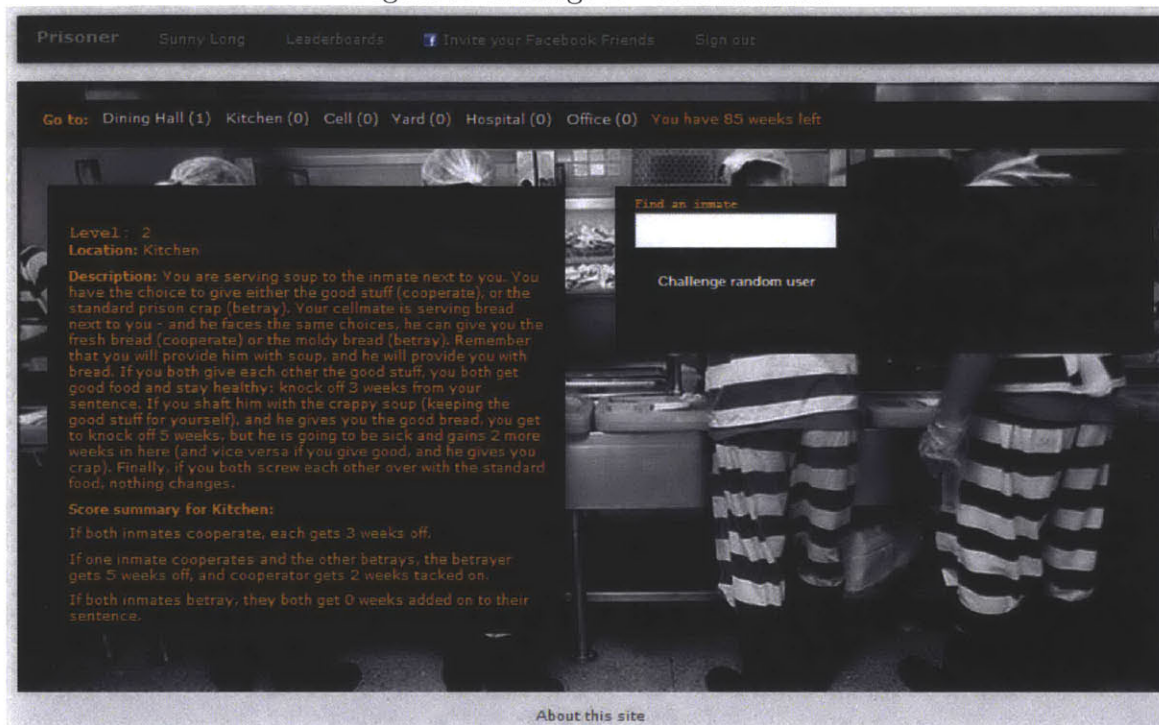
Figure B-5: Stage 3 - The Cell
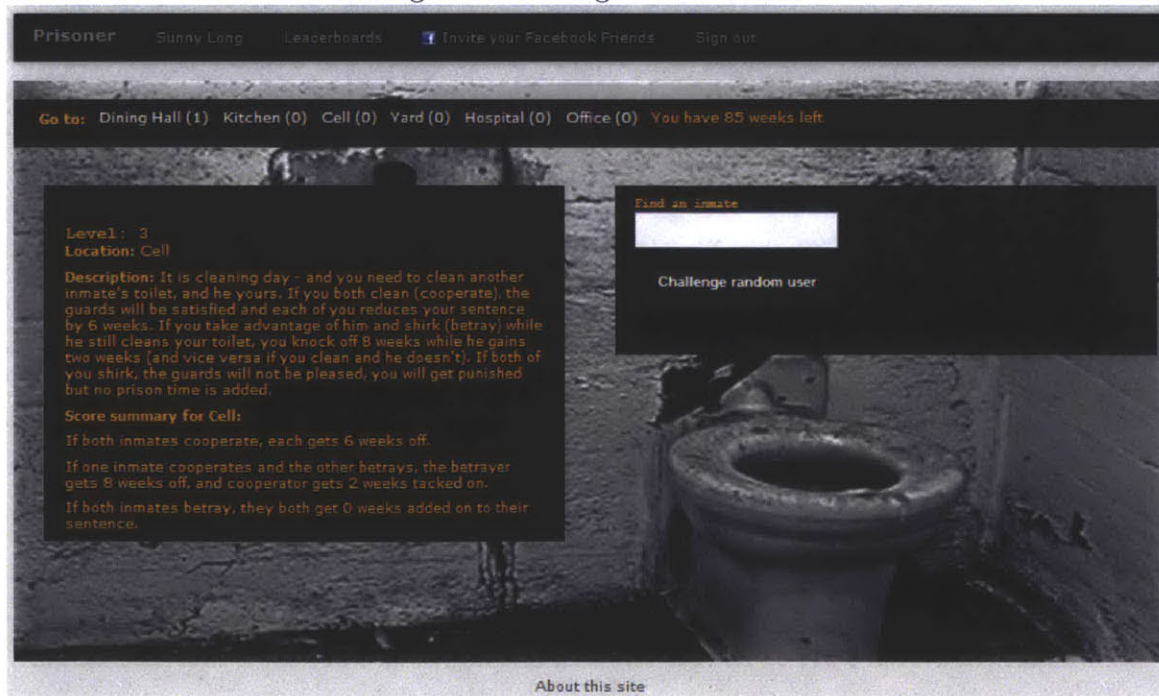
