Reputation in Indefinite Interactions: Experimental and Empirical Evidence

Lorenz Adams

Department of Economics

University of Exeter

Submitted for the degree of

Doctor of Philosophy in Economics

October 2018

I dedicate this thesis to my loving mother and father with thanks for their continuous and unconditional support on this journey...

Declaration

Thesis Title: Reputation in Indefinite Interactions: Experimental and Empirical Evidence

Submitted by Lorenz Adams to the University of Exeter as a thesis for the degree of Doctor of Philosophy in Economics, October 2018.

I certify that all material in this thesis which is not my own work has been identified and that no material has previously been submitted and approved for the award of a degree by this or any other university.

Chapter three of this thesis is partially based on work done in collaboration with Dr Xiaohui Zhang from the University of Exeter.

This dissertation contains fewer than 65,000 words including appendices, bibliography, footnotes, tables and equations and has fewer than 150 figures.

(Signature) ....

Lorenz Adams
October 2018

Acknowledgements

I thank my supervisors, Dr Rajiv Sarin, Dr Miguel Fonseca and Dr Xiaohui Zhang, for their never ending support and encouragement over the course of my whole PhD. Without them I could not have completed this project and I owe them a substantial debt of gratitude.

I would like to acknowledge my family for always standing by me, especially when things got difficult. Their guidance has made me the person I am today and I could not have persevered throughout this project without the values they have instilled in me.

My appreciation also goes to the PGR support staff at Exeter University's Business School for providing me with the assistance and funding necessary to undertake my research.

I thank the entire Economics Department at the University of Exeter for introducing me to the world of academia. I have relished the opportunity to interact with them and I am grateful for the time taken to give me feedback on my ideas.

A very special thanks goes to the following faculties: Dr David Kelsey, Dr Dieter Balkenborg, Dr Brit Grosskopf, Dr Jack Rogers, Dr Carlos Cortinhas, Dr Eva Poen, Dr Kaustav Das, Dr Rish Singhania, Dr Namhyun Kim, Dr Sebastian Kripfganz, Dr Surajeet Chakravarty and Dr Cecilia Chen.

Finally, I would like to thank my friends from Hope Hall: Rami Chehab, Aya Ghalayini, Haytem Troug, Lorenz Hartmann, Lutfi Rahimi, Hannes Titeca, Ben Balmford, Ana Cinta Gonzalez, Ernil Sabaj, Cliff Li, Bing Chao, Ozkan Haykir, Zezeng Li Leo, Elham Nik, Ali Akaak, Wan Mohd Adzim and Juliana Langgat.

Abstract

This thesis models the importance of reputation in indefinite interactions using experimental and empirical evidence. In modern society reputation plays an important role in many everyday life situations. I analyse the resulting incentive structures through microeconomic modelling and econometric regression analysis. Thus, I gauge the impact of reputational considerations on individual decision making.

Chapter one studies reputation in a random matching economy of indefinite duration. All subjects are long-lived agents, who are paired with a succession of anonymous strangers to play prisoner's dilemma games. I analyse the significance of reputation effects by varying whether or not behavioural histories are observable in the continuation game. In the absence of such information, individual interactions between the players are completely unrelated. However, when personal histories are observable, subjects can review all previous decisions of their current playing partner. The parametrization is such that cooperative behaviour can be maintained even in the absence of behavioural histories. Cooperation is enforced by a decentralized punishment mechanism, such that a single deviation will obliterate the continuation pay-offs of all players. This constructive equilibrium is highly fragile as the punishment regime cannot be reversed once it has been triggered. Personal histories highlight the interdependence of separate one shot pairings and should in theory produce a socially optimal outcome. To further illustrate the impact of reputational concerns, the experiment includes additional treatments where the stage game features an additional option not to play. The choice to abstain yields a guaranteed pay-off, which is below the one shot Nash equilibrium. Thus, subjects can preserve their reputation by accepting a reduced pay-off in the current period. The results show that agents are more likely to cooperate when their behavioural histories are observable. Yet, the option not to play has no significant effect on cooperation levels in the economy.

Chapter two focuses on a widely overlooked aspect of human interaction, which is the ability to terminate or abstain from relationships. In practice people often dedicate substantial effort to avoiding individuals they deem untrustworthy. The experiment features a modified version

of the conventional prisoner's dilemma, which makes interactions between the subjects entirely voluntary. The players are anonymous agents in a random matching economy of unknown duration. In each of their one shot matches, participants have the option to either play the prisoner's dilemma or select an outside option with a symmetric pay-off. The attractiveness of the option to abstain changes across the individual treatment configurations. Abstaining from a stage game therefore either yields a pay-off smaller than or equal to the Nash equilibrium. To fully uncover the effect of having the option not to play, I manipulate whether or not the subjects have reputational incentives. In particular, people either only observe the outcomes of their own pairings, or they can review the decision history of their current playing partner. The results suggest that the pay-off associated with the outside option is strongly linked with the individual decision to abstain. Whereas the costly outside option is never played, subjects utilize a costless outside option to avoid defectors and protect their own reputation. People only abstain from interactions when behavioural records are observable, which is in line with the theoretical predictions.

The third chapter examines the role of reputation in the Australian healthcare sector. The analysis focuses on birth episodes in the Melbourne metropolitan area. I study the decision behaviour of pregnant women with private health insurance in selecting one of 18 local hospitals. Australia's healthcare system enables holders of private health insurance to freely select their preferred healthcare provider from a wide range of options. By contrast, patients relying on Australia's universal health service are more restricted in their choice. Private patients account for approximately 40% of the Australian population and indicators of hospital performance are readily available from the Department of Health and Welfare. As a result, reputation effects should feature prominently in individual hospital selection. The analysis consists of two separate strands. Firstly, I study the distance women travel between their home and the chosen hospital. I therefore measure hospital reputation through the distance patients are prepared to travel to a particular institution. In the second part of the paper, I estimate the probability of individual hospitals being selected from a set of accessible options. I restrict the distance patients can travel and examine their choice behaviour. The results suggest that there are significant clinic specific differences in the willingness to travel, which implies that patients travel further to attend more reputable institutions. This result is complemented by substantial variation in the likelihood of individual hospitals being chosen. These differences are predominantly linked to hospital criteria and common performance indicators.

Contents

List of Figures

List of Tables

Chapter 1

Reputation in Indefinitely Repeated Games: Experimental Evidence

1.1 Introduction

In economics there is a growing interest in studying irregular interactions between anonymous agents. This interest can be explained by the structure of modern society, which is characterized by considerable globalization. In recent times, the probability of repeated contact between individuals has decreased significantly. Instead, we have an environment populated by agents, who are matched at random. When people interact anonymously as strangers, they effectively play a coordination game. They can either negotiate a mutually beneficial outcome or be forever trapped in a negative equilibrium. In the case of repeated interaction between the same individuals, achieving a positive outcome is generally straightforward. However, when subjects are paired randomly, sustaining cooperation among a group of strangers is far more complex ([11]).

The literature on supergames starts with [17]. That paper describes the repeated play of stage games in an infinite sequence. By contrast, [47] simulates an experimental environment with a random rule of continuation. Thus, each period of the interaction has a fixed probability of being the final round of play. [29] uses that framework to examine the case of agents being subjected to random matching in the prisoner's dilemma. In that set-up subjects only observe the outcomes of the one shot games they play. In addition, they have no means of communicating with or identifying each other. It is shown that even in the presence of such extreme informational barriers, cooperation can be sustained. The theoretical equilibrium entails that a single deviation by any player will trigger a total collapse of cooperation in the economy. The prospect of future pay-offs being obliterated, gives all players an incentive to cooperate with each other. [13] extends the model to adjust the severity of the community punishment. The addition is significant, as it makes the constructive equilibrium more robust. Hence, an individual deviation will not destroy cooperation.

[6] studies the dynamics of contagious punishment equilibria. The study identifies the theoretical minimum conditions required to maintain cooperation among a group of anonymous agents. The stability of the constructive equilibrium depends mainly on the opportunity cost of cooperating in each stage game, as well as the contact rate. This is the rate at which an individual defector expects to meet cooperators in the continuation game. These parameters are used to derive exact continuation pay-offs, both on and off the equilibrium path. The analysis allows for a precise characterization of the individual incentives, which are created by contagious punishment schemes.

In theory, even large groups of anonymous subjects can sustain cooperation. This result is conditional on the economy being indefinite in duration and players exhibiting sufficiently high discount rates. However, such an equilibrium is realistically unattractive, as it is fragile and requires the mass punishment of cooperators and defectors, alike. Reputation can help resolve this problem. Firstly, it conveys information regarding a subject's behaviour in the past. Secondly, since a negative reputation results in a worsening of future prospects, it gives individuals a strong incentive to behave cooperatively ([37]).

The study of reputational incentives famously begins with [31]. The work is based on the notion that in repeated interactions, agents may utilize the early stages to attain a reputation for being committed to a particular strategy. The paper revisits the chain-store paradox, which has a monopolist face a succession of potential entrants. The original result states that in any period the monopolist should accommodate entry into the market, since fighting a competitor is never optimal in the final round of play ([50]). However, [31] finds that a small amount of noise is enough to give rise to reputational concerns. Hence, the monopolist may choose to fight entrants early on to acquire the reputation of being a predator. Thus, entry into the industry can be deterred in subsequent rounds.

[29] develops a theoretical model, which allows for information about the behaviour of long lived players to be transmitted within a community of anonymous agents. As such, each subject is assigned a personal label, which accurately summarizes all decisions she has taken in the past. The label is updated at the start of each round, based on a subject's decision

in the preceding period. The revision of the label occurs without the knowledge of the entire community. As a result, each player only has local information about the state of the economy. The introduction of personal labels gives rise to a dynamically stable equilibrium, by which a socially optimal outcome can be achieved.

The prisoner's dilemma formalizes the conflict between individual gain and the common good, which can be attained through joint cooperation. This struggle is an integral aspect of human nature and the prisoner's dilemma is therefore useful in modelling interactions among strangers ([33]). Using the game as a platform for laboratory experiments offers important advantages. Firstly, due to the intuitive structure it is easy for subjects to grasp the nature of the interaction. Secondly, the prisoner's dilemma offers a salient link between the available action choices and the respective pay-offs.

[4] simulates indefinitely repeated prisoner's dilemma games with a stochastic endpoint. The experiment therefore creates an environment where there is always a possibility of continued interaction between the players. The experiment considers two scenarios with positive continuation probabilities. In addition, the study also includes treatments where the game is played a fixed number of times. The main purpose is to determine if the shadow of the future can remedy opportunism and deliver efficiency gains to the market. The most prominent result is that subjects perceive differences in the incentive structure of finitely repeated games, compared to those that continue indefinitely. Whereas players quickly learn to defect in the finitely repeated game, they cooperate when the time horizon is infinite.

[5] provides experimental evidence of random matching economies where anonymous agents play indefinitely repeated prisoner's dilemma games. For the given parametrization, complete cooperation by all subjects is an equilibrium of the continuation game. The paper reports evidence from four different treatments, whereby the duration of the economy is determined through a probabilistic continuation rule. The treatment configurations vary along two dimensions. Under private monitoring subjects only observe the outcomes of the matches they play, whereas under public monitoring they observe the state of the entire economy. Some treatments only allow subjects to punish opportunism by defecting in the stage game, while in others they can inflict a loss on defectors. For each experimental session, a group of twenty subjects is split up into five separate economies. People then only interact with the other members of their own group. The results show that generally high cooperation levels are achieved within each economy. The use of personal punishment is effective in enforcing cooperation. Moreover, cooperation levels are maximized when subjects observe

the decision histories of their playing partners.

This experimental study seeks to extend the analysis in [5]. I simulate an economy consisting of ten anonymous subjects, who repeatedly play prisoner's dilemma games. The duration of the economy is indefinite and the players are randomly and anonymously rematched after each round. In the absence of personal histories, subjects merely observe the results of their own stage games. However, when personal histories are available they can review all past decisions of their current playing partner. In some treatments subjects can react to a negative reputation only by defecting in the stage game. In other treatments an additional action choice is added to the prisoner's dilemma, which allows people not to play.

This chapter makes three significant contributions to the existent literature. Firstly, information about the state of the economy is transmitted locally, rather than at an aggregate level. Therefore, subjects only receive information about the behavioural history of the player they are currently paired with. This is a realistic approach to studying reputation, since people rarely have complete information about all other agents in their community. Secondly, I significantly increase the size of the experimental economy. Whereas [5] creates communities consisting of only four players, my study features economies of ten subjects. Thus, I am able to empirically examine whether very small communities are able to sustain cooperation more efficiently than others. Finally, I complement the theoretical framework of [29] and [53] to give subjects an option to abstain from one shot interactions in response to a negative reputation.

I attempt to answer the following questions: How robust is the result that communities can achieve cooperation even in the presence of severe informational constraints? Does an increase in the number of subjects make it more challenging to maintain a socially optimal equilibrium? Is the observability of personal histories effective in reducing the number of deviations in the community? Do subjects embrace the option to abstain from interactions with other players? Do they utilize the outside option to penalize deviations and preserve their own reputation?

The results of the experiment indicate that under private monitoring the level of cooperation in the economy is negligibly low. It is apparent that the introduction of personal histories substantially increases cooperation among the subjects. However, even when personal histories are observable, cooperative behaviour does not rise to a socially efficient level ([8]). In addition, players do not make use of the outside option, which enables them to punish opportunism whilst protecting their own label. The findings of this study therefore suggest that reputational considerations play a far less critical role in individual decision making than is commonly believed.

The remainder of this chapter is structured as follows: Section two provides a detailed overview of the theoretical framework. Section three describes the experimental design and stipulates the hypotheses. Section four analyses the data and discusses the results of the experiment. Section five summarizes the key findings and concludes the chapter.

1.2 The Model

I use generic notation to denote a subject's action choices and pay-offs. Throughout my analysis I assume that all players seek to maximize their earnings.

Table 1.1 Notation in the Theoretical Analysis

	Cooperate (C)	Defect (D)	Opt Out (O)
Cooperate (C)	π_{CC}, π_{CC}	π_{CD}, π_{DC}	π_{OO}, π_{OO}
Defect (D)	π_{DC}, π_{CD}	π_{DD}, π_{DD}	π_{OO}, π_{OO}
Opt Out (O)	π_{OO}, π_{OO}	π_{OO}, π_{OO}	π_{OO}, π_{OO}

The possible stage game pay-offs are $\pi_{CC}, \pi_{CD}, \pi_{DC}, \pi_{DD}$ and π_{OO} when the outside option is available. The first subscript refers to the action choice of the row player and the latter subscript to that of the column player. In the prisoner's dilemma joint cooperation produces the socially optimal outcome. However, as $\pi_{DC} > \pi_{CC}$ and $\pi_{DD} > \pi_{CD}$, the game constitutes a social dilemma where each player has an incentive to deviate. The guaranteed pay-off associated with the outside option is such that $\pi_{DD} > \pi_{OO} > \pi_{CD}$.

1.2.1 Equilibrium in the Stage Game (No Outside Option)

In the one shot interaction, an individual should always choose to defect in order to maximize her expected pay-off. This is easy to see, since cooperation constitutes a dominated strategy. By opting to play D, a player is guaranteed to attain π_{DD}. This pay-off is greater than π_{CD} and defection therefore represents an optimal response to a play of D by the opponent. Should the playing partner cooperate, defection will yield π_{DC}, which exceeds the individual pay-off from joint cooperation. It follows therefore that in the stage game both players should choose to defect and earn π_{DD}, the mutual minimax pay-off. This outcome constitutes a unique Nash equilibrium.

Proposition 1. *Mutual defection is a unique Nash equilibrium in the one shot prisoner's dilemma.*

1.2.2 Equilibria in the Repeated Game (No Reputation/No Outside Option)

To start with, I develop a framework for the private monitoring treatment without the outside option. This treatment configuration forms a theoretical baseline for all further analysis. All individual meetings between the subjects are independent, as subjects observe only the action profiles and outcomes of the matches they play. Since mutual defection is a stable outcome within each of the stage games, it follows that defection remains an equilibrium strategy in the repeated game, as well.

Proposition 2. *A scenario where subjects defect in all their matches constitutes a sequential equilibrium of the continuation game.*

However, unlike the stage game, the repeated interaction exhibits a multiplicity of equilibria. In particular, the socially optimal outcome where all subjects cooperate in every period is an equilibrium of the indefinitely repeated game. This cooperative outcome is sustained by a contagious punishment mechanism, which is first introduced in [29]. In the following analysis I use the results in [6] to define a constructive equilibrium. [1]

I consider a scenario where each subject abides by a social norm, which consists of two components: a rule of conduct (always play C), as well as a punishment mechanism: (always play D). Subjects revert to the punishment mechanism as soon as a deviation from the rule of conduct is observed, which is similar to a grim trigger:

If subject i adheres to the rule of conduct (always cooperate) in period t, then she will switch to the punishment mechanism (always defect) in period $t+1$, if and only if a defection was witnessed. Otherwise, subject i will continue to follow the rule of conduct.

If all subjects pursue the above strategy, a single defection will trigger a collapse of cooperation in the economy. The initial deviation will spread through the community much like a contagious disease and erode the earnings of all agents (Figure 1.1). It is that erosion of future pay-offs, which provides the individual incentive to sustain cooperative behaviour. Therefore, if agents display sufficiently high discount rates, none of them will have any incentive to defect. This can be formally expressed as follows:

$$\frac{\pi_{CC}}{1-\beta} \geq \frac{1}{1-\beta}[\pi_{DC}\phi_k + \pi_{DD}(1-\phi_k)] \tag{1.1}$$

Subject i's discount rate is given by $\beta \; \varepsilon \; (0,1)$, the continuation probability of the repeated game. The parameter ϕ_k is the contact rate in the economy, which I define below. The left hand side of equation (1.1) represents the continuation pay-off to subject i, if she cooperates. The right side of that equation gives her expected pay-off if she chooses to defect. In order for a constructive equilibrium to exist, the pay-off from cooperation must be greater than the expected gain from defection.

[1]The proof is given in Section 2 of that paper.

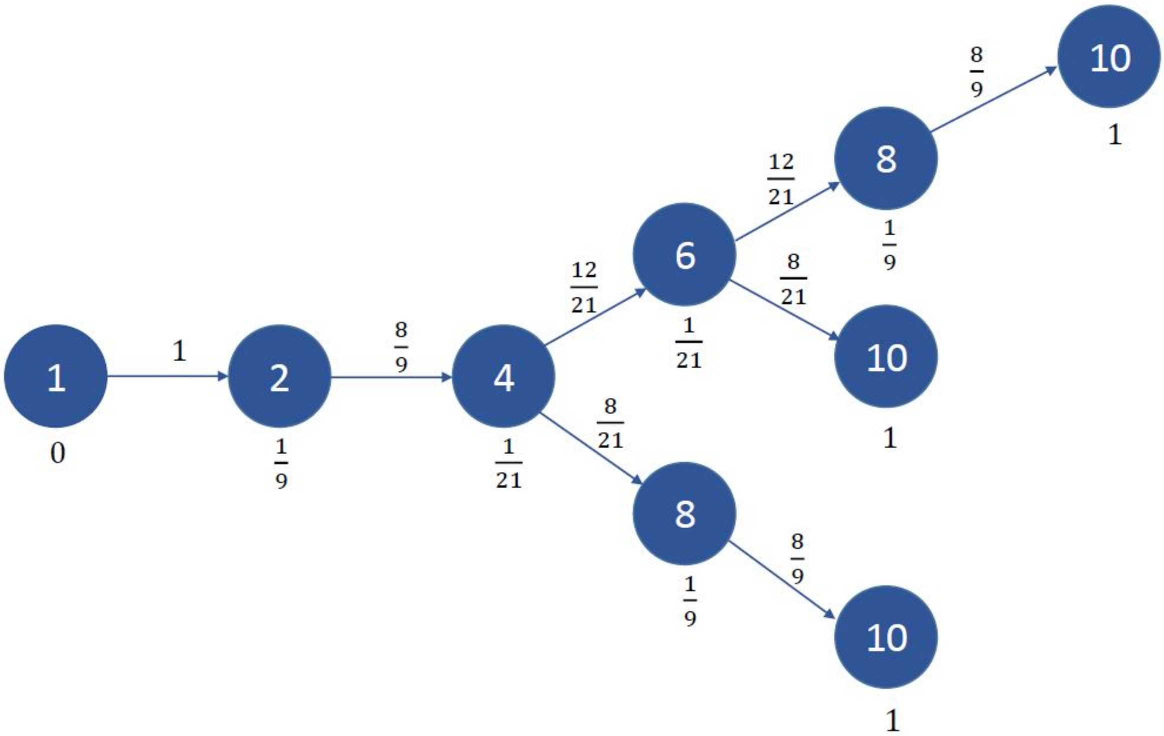

Figure 1.1 The Spread of Deviations in the Random Matching Community

In equation (1.1), ϕ_k represents the contact rate in the economy. This is the rate at which a given defector, subject i, anticipates to meet cooperators in the continuation game. To derive the contact rate, one must first calculate the number of cooperators that the defector expects to be paired with:

$$\sigma_k + \beta \sum_{k'=k}^{N} Q_{kk'}\sigma_{k'} + \beta^2 \sum_{k'=k}^{N} Q^2_{kk'}\sigma_{k'} + .. = e_k^T (I - \beta Q)^{-1}\sigma \tag{1.2}$$

The economy consists of N subjects. In a given period t, there are k defectors in the community, as well as $N-k$ co-operators. Therefore, σ_k is the probability that subject i is matched with a cooperator in period t, given the current state of the economy. Since subjects are randomly and bilaterally rematched after each stage game, $Q_{kk'}$ is the probability that in the next round the number of defectors increases from k to k'. More specifically, $Q_{kk'}$ is the probability that in the next period $k'-k$ mixed pairings are realized. A mixed pairing occurs whenever a defector is matched with a cooperator under the random matching protocol. It is now straightforward to interpret the meaning of $Q_{kk'}\sigma_{k'}$. This term simply yields the probability that subject i will be matched with a cooperator in any future round of the interaction. That probability depends on a pairing being realized, which increases the number of defectors from k to k'. Finally, as the duration of the economy is not infinite, it is necessary to discount future rounds of play in accordance with the continuation probability β. The vector e_k^T represents the current state of the economy, whereas Q is a triangular diffusion matrix with non zero entries above the main diagonal. Also, σ is a vector containing the conditional probabilities of being involved in mixed pairings ([6]).

The economic intuition captured in equation (1.2) is that even though the horizon of the continuation game is infinite, each cooperator in the economy can only be met once. This is a key characteristic of any contagious punishment process. Whenever a defector meets a cooperator, the cooperator subsequently reverts to the punishment protocol. It follows that once an initial defection has occurred, the economy will, in finite time, transition to a state where all subjects follow the punishment protocol. The number of cooperators that any individual defector can meet, is bounded above by $(1-\beta)^{-1}$, the expected duration of the economy. This is because the expected number of meetings with cooperators can logically never be greater than the projected duration of the community. Therefore, in order to obtain ϕ_k, I divide the expected number of cooperators by the predicted number of rounds to be played. Thus, I obtain the average probability of an individual defector being paired with a cooperator in any given period:

$$\phi_k = (1-\beta)e_k^T(I-\beta Q)^{-1}\sigma \tag{1.3}$$

The term ϕ_k in equation (1.3) captures the rate at which subject i meets cooperators throughout the indefinitely repeated interaction. This is calculated on the basis that there are currently $k \geq 1$ defectors in the economy. The contact rate can be interpreted as follows: if there are currently few defectors in the community, the probability of being matched with cooperators will be high. This results in ϕ_k being close to one. Similarly, if there are many defectors in the community, the probability of being matched with cooperators in the continuation game will be low, which means that ϕ_k is close to zero.