Figure B-6: Stage 4 - The Yard
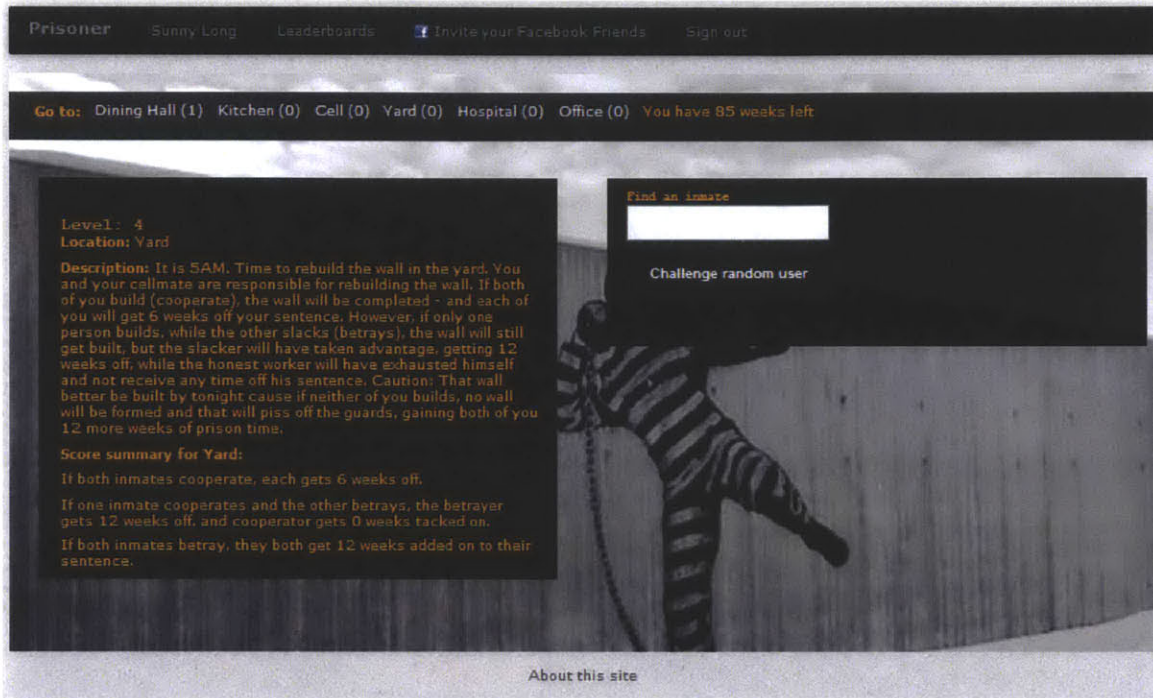
Figure B-7: Stage 5 - The Hospital