It is apparent that in order for cooperative behaviour to be maintained by a contagious punishment mechanism, it must be the case that subject i has no incentive to defect in an environment where there are currently no defectors. As such, none of the players should have any incentive to commit a deviation from the rule of conduct. Formally, this condition can be expressed as follows:

$$\frac{\pi_{CC}}{1-\beta} \geq \frac{1}{1-\beta}[\pi_{DC}\phi_1 + \pi_{DD}(1-\phi_1)] \tag{1.4}$$

Equation (1.4) stipulates the condition required for the existence of a constructive equilibrium. It states that subject i's pay-off from continued cooperation must be greater than her expected gain, should she choose to commit the first defection.

In order to verify the existence of a contagious punishment equilibrium, it is necessary to show that none of the subjects have an incentive to stop the collapse of cooperation, once a contagious response has been triggered. Therefore, let subject i be one of $k \geq 2$ defectors currently in the community. The following condition must be satisfied for no subject to have any reason to slow down the punishment response:

$$\sum_{k'=k}^{N} Q_{kk'}[\sigma_{k'}(\pi_{CC}+\beta v_{k'-1})+(1-\sigma_{k'})(\pi_{CD}+\beta v_{k'})] \leq \sum_{k'=k}^{N} Q_{kk'}[\sigma_{k'}\pi_{DC}+(1-\sigma_{k'})\pi_{DD}+\beta v_{k'}] \quad (1.5)$$

The right side of equation (1.5) gives the pay-off earned by subject i if she participates in the punishment regime. Given that the number of defectors in the economy increases from k to k' with probability $Q_{kk'}$, there are two possible scenarios to be considered. Firstly, subject i is involved in a mixed match with probability $\sigma_{k'}$. In this case subject i will earn π_{DC} by choosing to defect. Alternatively, subject i is paired with another defector with probability $(1-\sigma_{k'})$. If that outcome is realized, the pay-off she earns is π_{DD}. The variable $\beta v_{k'}$ gives subject i's discounted pay-offs in the continuation game, given there are now k' defectors in the economy.

The left hand side of equation (1.5) illustrates the case of subject i attempting to slow down the contagious punishment regime. In particular, subject i deviates from the punishment protocol by choosing to play C in the current period (instead of D, which is the designated punishment action). Again there are two possible cases to be considered: subject i is either involved in a match with a cooperator, or otherwise she is paired with a defector. If matched with a cooperator, subject i's play of C will earn her π_{CC} instead of π_{DC}. However, since subject i's partner is not turned into a defector at this stage, the future pay-offs in the continuation game are now greater than before. In particular, they are $\beta v_{k'-1}$ instead of $\beta v_{k'}$. On the other hand, if subject i is paired with a defector, her C choice will earn her π_{CD} instead of π_{DD}. Yet, in this particular case there is no effect on the continuation pay-offs, as no new infection would have occurred anyway.

The conditions specified above are sufficient to confirm the existence of a constructive equilibrium, which is supported by a contagious punishment scheme. The conditions stipulated by equations (1.4) and (1.5) are satisfied for the parametrization in this study.

Proposition 3. *An outcome where all subjects cooperate in every match constitutes a sequential equilibrium even in the absence of personal histories.*

1.2.3 Equilibria in the Stage Game (Outside Option)

Adding an outside option to the prisoner's dilemma results in an additional stage game Nash equilibrium. A scenario where both players opt out from the interaction is now also a stable outcome. However, for the stage game outlined in table 1.1, this Nash equilibrium is Pareto dominated. This is due to the fact that $\pi_{OO} < \pi_{DD}$. Therefore, in a one shot interaction no subject should ever choose to opt out.

Proposition 4. *When the outside option is available there exists an additional Nash equilibrium, such that both players abstain from the interaction.*

1.2.4 Equilibria in the Repeated Game (No Reputation/Outside Option)

Next, I construct a framework for the private monitoring treatment with the outside option. As before, all individual meetings between the players are completely unrelated and subjects only observe the outcomes of the matches they are directly involved in. It is obvious that the two sequential equilibria, which existed in the absence of the outside option, are also equilibria in this treatment configuration. However, given that the case of both players opting out constitutes a Nash equilibrium in the stage game (albeit a Pareto dominated one), it is straightforward to extend this to the repeated game. In particular, all players opting out in every match they play constitutes an additional sequential equilibrium. This is due to the fact that opting out is always a best response to a play of O by the opponent.

Proposition 5. *When the outside option is included, an additional sequential equilibrium exists where all subjects abstain from all their matches.*

1.2.5 Equilibria in the Repeated Game (Reputation/No Outside Option)

The analysis has so far focused on how even under limited observability a community can achieve a socially optimal outcome by imposing strict rules of conduct and severe sanctions. However, such an equilibrium is highly fragile and requires the punishment of innocent players. I now extend the analysis to the case where personal histories are observable. Unlike before, individual encounters between the subjects are not unrelated in this environment. In particular, subjects can now observe a complete decision history for the individual they are paired with in the current round of play. I first consider the case without the outside option.

[53] develops a model for exactly this framework. Each subject is assigned a label, which gives an accurate account of all her previous decisions. As a result, when two players are randomly matched, they first observe each other's history and then select their actions. An individual's label is updated every period, based only on the decision she makes. The opponent's choice therefore plays no direct role in updating her individual label. Instead, the partner's behavioural record affects the decision made in the current round. Formally, we have the following set up:

1. In period t, subject i has a label $z_i(t)$, such that $z_i(t) = (a_i(1), ..., a_i(t-1))$.
2. When subjects i and j are matched in period t and choose their actions $a_i(t)$ and $a_j(t)$, their respective labels will subsequently be updated as follows:

$$(z_i(t+1)) = f(z_i(t), a_i(t)) \tag{1.6}$$

$$(z_j(t+1)) = f(z_j(t), a_j(t)) \tag{1.7}$$

3. At the start of period t, each subject is aware of her own label and can observe that of her opponent.

I construct an equilibrium based on the stage game parametrization in my study. First, I reduce the pay-off matrix to its normal form as shown in [52]. Thus, I obtain a reduced form of the stage game, which is shown in table 1.2:

Table 1.2 The Reduced Form Game

	Cooperate	Defect
Cooperate	1,1	$-l, 1+g$
Defect	$1+g, -l$	0,0

In the above stage game, the variable g captures the gain from unilateral deviation in the stage game. Similarly, l represents the loss incurred when being defected upon. If $g > l$, the game is said to be submodular. Thus, the incentive to defect is greater, the more likely a partner is to cooperate. By contrast, if $g < l$ the game is supermodular. Hence, the incentive to cooperate is greater, the more likely a partner is to cooperate. Under my parametrization, $g = \frac{2}{3} = l$, which implies that the stage game is neither submodular nor supermodular ([52]).

A strategy profile satisfies the property of indifference if a player is indifferent between cooperation and defection under all behavioural histories. Independence of own play requires the subject to choose her actions independently of her own past play. A strategy is independent and indifferent, if it satisfies independence of own play and indifference under all behavioural histories.

[53] shows that if $g = l$ and $\beta \geq \frac{g}{1+g}$, there exists an independent and indifferent equilibrium with records of length one [2]. As a result, subjects choose to cooperate or defect solely based on the current partner's decision in the previous round. There are three possible scenarios:

[2]The proof is given in Section 4.5 of [53].

1. If two players are matched, who have defected in the previous round, they will mutually minimax each other, i.e. they will both defect and earn π_{DD}.
2. If two players are matched, who have cooperated in the previous match, they will both cooperate and therefore attain π_{CC}.
3. If a player, who cooperated in the previous period, is matched with a subject, who defected in the preceding round, the cooperator will minimax the defector by playing D. However, the deviator will cooperate.

It is apparent from the above analysis that a defection results in the deviator being punished only for one period. However, a single deviation irreversibly damages the economy even when personal histories are available. Specifically, an initial defection triggers a chain reaction, such that at least one player is subject to punishment action at all times. Without the outside option there is no mechanism by which to distinguish deviators from those who penalize deviations.

Proposition 6. *In the reputation treatment without the outside option there exists an independent and indifferent equilibrium with records of length one.*

1.2.6 Equilibria in the Repeated Game (Reputation/Outside Option)

Finally, I construct a theoretical framework for the reputation treatment with the outside option. Recall that playing O is a dominated strategy in the stage game. Therefore, one would not expect the outside option to be used in the treatment with private monitoring. However, in the presence of personal histories, the decision to abstain is game theoretically meaningful. As before, subjects are assigned labels, which give a true account of all their past actions. The main difference is that with the outside option included, the minimax strategy is to abstain, rather than defect. Hence, π_{OO} is the minimax pay-off in this framework. A cooperative equilibrium exists, as follows:

1. If two players are matched, who have defected in one or both of the two preceding rounds, they will both opt out and earn π_{OO}.
2. If two players are matched, who have cooperated in both preceding rounds, they will both cooperate and attain π_{CC}.
3. Finally, if a player, who cooperated in the previous two periods, is matched with a subject, who defected in one or both of the two previous rounds, the co-operator will minimax the defector by opting out. However, the other player will cooperate. Unlike before, both players will earn the minimax pay-off π_{OO}.

The outside option changes the designated punishment strategy. Punishing a defection is now costly, as both the punisher and the deviator incur the minimax pay-off. However, the punisher will not damage her reputation by having a defection added to her behavioural history.

I now show that the above strategy profile constitutes an equilibrium. My analysis follows directly from [2]. I begin by ensuring that it is individually optimal for subject i to repent, given that she defected in the previous period $t-1$:

$$(1-\beta^2)\pi_{OO}+\beta^2\pi_{CC}\geq\pi_{OO} \tag{1.8}$$

The right hand side of equation (1.8) captures the pay-off for subject i if she continues to defect. In that case she will earn π_{OO} in every round, as her playing partners choose to abstain from the pairing. The left side of the expression gives the expected pay-off to subject i if she repents and her defection is forgiven in period $t+2$.

Next, I verify that for a player, who cooperated in her two previous matches, the prescribed actions are in her best interest:

$$(1-\beta)\min[\pi_{OO},\pi_{CC}]+\beta\pi_{CC}\geq(1-\beta)\pi_{DC}+\beta[(1-\beta^2)\pi_{OO}+\beta^2\pi_{CC}] \tag{1.9}$$

In equation (1.9) the left side is the lower bound of an innocent subject's average pay-off if she follows the equilibrium strategy. Given that the outside option is available, subject i will only ever choose to cooperate or opt out. It is apparent that opting out is costly. However, the possibility to abstain allows the punisher to preserve her label. Hence, she will not be penalized in the future. The right side of the expression gives the maximum pay-off subject i can attain if she decides to deviate from the equilibrium strategy, given that the duration of the punishment phase is two periods.

Proposition 7. *With the outside option available, a single defection will result in the deviator being punished for two rounds. The economy will subsequently return to its original state.*

1.3 The Experiment

1.3.1 Experimental Design

I manipulate the observability of personal histories for all subjects. Also, the stage game participants play varies across different treatments, depending on whether or not the outside option is included.

Table 1.3 Prisoner's Dilemma with Outside Option

	Left (Cooperate)	Right (Defect)	Opt Out
Left (Cooperate)	30,30	5,40	10,10
Right (Defect)	40,5	15,15	10,10
Opt Out	10,10	10,10	10,10

The Stage Game: Each of the stage games consists of two players independently and simultaneously selecting an action from the available set of choices. Whenever the outside option is not included, subjects are faced with a conventional prisoner's dilemma. Otherwise, participants play the stage game shown in table 1.3, which includes an additional action choice allowing them to abstain. The numerical entries in each cell give the pay-offs for the row and column player, respectively. For the sake of neutrality the action choices Cooperate and Defect are labelled as Left and Right. When personal histories are available, subjects can observe all previous decisions of the person they are currently matched with before selecting an action choice. No other information about an individual's identity is revealed at any point.

The Indefinitely Repeated Game: Each indefinitely repeated game consists of an unknown number of stage games being played by the subjects. This infinite horizon set up is achieved by employing a random rule of continuation. As the continuation probability is constant, the expected number of additional rounds is given by $(1-\beta)^{-1}$. Since β equals 0.9 in the economy, each repeated game is expected to last for eleven rounds in total. The continuation probability and the random matching regime are identical across all treatments. The exact number of rounds in each repeated game is determined by drawing random integers between 1 and 100 from a uniform distribution ([5]). A draw of 91 or any higher number ends the interaction. The players are randomly and bilaterally matched in every period of the interaction. There are ten subjects in each economy, which results in five pairings being created every time. These pairings last for only one period, are equally likely, and independent over time.

The Experimental Session: Each of the experimental sessions involves twenty subjects, who are randomly split up into two groups of ten people. This is done by having the participants draw numbers between one and twenty from an urn. Subjects only interact with the other nine people in their own group and never with members of the other group. This is common knowledge. In every session participants play six indefinitely repeated games (Henceforth, I call them cycles). By design, every cycle terminates simultaneously for both groups. Subjects first play three cycles without personal histories. These are then followed by another three cycles where personal histories are observable. Within each session subjects therefore only encounter one of the two stage game configurations.

The Experimental Treatments: This experiment uses a 2x2 design and therefore consists of four treatments. The individual treatments differ in the observability of personal histories and the availability of the outside option (Table 1.4). For each treatment configuration five sessions are conducted in the laboratory.

Table 1.4 Comparison of Experimental Treatments

	Classic Prisoner's Dilemma	Outside Option Included
No Reputation	5 Sessions	5 Sessions
Reputation	5 Sessions	5 Sessions

In the two private monitoring treatments, subjects merely observe the outcomes of the stage games they are directly involved in. One of those treatments, where subjects play the prisoner's dilemma, is the theoretical benchmark of [29]. In the other private monitoring treatment, subjects have the additional option to abstain from any of their one shot interactions by choosing to opt out.

The two remaining treatments make personal histories observable for all players. This means that each subject can review the previous decisions made by the opponent they are matched with in the current round of play. The treatment, which implements personal histories in the classic prisoner's dilemma, constitutes an experimental test of [53] where information about subjects' behaviour is transmitted through personal labels. The other treatment, which combines reputational incentives with the option not to play, tests how subjects respond to the possibility of punishing deviations whilst maintaining a positive reputation.

To summarize, this study varies whether personal histories are observable in the economy, and the set of action choices available in the stage game. Firstly, subjects either only observe the results of the matches they play, or they can observe a complete decision history for every opponent they are paired with (No Reputation vs. Reputation). Secondly, participants either play a classic prisoner's dilemma or they are faced with a similar stage game, which includes an outside option (Classic Prisoner's Dilemma vs. Outside Option Included).

Experimental Procedures: A total of 100 people participated in this study. The subjects were undergraduate students at the University of Exeter, who were recruited from all disciplines other than Economics and Psychology. The students signed up online using the ORSEE recruitment system ([22]) and all sessions were conducted at Exeter University's FEELE laboratory. Dividers were in place between desks to ensure that subjects could not communicate with each other in any way. The participants interacted with one another anonymously through computer terminals. The experimental interface for this study was programmed using z-Tree ([16]).

Experimental instructions were laid out on all desks, and they were also read aloud [3]. After the instructions had been read by the experimenter, subjects were asked to complete a short quiz aimed at testing their understanding. All answers were checked by the experimenter and any additional questions were answered in private. The correct answers to the quiz were subsequently announced to the room and were therefore common knowledge.

Every experimental session consisted of six cycles being played. The average cycle continued for 10.24 additional periods. The variance for the distribution of randomly drawn numbers was 2.29. People's earnings were determined solely by their own actions and the actions of their playing partners. Subjects were paid their cumulative earnings for the second and third cycles with and without reputation. The experimental earnings were converted to money using a fixed exchange rate, such that 60 tokens = £1. For a session, which lasted for one hour in total, subjects earned £10.35 on average. That amount included a £3 show-up fee, which was paid to all participants. People were paid individually and advised that they were not obliged to disclose how much money they had earned.

[3] The experimental instructions can be found in Appendix A.

1.3.2 Experimental Hypotheses

To start with, I consider the significance of the outside option when subjects only observe the results of their own pairings. It is apparent that in this environment choosing not to play can never be optimal. This is due to the fact that opting out yields a lower pay-off than defection. Also, there are no reputational incentives, since personal histories cannot be observed.

Hypothesis 1. *None of the subjects should opt out in the private monitoring treatment.*

The outside option is not game theoretically meaningful under private monitoring. As a result, cooperation levels should be identical across the two treatments with private monitoring.

Hypothesis 2. *No significant difference in cooperation should be observed across the private monitoring treatments.*

Cooperation can be sustained even when personal histories are not observable (Proposition 3). However, the constructive equilibrium is highly unstable. On the other hand, when personal histories are available, each agent is assigned a label, which gives an accurate account of all her previous choices. This aspect of observability allows cooperators to identify each other and penalize deviations. As a result, a single defection will not trigger a collapse of cooperation in the economy. Instead, the community will target the defector (Proposition 6). This can only have the effect of promoting sociable behaviour.

Hypothesis 3. *Cooperation levels should be significantly higher in the reputation treatments.*

The dynamic stability of the constructive equilibrium increases with the introduction of personal histories. Yet, without the outside option the economy will be damaged irrevocably by a single defection. In particular, a deviation by any subject will result in that player being punished. However, the punishment can only be administered by defecting on the deviator in the subsequent match. As a result, the economy will transition to a state where there is always at least one player that is currently being penalized (Proposition 6).

Hypothesis 4. *Without the outside option a single defection means that the economy can never return to complete cooperation, even when personal histories are observable.*

Whereas the outside option should never be played under private monitoring, it is part of the punishment strategy in the reputation treatment (Proposition 7). Whenever a player, who cooperated in the two preceding periods, is matched with a subject, who defected in one or both of the two previous rounds, she will abstain from the match. Although this is costly, it allows the punisher to preserve her behavioural record. Therefore, no permanent damage comes to the community as a result of an individual deviation.

Hypothesis 5. *The outside option should be used to punish defections in the reputation treatment.*

Opting out in response to a defection results in a deviation not triggering a chain reaction in the economy. Instead, once the defector has been punished, the community will return to its original state (Proposition 7). Therefore, cooperation levels should be higher when the outside option is available.

Hypothesis 6. *Across the two reputation treatments, cooperation should be maximized when the outside option is available.*

In the reputation treatment without the outside option, a defection will result in the deviator being punished in the subsequent round of play (Proposition 6). Subjects should therefore monitor the current partner's decision in the preceding period. That decision should have a significant effect on the action chosen.

Hypothesis 7. *The partner's decision in the preceding match should significantly influence a subject's decision behaviour in the reputation treatment without the outside option.*

When the outside option is included and behavioural histories are available, the duration of the punishment phase increases to two periods (Proposition 7). Therefore, subjects should monitor a partner's decision two periods ago only when the outside option is included.

Hypothesis 8. *The current partner's action choice two periods ago should affect an individual's decision only in the presence of the outside option.*

Finally, it follows from the theoretical analysis that subjects should cooperate after committing a defection in order for their deviation to be forgiven.

Hypothesis 9. *When personal histories are observable subjects should repent following a deviation.*

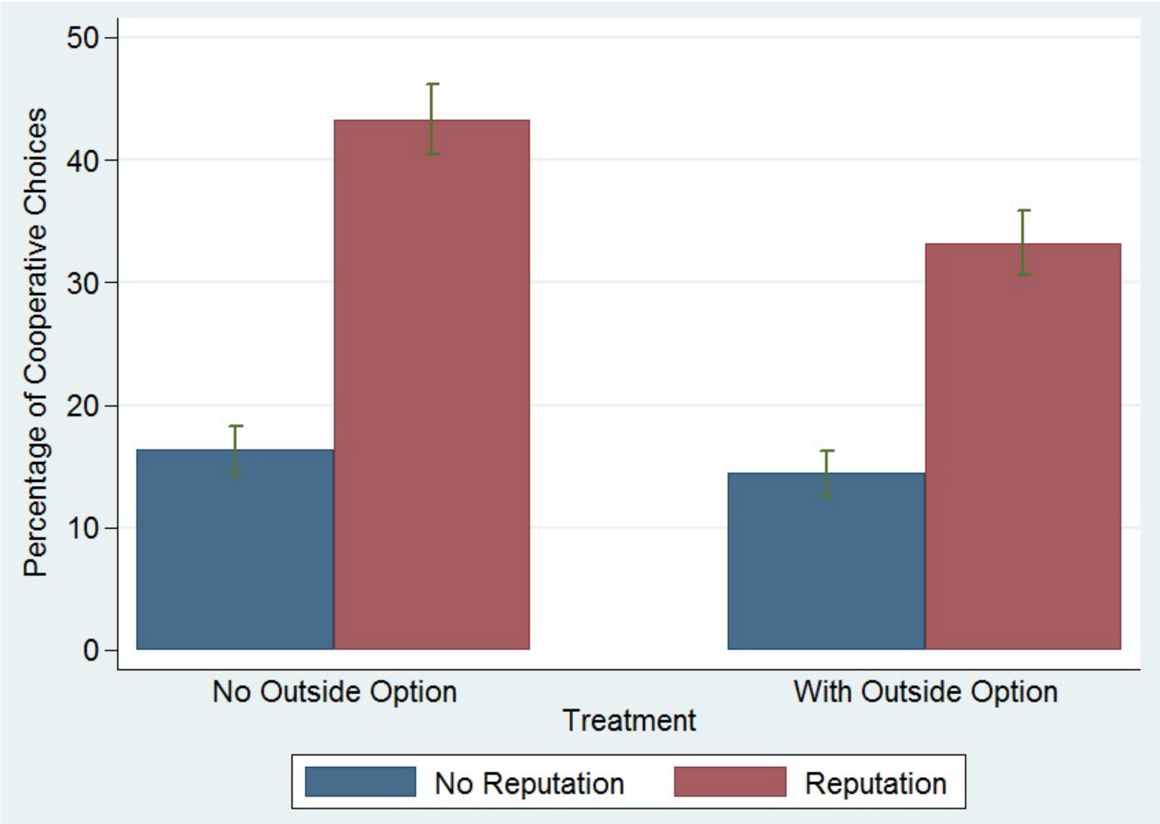

Figure 1.2 Average Cooperation by Treatment

1.4 The Results

1.4.1 Non-parametric Analysis

I begin the analysis of the data by studying the overall differences between individual experimental treatments. I aggregate all cycles by treatment configuration and analyse the differences in cooperation at an individual level (Figure 1.2).

Cooperation levels across the two private monitoring treatments are nearly identical (16.3% without the outside option and 14.25% when the outside option is included). A Mann-Whitney test fails to reject the null-hypothesis of equality (z=0.104, p=0.917).

Result 1. *There is no significant difference in cooperation across the private monitoring treatments.*

As expected, subjects do not use the outside option in the absence of behavioural histories. This is not a surprising result, since opting out is likely perceived as a dominated action choice. Also, there is no strategic value in choosing not to play when subjects only observe the outcomes of their own matches.

Result 2. *The outside option is never used under private monitoring.*

Next, I examine the differences in cooperation between the two private monitoring treatments and those with reputation. The observability of personal labels is manipulated in a within-subjects design. In the prisoner's dilemma, cooperation averages 42.6% for the reputation treatment. This represents a significant increase in cooperative behaviour over the corresponding treatment with private monitoring (z=2.023, p=0.043, Wilcoxon Sign-Rank test). A similar trend is also apparent for the framework with the outside option. Here the reputation treatment achieves a cooperation level of 32.6%, which is significantly higher than the level of cooperation observed in the absence of personal histories (z=2.023, p=0.043, Wilcoxon Sign-Rank test).

Result 3. *Cooperation increases significantly in the presence of reputational concerns.*

It is surprising that cooperative behaviour is maximized in the prisoner's dilemma environment. Based on the theoretical analysis, cooperation should be easier to maintain when the outside option is available (Proposition 7). I evaluate the difference between the two reputation treatments using a Mann-Whitney test, which fails to reject the null-hypothesis of equality (z=1.149, p=0.251). However, it is interesting that the difference between the two reputation treatments is not caused by subjects choosing not to play. On the contrary, the option to abstain is disregarded even in the presence of behavioural histories. It therefore seems that the outside option actually induces relatively higher levels of defection. This is an unexpected result, which the theory fails to account for.