Figure B-8: Stage 6 - The Office

Figure B-9: User's statistics on games played
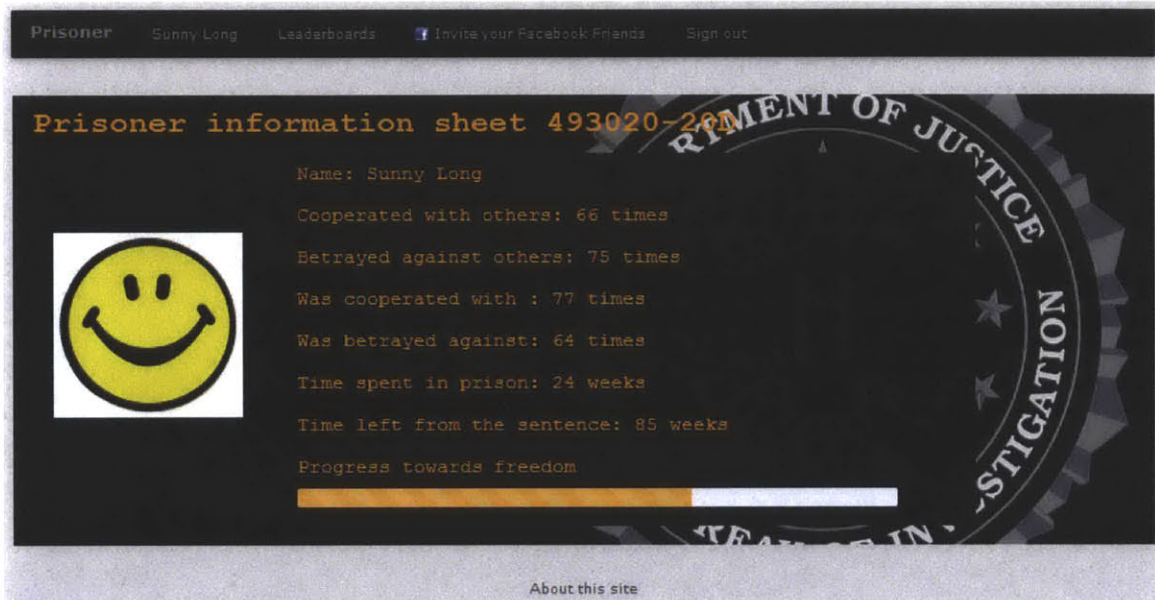
Figure B-10: Leaderboard of all users

Figure B-11: Recommended challenges

Bibliography

[1] Paul R. Abramson. Using intrade.com to teach campaign strategies in the 2008 u.s. presidential election. *PS: Political Science and Politics*, 43(1):pp. 139–144, 2010.

[2] Robert Axelrod. More effective choice in the prisoner's dilemma. *The Journal of Conflict Resolution*, 24(3):pp. 379–403, 1980.

[3] Robert Axelrod. The emergence of cooperation among egoists. *The American Political Science Review*, 75(2):pp. 306–318, 1981.

[4] Robert Axelrod and Robert O. Keohane. Achieving cooperation under anarchy: Strategies and institutions. *World Politics*, 38(1):pp. 226–254, 1985.

[5] Robert M. Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.

[6] B. Beaufils, J. P. Delahaye, and P. Mathieu. Our meeting with gradual, a good strategy for the iterated prisoner's dilemma. In *Artificial Life V: Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems*, 1996.

[7] P. J. Darwen and X. Yao. On evolving robust strategies for iterated prisoner's dilemma. In *IN PROGRESS IN EVOLUTIONARY COMPUTATION*, pages 276–292. Springer-Verlag, 1995.