Result 4. *There is no significant difference between the two reputation treatments. The outside option therefore does not promote cooperation.*

Given that subjects do not use the outside option even in the reputation treatment, it seems that players do not value the option to punish deviations without tarnishing their own label. Instead, they reject the outside option on the basis that choosing not to play entails a reduced stage game pay-off. Indeed, people prefer the indirect cost of having a defection added to their behavioural record ([26]).

Result 5. *Subjects reject the outside option as a means of punishment.*

1.4.2 Treatment Specific Analysis

I use a random effects probit regression to examine the dynamics of the individual decision to cooperate ([23]). A random effects model is appropriate, as a Hausman test leads to the non-rejection of the null hypothesis. In the regression model, the dependent variable is the decision taken by a subject in each of her stage game pairings. This variable takes on the value of one when a person cooperates, and it is otherwise equal to zero [4].

Table 1.5 Random Effects Probit Regression on the Choice to Cooperate

	Control (Cooperate = 1)	Reputation (Cooperate = 1)	Opt Out (Cooperate = 1)	Opt-Out Reputation (Cooperate = 1)
2nd Cycle	−0.10	0.51***	0.37	0.20
	(0.31)	(0.15)	(0.77)	(0.27)
3rd Cycle	−0.33	0.72*	−0.63**	0.65**
	(0.39)	(0.39)	(0.30)	(0.32)
Period	−0.17***	−0.20***	−0.07*	−0.13**
	(0.04)	(0.04)	(0.04)	(0.05)
2nd Cycle x Period	0.00	0.13***	−0.02	0.07**
	(0.01)	(0.03)	(0.04)	(0.03)
3rd Cycle x Period	0.01	−0.02	−0.05	−0.01
	(0.05)	(0.04)	(0.06)	(0.09)
Length Previous Cycle	−0.00	0.22***	−0.38	−0.03
	(0.13)	(0.07)	(0.44)	(0.06)
Restart	0.28*	0.31***	0.44	0.41**
	(0.15)	(0.10)	(0.27)	(0.18)
Intercept	−0.39	−1.44**	2.70	−0.24
	(1.51)	(0.68)	(4.58)	(0.35)
lnsig2u				
_cons	−0.50*	0.29	0.29	0.41
	(0.21)	(0.48)	(0.23)	(0.29)
N	1,330	1,180	1,379	1,235

Robust Standard Errors in Parentheses — * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

It is apparent from table 1.5 that the dynamics of cooperative behaviour vary significantly across different treatments. In the communities with private monitoring subjects quickly learn to defect and

[4]The use of a random effects logit model would be equally appropriate and produces very similar results. The choice of a probit model is due to the underlying standard normal distribution.

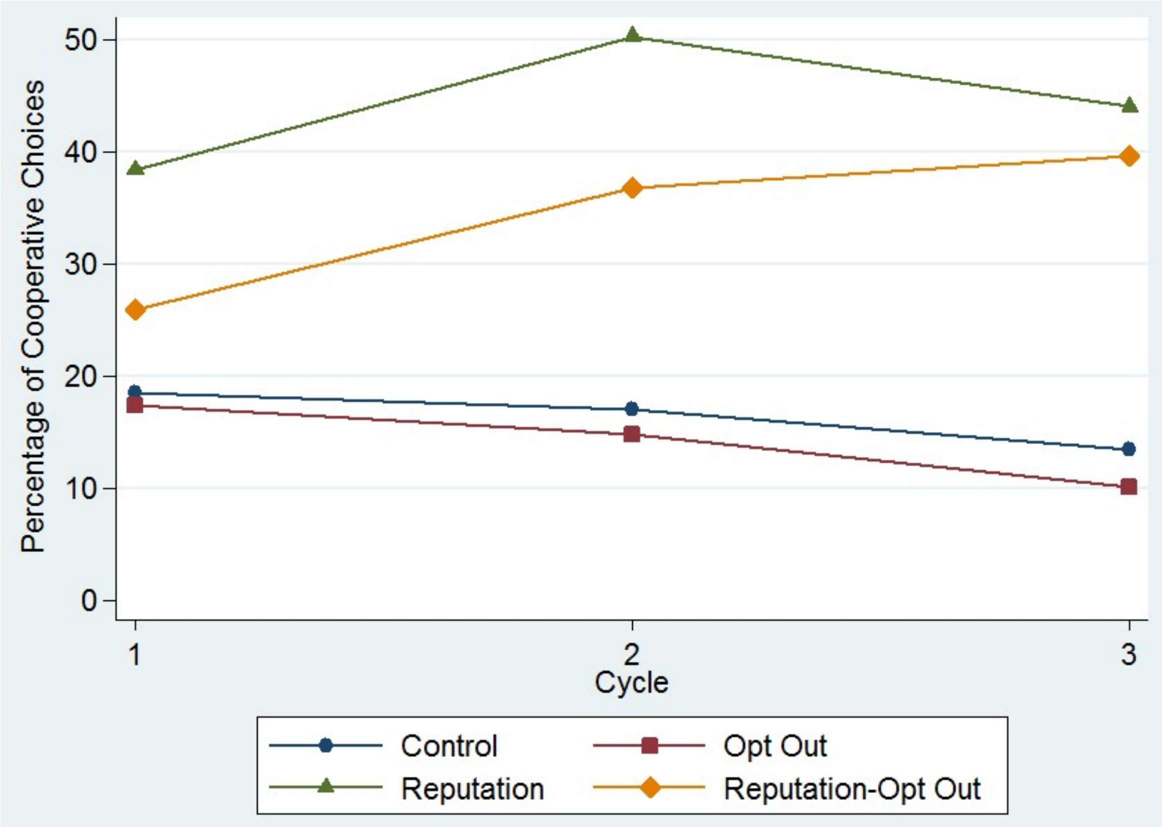

Figure 1.3 Average Cooperation across Consecutive Cycles

cooperation therefore decreases across consecutive periods. By contrast, average cooperation has an increasing trend for the treatments that feature personal histories. This behavioural pattern offers support for the theoretical result that reputational incentives deliver efficiency gains to the economy ([18]).

Result 6. *Some cooperation emerges in the economies with personal histories. Cooperation levels in later cycles are higher compared to the first cycle.*

The differences in cooperation across consecutive cycles shown in figure 1.3, suggest that in the presence of reputational concerns subjects embrace the incentive structure of the indefinitely repeated game. On the other hand, people do not recognise these incentives in the private monitoring communities. In fact, it seems that under private monitoring subjects regard the various stage games as separate one shot matches ([51]). This explains not only the very low cooperation levels within each cycle, but also the steady decline in cooperation over consecutive cycles.

Result 7. *Little cooperation is evident in the private monitoring economies. It appears that subjects in those communities treat the repeated game as a mere sequence of one shot interactions.*

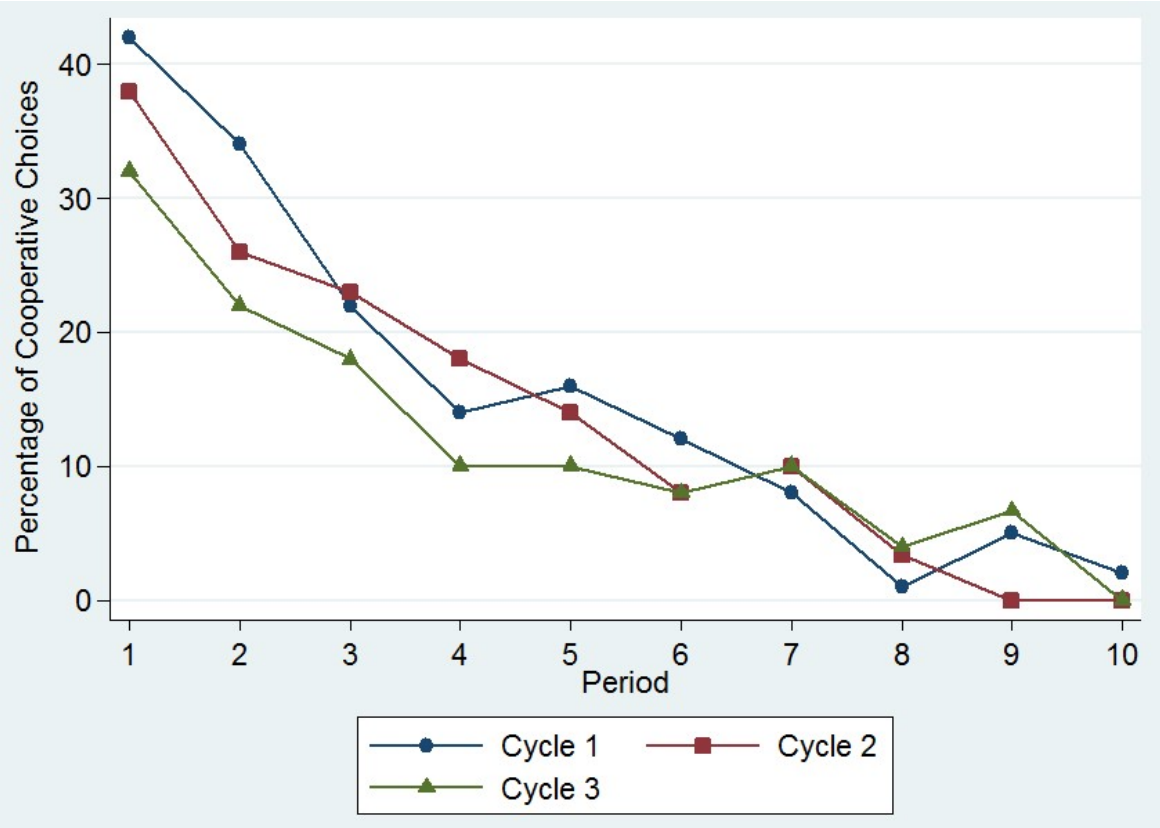

Figure 1.4 Cooperation in the Control Treatment

The Private Monitoring Treatments

I analyse each of the treatment configurations separately to examine any dynamic changes in co-operative behaviour both within and across individual cycles. I start with the control treatment, which provides a behavioural benchmark. It appears from figure 1.4, that there are no differences in the observed behavioural patterns across consecutive cycles. This implies that in the absence of reputational considerations there are no significant learning effects associated with repeated play ([8]).

It is evident that cooperation declines rapidly within each cycle. However, there is a significant increase in cooperation at the start of every new cycle. The substantial restart effect indicates that several subjects cooperate at the beginning of a new cycle and revert to defection shortly afterwards ([46]). This trend is consistent with a contagious punishment regime, which stipulates that subjects should cooperate initially and switch to defection as soon as a single deviation is witnessed.

Result 8. *There is evidence of subjects using grim trigger strategies in the control treatment. That observation is consistent with people following a community punishment scheme.*

The private monitoring treatment with the outside option strongly resembles the control treatment in terms of the behavioural dynamics. This is not a surprising result, seeing as the outside option is

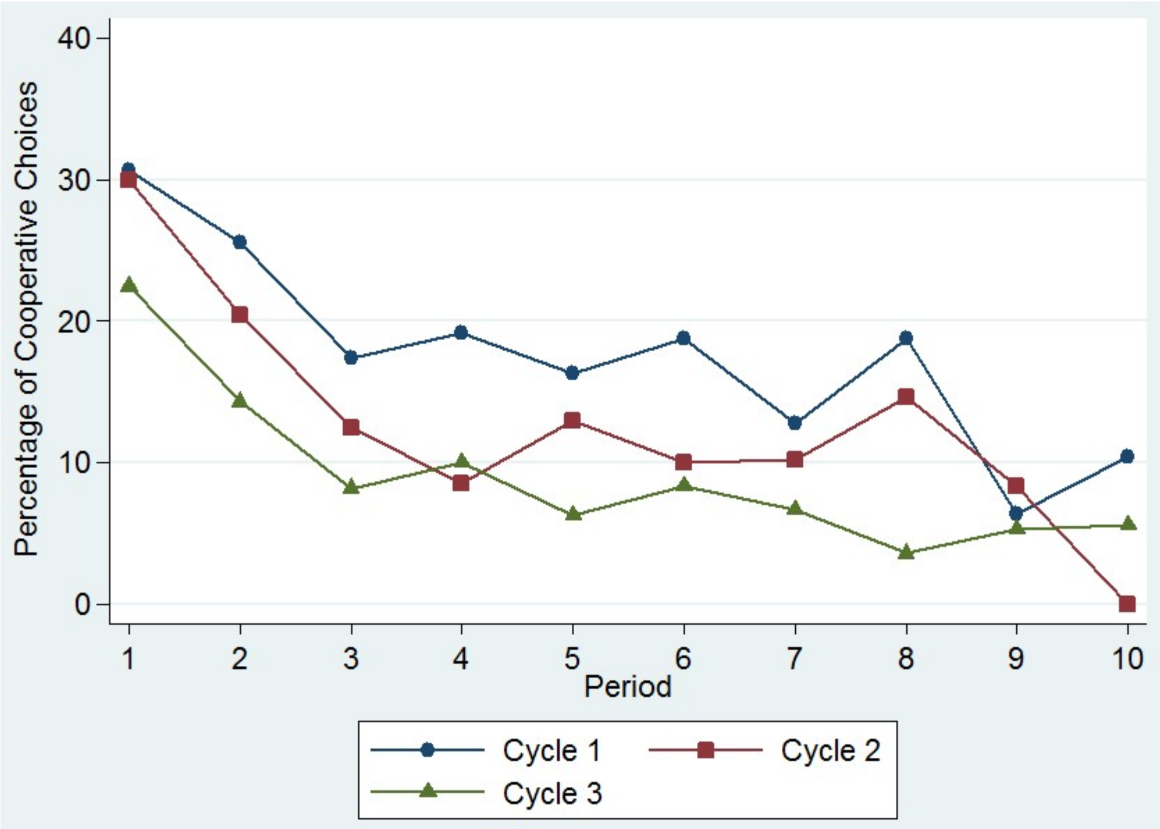

Figure 1.5 Cooperation in the Opt Out Treatment

never played by any of the subjects. Figure 1.5 shows that cooperation levels decline within each of the cycles. However, the restart effect, which is apparent in the prisoner's dilemma framework, is far less evident when the outside option is included. As a result, cooperation levels do not recover as significantly at the start of each new cycle. This is likely due to the pay-off for mutual defection seeming relatively more attractive in the presence of the outside option.

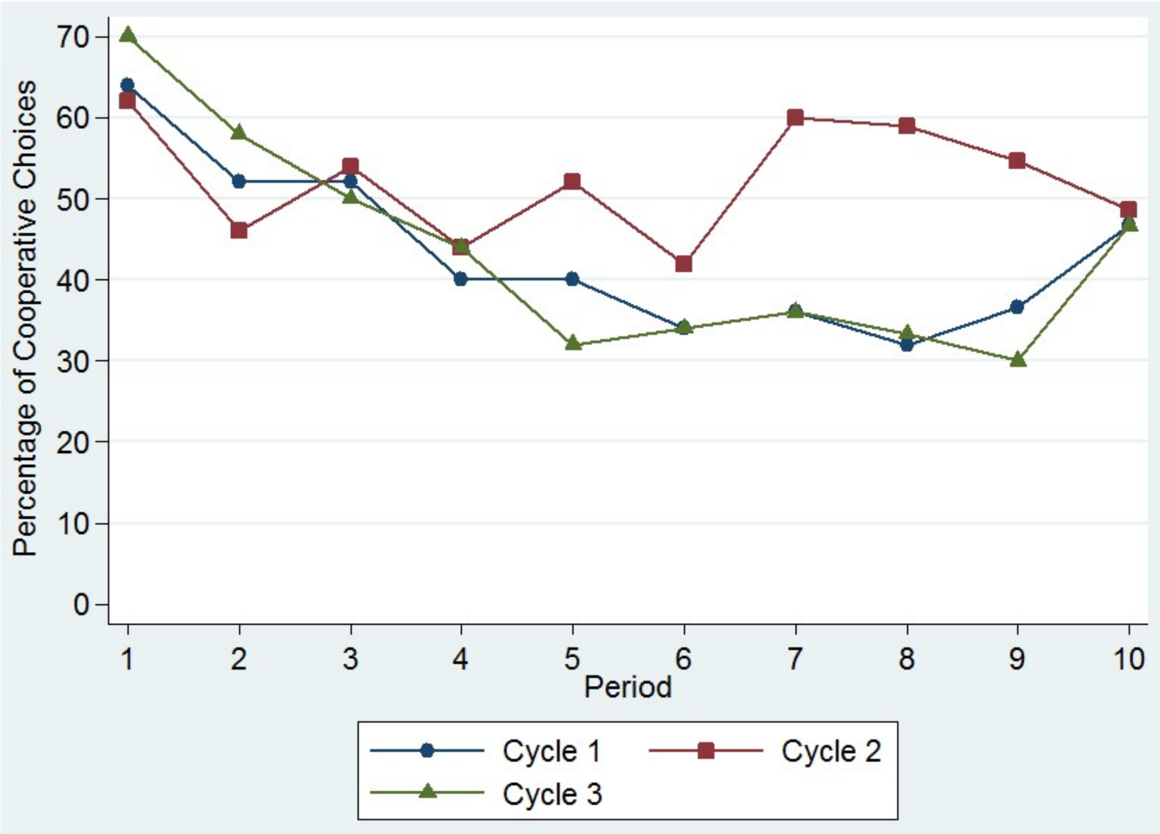

Figure 1.6 Cooperation in the Reputation Treatment

The Reputation Treatments

Cooperation levels are significantly more robust in the communities that feature personal histories. For the reputation treatment, which introduces personal labels in the prisoner's dilemma framework (Figure 1.6), the sharp decline in cooperation, apparent in both private monitoring treatments, is not observable. Instead, cooperation only declines very modestly over the course of consecutive periods. Also, the data shows that there is a significant increase in cooperation in the second of the three cycles.

It is evident that over time subjects realize the potential benefits associated with sustained cooperation in the economy. This illustrates that people comprehend the interdependence of individual one shot matches when their past behaviour is observable. When personal histories are available, subjects' behaviour is not only affected by the start of a new cycle, but also by the duration of the preceding interaction. This trend is non-existent in the private monitoring treatments, where the duration of the previous cycle has no discernible impact. This again confirms that subjects do not consider themselves to be part of a wider community when they only observe the action profiles and outcomes of the matches they play.

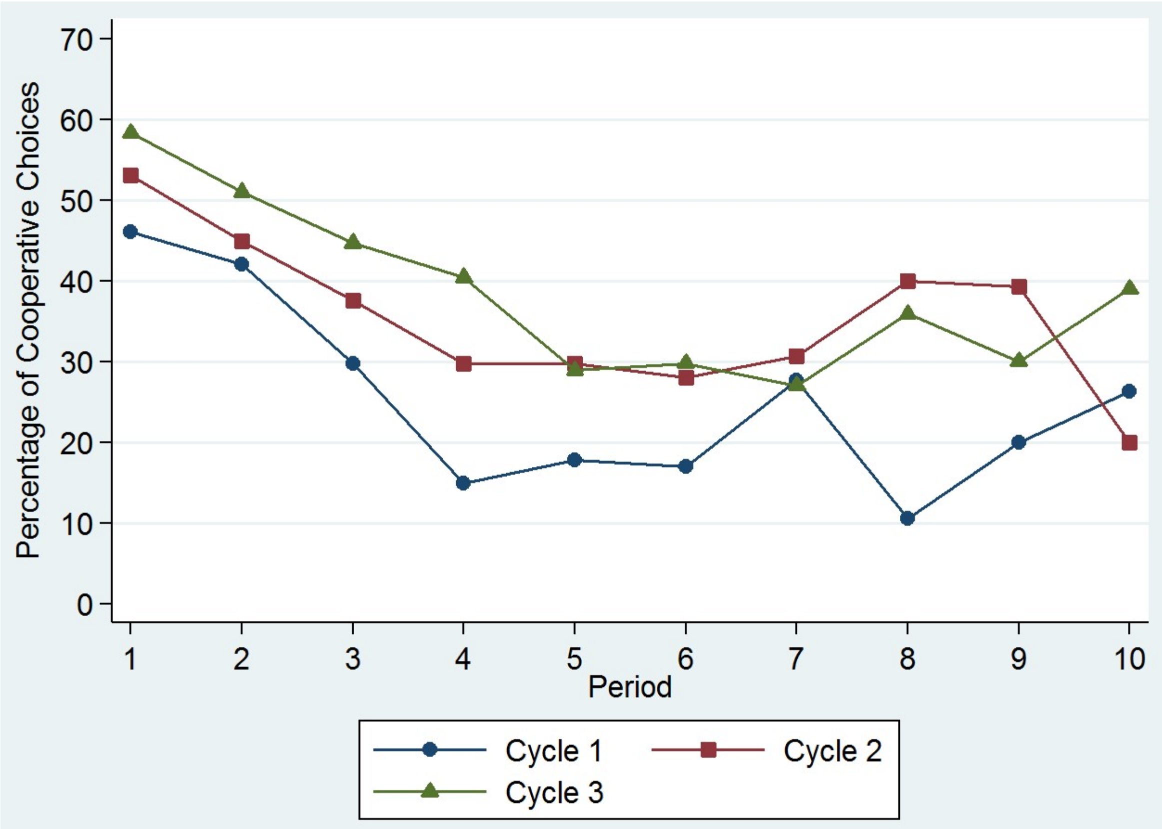

Figure 1.7 Cooperation in the Reputation-Opt Out Treatment

Result 9. *In the presence of personal histories cooperation levels only decline modestly within each cycle. Cooperation increases in the later cycles, as subjects embrace the benefits of achieving a mutually beneficial outcome.*

For the reputation treatment with the outside option depicted in figure 1.7, the dynamics of cooperative behaviour resemble those observed in the prisoner's dilemma. Specifically, cooperation is again far more stable than under private monitoring. As before, there is an increase in sociable behaviour across the consecutive cycles. Interestingly, the increase in cooperation is again most obvious in the second of the three cycles. However, cooperation levels are consistently below the prisoner's dilemma environment. As mentioned previously, this difference between the two treatments is not directly linked with the outside option, which is completely disregarded by the subjects.

1.4.3 Individual Reputation Effects

I now examine the two reputation treatments in detail. I seek to measure the effect of a person's behavioural history on their partner's decision making. In order to conduct this analysis, I pool the data for both reputation treatments. This allows me to study how subjects respond to a partner's decision behaviour in the previous rounds. In addition, I investigate whether the outside option has any discernible effect on how subjects respond to a partner's reputation. As before, a Hausman test performed on the data fails to reject the null hypothesis and I therefore opt for a random effects probit model in my analysis. However, a random effects logit regression produces very similar parameter estimates and supports the same conclusions [5].

The variable *Partner's Reputation* in the regression model gives the number of cooperative choices relative to the total number of matches played. As a result, the variable strictly takes on values between zero and one. It is equal to zero if a person has not previously cooperated. On the other hand, the variable equals one if a player has cooperated in each of their previous matches. The effect of the two most recent decisions is measured in a separate regression model, to avoid issues of multicollinearity. The variable *Partner Lag 1* refers to the current partner's choice in the previous period, whereas *Partner Lag 2* is their decision two rounds ago.

In the model, I control for whether a subject has defected in her previous match. This is because a deviation in the preceding round would strongly affect her decision in the current period of play. It follows from the theoretical analysis that subjects should repent after a committing a defection, so as to be forgiven in the subsequent period (Hypothesis 9). In addition, I include a player's decision in the first period of the cycle. I do this in order to address the endogeneity problem, which arises from the variable *Having Defected in Previous Match* being perfectly correlated with the dependent variable at lag one.