[8] Jr. Dolbear, F. Trenery and Lester B. Lave. Risk orientation as a predictor in the prisoner's dilemma. *The Journal of Conflict Resolution*, 10(4):pp. 506–515, 1966.

[9] Ray C. Fair. Interpreting the predictive uncertainty of elections. *The Journal of Politics*, 71(2):pp. 612–626, 2009.

[10] Morris F. Friedell. A laboratory experiment in retaliation. *The Journal of Conflict Resolution*, 12(3):pp. 357–373, 1968.

[11] Drew Fudenberg, David G. Rand, and Anna Dreber. Slow to anger and fast to forgive: Cooperation in an uncertain world. *American Economic Review*, 102(2):720–49, September 2012.

[12] James P. Gahagan and James T. Tedeschi. Strategy and the credibility of promises in the prisoner's dilemma game. *The Journal of Conflict Resolution*, 12(2):pp. 224–234, 1968.

[13] Ch. Hauert and H. G. Schuster. Effects of increasing the number of players and memory size in the iterated prisoner's dilemma: A numerical approach. *Proceedings: Biological Sciences*, 264(1381):pp. 513–519, 1997.

[14] Ian Hay. *The First Hundred Thousand*. Blackwood, 1916.

[15] Marek M. Kaminski. *Games Prisoner's Play: The Tragicomic Worlds of Polish Prison*. Princeton University Press, 2010.

[16] Lester B. Lave. An empirical approach to the prisoners' dilemma game. *The Quarterly Journal of Economics*, 76(3):pp. 424–436, 1962.

[17] Robert Duncan Luce and Howard Raïffa. *Games and Decisions: Introduction and Critical Survey*. Courier Dover Publications, 1957.

[18] Bonaventura Majolo, Kaye Ames, Rachel Brumpton, Rebecca Garratt, Kate Hall, and Natasha Wilson. Human friendship favours cooperation in the iterated prisoner's dilemma. *Behaviour*, 143(11):pp. 1383–1395, 2006.

[19] J. Keith Murnighan and Alvin E. Roth. Expecting continued play in prisoner's dilemma games: A test of several models. *The Journal of Conflict Resolution*, 27(2):pp. 279–300, 1983.

[20] John F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Science*, 36:48–49, January 1950.

[21] Martin A. Nowak and K. Sigmund. A strategy of win-stay, lose-shift that outperforms tit-for-tat in prisoner's dilemma. *Nature*, 364:56–58, 1993.

[22] Andrzej Pelc and Krzysztof J. Pelc. Same game, new tricks: What makes a good strategy in the prisoner's dilemma? *The Journal of Conflict Resolution*, 53(5):pp. 774–793, 2009.

[23] Amnon Rapoport. Optimal policies for the prisoners dilemma. In *Experimental Studies of Interactive Decisions*, volume 5 of *Theory and Decision Library*, pages 1–18. Springer Netherlands, 1990.

[24] Anatol Rapoport. *Strategy and Conscience*. Harper & Row, 1964.

[25] Anatol Rapoport and Phillip S. Dale. The "end" and "start" effects in iterated prisoner's dilemma. *The Journal of Conflict Resolution*, 10(3):pp. 363–366, 1966.

[26] David Rothschild. Forecasting elections: Comparing prediction markets, polls, and their biases. *The Public Opinion Quarterly*, 73(5):pp. 895–916, 2009.

[27] Philip Streich and Jack S. Levy. Time horizons, discounting, and intertemporal choice. *The Journal of Conflict Resolution*, 51(2):pp. 199–226, 2007.

[28] Martijn J. van den Assem, Dennie van Dolder, and Richard H. Thaler. Split or steal? cooperative behavior when the stakes are large. *Management Science*, 58:2–20, January 2012.

[29] Chen-Bo Zhong, Jeffrey Loewenstein, and J. Keith Murnighan. Speaking the same language: The cooperative effects of labeling in the prisoner's dilemma. *The Journal of Conflict Resolution*, 51(3):pp. 431–456, 2007.