It is evident from table 1.6 that there is a significant relationship between a subject's decision to cooperate and the behavioural history of their playing partner. The more times a player has cooperated in the past, the more likely her partner is to cooperate in the current match. The magnitude of this relationship increases over time, which is an interesting observation. As each cycle continues, the number of decisions made by any subject increases steadily. Therefore, it is intuitive that a person's reputation would have a more significant effect in the later periods of a cycle, when more information about their past behaviour is available.

Result 10. *Behavioural histories have a significant effect on individual decisions. The magnitude of this relationship increases over time.*

[5] The probit model is favoured due to the underlying standard normal distribution.

Table 1.6 Random Effects Probit Regression on the Choice to Cooperate

	Coefficient (Cooperate = 1)	dy/dx (Cooperate = 1)	Coefficient (Cooperate = 1)	dy/dx (Cooperate = 1)
2nd Cycle	0.08	0.06***	0.21	0.05***
	(0.20)	(0.02)	(0.18)	(0.02)
3rd Cycle	0.18	−0.00	0.10	−0.01
	(0.30)	(0.03)	(0.35)	(0.03)
Period	−0.13***	−0.03***	−0.05	−0.01***
	(0.04)	(0.01)	(0.03)	(0.00)
2nd Cycle x Period	0.07*		0.01	
	(0.03)		(0.03)	
3rd Cycle x Period	−0.04		−0.02	
	(0.05)		(0.05)	
Outside Option Treatment	−0.16	−0.04	−0.26	−0.06
	(0.17)	(0.04)	(0.16)	(0.04)
Length Previous Cycle	0.00	0.00	−0.02	−0.00
	(0.03)	(0.01)	(0.03)	(0.01)
Partner's Reputation	1.03***	0.24***		
	(0.10)	(0.03)		
Partner's Reputation x Period	0.07*	0.02**		
	(0.04)	(0.01)		
Having Defected in Previous Match	−0.51***	−0.12***	−0.74***	−0.16***
	(0.09)	(0.02)	(0.11)	(0.02)
Cooperation in Period One	0.73***	0.17***	0.89***	0.20***
	(0.11)	(0.02)	(0.15)	(0.03)
Partner Lag 1			0.93***	0.20***
			(0.11)	(0.02)
Partner Lag 2			0.35	0.08
			(0.29)	(0.07)
Partner Lag 2 x Outside Option Treatment			0.25	0.05
			(0.21)	(0.05)
Intercept	−0.65***		−0.86***	
	(0.25)		(0.19)	
lnsig2u				
_cons	−0.34		−0.67	
	(0.29)		(0.35)	
N	2,115	2,115	1,750	1,750

Robust Standard Errors in Parentheses — * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

It follows from the model that when personal histories are observable, a defection should result in the deviator being penalized. In the conventional prisoner's dilemma framework, the only means of punishment is a defection in the stage game. Therefore, the decision made by the current partner in the preceding match should have a significant effect on an individual's action choice. The regression shows that this trend is in fact apparent in the data. Specifically, the current partner's play in the previous period is one of the key determinants of a person's decision. Furthermore, it seems that the current partner's action choice two periods ago has no statistically significant effect. Therefore, participants do indeed follow the strategy profile outlined in proposition six. This strategy stipulates that a defection by any player should be punished for the duration of one period only.

Result 11. *There is evidence that subjects punish a defection by their partner in the previous period. However, the current partner's choice two rounds ago is largely ignored.*

When the outside option is available, the duration of the punishment phase increases to two periods (Proposition 7). Therefore, the current partner's action choice two periods ago should have a significant effect whenever the outside option is included. However, it is apparent from table 1.6 that no such trend exists. In particular, the choice made by the current partner two periods ago seems to play no role in any of the treatments. It therefore appears that subjects generally have a short memory when it comes to penalizing deviations.

Result 12. *The data shows that even when the outside option is available, subjects have no interest in penalizing defections more than one period ago.*

Finally, people show no signs of repenting after committing a deviation. Instead it seems that an individual is far less likely to cooperate after defecting in her previous match.

1.5 Concluding Remarks

This experiment seeks to study long run equilibria in communities made up of strangers. The subjects are anonymous agents, who are randomly matched for the duration of one period at a time to play the prisoner's dilemma. Players have no means of communicating with one another and they face a variety of opponents, whose identities they do not observe. Achieving significant levels of cooperation in this setting is extremely difficult, as subjects have no means of enforcing cooperation without committing a defection themselves. [29] devises a theoretical framework by which cooperation in a community can be sustained even in the presence of these extreme informational barriers. The model shows how cooperative behaviour can be maintained through the adoption of a social norm, which specifies a desirable rule of conduct and a punishment mechanism. [5] produces behavioural evidence in support of [29]. The paper finds that economies, which are made up of small numbers of people, can effectively overcome significant hurdles and achieve high levels of social efficiency.

However, the results of this experiment cast significant doubt over the robustness of that result. In particular, I find that for larger groups of anonymous agents, cooperation levels in the economy are generally very low ([43]). In the absence of personal histories, subjects quickly learn to defect in the repeated game. This is likely due to the fact that they do not consider themselves as part of a community when they observe only the outcomes of their own matches. Instead, it appears that in the presence of such informational barriers, agents treat the repeated game as a series of one shot encounters. Also, there is no evidence of cooperation rates in the community increasing over consecutive cycles. Therefore, it seems that the trends observed in [5] are applicable only in the context of very small communities.

I further develop this result by conducting an experimental test for the theoretical model in [53]. I investigate whether individual labels, which allow for the observation of a subject's previous decisions, promote sociable behaviour. The label allows for information about a player's history to be transmitted locally. Therefore, a person's behavioural record can be updated without the knowledge of the entire community. Based on the theoretical framework of [53], the introduction of such labels results in an equilibrium with records of length one, by which cooperation can be maintained. The display of personal histories allows for defectors in the community to be identified, and for opportunism to be punished accordingly.

Indeed there is evidence in the data that subjects closely monitor the partner's choice in the preceding round, whereas no attention is paid to their decision two periods ago. This individual behaviour is very much in line with the theoretical result that any defection by a subject should be punished for one period only. The experimental evidence shows that the introduction of behavioural histories delivers significant efficiency gains to the economy. Yet, it is noteworthy that even in the presence of personal labels, cooperation only accounts for 42.6% of all individual choices. This result implies that in large communities reputational incentives alone are not sufficient to uphold cooperative behaviour.

In order to further analyse the dynamics of reputational concerns, I add an outside option to the prisoner's dilemma. Whereas that outside option has little value under private monitoring, it is game theoretically relevant when behavioural histories are displayed. In particular, the outside option allows subjects to penalize defections without the punishment having a detrimental effect on their own reputation. In the prisoner's dilemma a deviation can only be punished by defecting on the deviator. However, the outside option enables a player to penalize opportunism without being subjected to punishment action in the future. Therefore, the outside option effectively eliminates the reputational cost of punishing deviations. This is reflected in the opt out pay-off, which is lower than that assigned to mutual defection in the prisoner's dilemma. As a result, the punisher pays a price when opting out, which permits her to preserve her own label. However, it is evident that the participants are strongly opposed to bearing a direct cost in order to maintain a favourable reputation ([26]). Instead, subjects choose to pay the implicit price of having a defection added to their behavioural history. This trend is also reflected in the fact that there is no significant difference in cooperation between the two reputation treatments.

Overall, the results of this study show that high levels of cooperation are difficult to achieve in an economy consisting of a large number of strangers. My findings therefore contradict the observation that cooperation can emerge even in the presence of extreme informational restrictions. This experiment has shown that reputational incentives influence individual behaviour in the way economic theory would predict. However, it seems that information about the past behaviour of others has a less significant effect on individual decision making than is commonly understood.

Chapter 2

Should I Play or Walk Away: Outside Options in the Prisoner's Dilemma

2.1 Introduction

The threat not to interact with someone is commonly used to achieve socially optimal equilibria. The option to unilaterally terminate a partnership is therefore a force in many relationships studied by economists ([12]). Employees can quit their jobs, couples can split up, and customers can withhold their business following a bad experience. Such interactions have traditionally been modelled using the prisoner's dilemma ([28]). However, an obvious limitation of this approach is that the option to abstain from a relationship is not incorporated. Instead, players in repeated interactions only have access to in-relationship punishments. For example, workers can take industrial action until their demands are granted, whereas couples can argue and subsequently re-conciliate. Although these outcomes are typically socially optimal, it is nevertheless important to incorporate the option not to play in the study of indefinite partnerships ([55]).

The economic analysis of being able to opt out starts with [42]. That paper shows the ability to choose between playing the prisoner's dilemma and a risk free outside option can increase market efficiency and benefit cooperators. When cooperators are more willing to play, they will interact more regularly with one another. As mutual cooperation produces the socially efficient outcome, optional participation will therefore increase market efficiency. The experiment presents six strangers with a set of one shot prisoner's dilemma games. They interact only once with each of the other five individuals. Under the conventional choice regime subjects are obliged to either cooperate or defect, whereas under a ternary decision rule they have the additional option not to play. The stage game pay-offs vary across the individual matches each subject plays. However, the incentive structure is always such that a player should abstain, if they expect their opponent to defect.

[24] examines the effect of outside options in a finitely repeated game. In particular, seven people, in full view of each other, play ten consecutive prisoner's dilemma games with each of the other subjects. Hence, the experiment features a finite horizon framework with fixed matching. The attractiveness of the outside option is varied, as the pay-off associated with choosing to abstain is either greater or smaller than the Nash equilibrium pay-off in the prisoner's dilemma. The results of the experiment suggest that an outside option, which yields a reward greater than mutual defection, is used by the subjects to systematically exclude defectors and enforce cooperation. By contrast, the use of the unattractive outside option is negligibly low.

[55] studies the effects of implementing an outside option in a fixed partnership, which continues indefinitely. The experiment has two players engage in a joint-effort task with imperfect monitoring. In each round the two partners simultaneously select an effort level. Based on the resulting action profile, a public signal indicating success or failure is observed by both players. The probability of success is a function of the selected actions, such that each player has an incentive to make little effort. However, the probability of the socially optimal outcome is maximized when both players exert high effort. The experiment modifies the repeated game by adding the additional option to walk away. This allows either of the two players to terminate the partnership. The value and relative distribution of the outside option vary across the experimental treatments. In particular, choosing to opt out either results in both partners receiving symmetric pay-offs, or alternatively the pay-offs are asymmetric. Importantly, the findings show that cooperation rates remain unaffected by the value of the outside option.

The effect of voluntary participation on cooperative behaviour is also examined in [41]. That experiment features a one-shot public goods game played by groups of two players. Within each group, subjects receive an initial endowment, which they can allocate between a joint and a personal account. For each token paid into the group account both subjects earn two points. However, each token invested into the private account yields three points to the individual alone. The experiment simulates three versions of this game. Whereas there is no outside option available in the control treatment, the two experimental treatments feature the option to abstain. It varies whether the subjects can choose not to play, either before or after the allocation of tokens has been finalized in their group. The two opt out treatments therefore differ in the level of information available when deciding to either play or abstain. The experiment shows that the prospect of non-interaction is powerful in eliminating free riding.

The basic structure of this study resembles that in [49]. That study aims to determine the economic impact of public disclosure in a modified prisoner's dilemma with anonymous random matching. The experiment varies the extent to which subjects observe the past behaviour of their current playing partner. Thus, subjects either have no information about the previous decisions of an opponent, or they can observe their behavioural history. However, the degree of disclosure changes across individual treatments, as people either observe a complete decision history, or a restricted information set. The outside option yields the Nash equilibrium pay-off for the prisoner's dilemma, and each subject is exposed to only one treatment configuration. Thus, the experiment does not study the economic importance of having the option not play.

This chapter adds to the existing literature by analysing the impact of two different outside options in an indefinite interaction with random matching. I simulate an economy, which consists of ten anonymous strangers repeatedly playing prisoner's dilemma games. The economy features a stochastic endpoint and the agents are randomly rematched after each stage game they play ([29] and [13]). I modify the conventional prisoner's dilemma by making all interactions between the players entirely optional. In particular, subjects can choose to cooperate, defect or abstain. The outside option is designed to give a symmetric pay-off at all times. However, depending on the treatment specification, the option to abstain either produces a pay-off smaller than or equal to the Nash equilibrium. When it delivers a pay-off lower than mutual defection, I refer to the outside option as costly. Otherwise, the option not to play is said to be costless. In addition, I manipulate whether or not behavioural histories are observable in the continuation game. In the absence of personal labels, subjects only observe the action profiles in their own matches. On the other hand, when behavioural records are observable, subjects can review all past decisions of their current playing partner.

The reputation parameter in this experiment is cumulative. Therefore, as the game continues people develop a behavioural history, which gives a true account of all previous decisions. When two players meet they only receive first order information about each other's behavioural history, like in [53][1]. I deliberately choose to model reputation in this way, as it would intuitively be difficult to maintain cooperative behaviour. This is because without the outside option it is impossible to differentiate defectors from those, who punish deviations. However, by introducing the option to walk away, I permit subjects the opportunity to penalize bad behaviour without incurring a reputational cost ([19]).

This research seeks to address the following questions: Do people comprehend the strategic value of the option to walk away? Does the option to abstain encourage cooperation among the subjects? To what extent does the use of the outside option depend on the guaranteed pay-off? Are cooperation rates in the community affected by the pay-off associated with the outside option? Do cooperators use the option not to play to exclude defectors? How does the option to abstain influence the perception of behavioural histories?

The results of this study show that without reputational incentives cooperation levels in the economy are very low. Also, subjects make no use of the outside option when their past actions are not observable. It is evident that the introduction of personal labels changes that behavioural pattern ([9]). Cooperation rates are significantly higher for those communities, which feature behavioural records.

[1]Players only observe their current partner's decisions in all previous rounds of the game.

However, it is noticeable that players only take advantage of the outside option when it is costless not to play. The costly outside option is never utilized at all. The use of the costless outside option is largely in line with theoretical predictions, as subjects only abstain from interactions with players, who have recently defected. Finally, it is apparent that a costless option to abstain causes subjects to review behavioural histories in more detail.

This chapter is organized as follows: Section two outlines the theoretical model in detail. Section three describes the experimental environment and formulates the hypotheses. Section four analyses the experimental data and discusses the findings. Section five summarizes the key results and concludes the study.

2.2 The Model

I use generic notation to denote a subject's action choices and pay-offs. Throughout my analysis I assume that all players are rational agents, who seek to maximize their utility.

Table 2.1 Notation in the Theoretical Analysis

	Cooperate (C)	Defect (D)	Opt Out (O)
Cooperate (C)	π_{CC},π_{CC}	π_{CD},π_{DC}	π_{OO},π_{OO}
Defect (D)	π_{DC},π_{CD}	π_{DD},π_{DD}	π_{OO},π_{OO}
Opt Out (O)	π_{OO},π_{OO}	π_{OO},π_{OO}	π_{OO},π_{OO}

The possible stage game pay-offs are $\pi_{CC},\pi_{CD},\pi_{DC},\pi_{DD}$ and π_{OO}. The first subscript refers to the action choice of the row player and the latter subscript to that of the column player. In the conventional prisoner's dilemma joint cooperation leads to the Pareto efficient outcome being realized. However, as $\pi_{DC} > \pi_{CC}$ and $\pi_{DD} > \pi_{CD}$, that game constitutes a social dilemma, such that each player has an incentive to unilaterally deviate. The guaranteed pay-off associated with choosing not to play depends on whether the outside option is either costly or costless. Whenever the outside option is costly $\pi_{DD} > \pi_{OO} > \pi_{CD}$, whereas for the costless outside option $\pi_{DD} = \pi_{OO}$.

2.2.1 Equilibria in the Stage Game (Costly Outside Option)

Incorporating a symmetric outside option into the prisoner's dilemma results in an additional stage game Nash equilibrium. Apart from mutual defection, a scenario where both players abstain from the interaction is also a stable outcome[2]. However, for the stage game with the costly outside option that Nash equilibrium is strictly Pareto dominated, given that $\pi_{DD} > \pi_{OO}$. It is therefore apparent that in a one shot interaction both players should always choose to defect rather than abstain.

Proposition 8. *When the option not to play is costly, there exists a Pareto dominated Nash equilibrium in the stage game, such that both players opt out.*

2.2.2 Equilibria in the Repeated Game (No Reputation/Costly Outside Option)

I first develop a theoretical framework for the experimental treatment where personal histories are not observable and the outside option is costly. Under this configuration all individual pairings are entirely independent, as players only observe the outcomes of the matches they are directly involved in. Since mutual defection is an equilibrium in each of the stage games, it follows that defection by all players is an equilibrium even in the continuation game.

Proposition 9. *A scenario where all subjects defect in all their matches constitutes a sequential equilibrium of the repeated game.*

[2] See [49] for the generic proof.

Given that the case of both players opting out also constitutes a Nash equilibrium in the stage game (albeit a Pareto dominated one), it is straightforward to extend this outcome to the repeated game. Therefore, all players abstaining from all their matches constitutes an equilibrium in the repeated game, as well.

Proposition 10. *A sequential equilibrium exists in the repeated game, such that all subjects choose to play O in all their stage game matches.*

The socially optimal outcome, where all subjects cooperate in every match, can also be achieved when the game is repeated indefinitely. This cooperative equilibrium is sustained by a contagious punishment mechanism ([29]). In the following analysis I make use of the results in [6] to construct a cooperative equilibrium for my parametrization[3].

I consider a scenario where all subjects follow a strategy, which consists of two components: a desirable rule of conduct (always play C), as well as a punishment mechanism: (always play D). Each of the subjects will revert to the punishment mechanism once a deviation from the behavioural norm has been witnessed. The strategy therefore follows the notion of a grim trigger. Based on this structure, I define the following strategy profile for a given subject *i*:

If subject i adheres to the rule of conduct (always cooperate) in period t, then she will switch to the punishment mechanism (always defect) in period $t+1$, if and only if a defection was witnessed. Otherwise, subject i will continue to follow the rule of conduct.

If all subjects pursue the above strategy, then a single defection will trigger the collapse of cooperative behaviour in the economy. The initial deviation will spread in the community akin to a contagious disease and erode the future pay-offs of all subjects, as shown in figure 2.1. The erosion of earnings in turn provides an incentive to maintain cooperation. Therefore, if agents are patient and exhibit sufficiently high discount rates, none of them will have any motivation to deviate from the rule of conduct. Formally, this can be formally expressed as follows:

$$\frac{\pi_{CC}}{1-\beta} \geq \frac{1}{1-\beta}[\pi_{DC}\phi_k + \pi_{DD}(1-\phi_k)] \tag{2.1}$$

Subject i's discount rate is given by $\beta \ \varepsilon \ (0,1)$, the continuation probability of the repeated game. The parameter ϕ_k is the contact rate in the economy, which I define below. The left hand side of equation (2.1) represents the continuation pay-off to subject *i*, if she always cooperates. On the other hand, the right side of that equation gives her expected pay-off if she chooses to defect. In order for a constructive equilibrium to exist, the pay-off from cooperation must be greater than the projected gain from defection.

[3]The proof is available in Section 2 of that paper.

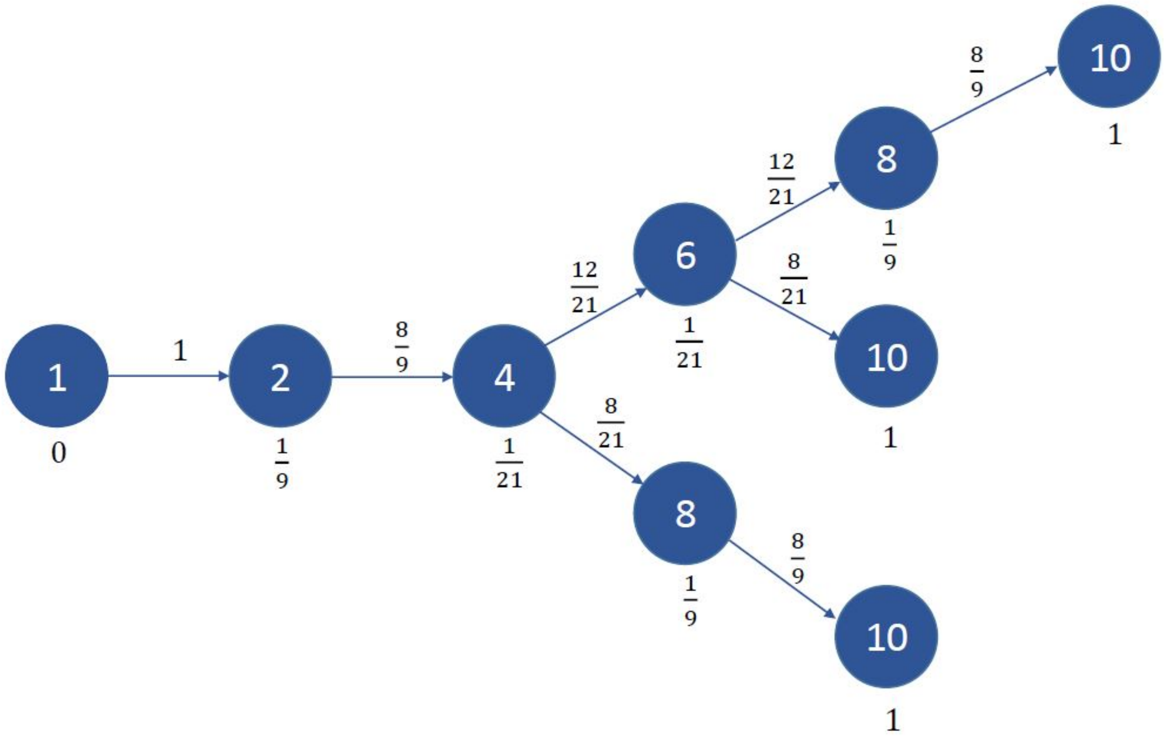

Figure 2.1 The Spread of Deviations in the Random Matching Community

In equation (2.1), ϕ_k represents the contact rate in the economy. This is the rate at which a given defector, subject i, anticipates to meet cooperators in the continuation game. To obtain this rate, I first calculate the number of cooperators the defector expects to be paired with:

$$\sigma_k + \beta \sum_{k'=k}^{N} Q_{kk'}\sigma_{k'} + \beta^2 \sum_{k'=k}^{N} Q_{kk'}^2 \sigma_{k'} + .. = e_k^T (I - \beta Q)^{-1} \sigma \tag{2.2}$$

The economy is comprised of N subjects. In a given period t, there are k defectors in the community, as well as $N-k$ cooperators. Therefore, σ_k is the probability that subject i is matched with a cooperator in period t, given the current state of the economy. Since subjects are randomly and bilaterally rematched after each round of the interaction, $Q_{kk'}$ gives the probability that in the subsequent round the number of defectors increases from k to k'. More specifically, $Q_{kk'}$ is the probability that in the following period $k'-k$ mixed pairings are realized. A mixed pairing occurs when a defector is matched with a cooperator under the random matching protocol. It is now straightforward to interpret the meaning of $Q_{kk'}\sigma_{k'}$. This expression simply yields the probability that subject i will be paired with a cooperator in any future round of the interaction. However, this probability hinges on a pairing being realized, which increases the number of defectors from k to k'. Finally, as the economy is not infinite in duration, it is necessary to discount any future stage game in accordance with the continuation probability β. The vector e_k^T captures the present state of the random matching community. Q is a triangular diffusion matrix, which describes the spread of deviations in the economy. The parameter σ is a vector containing the conditional probabilities of being involved in mixed matches ([6]).

The economic intuition captured in equation (2.2) is that even though the horizon of the interaction is infinite, each cooperator in the economy can only be met once. This is a key characteristic of any contagious punishment protocol. Whenever a defector meets a cooperator, the latter subsequently turns into a defector. It follows therefore that once an initial defection has occurred, the community will, in finite time, reach a state where all subjects follow the punishment protocol. It is apparent that the number of cooperators any individual defector can meet is restricted by $(1-\beta)^{-1}$. This is simply because the expected number of meetings with cooperators can intuitively not exceed the expected lifetime of the community. Therefore, in order to obtain ϕ_k, I divide the expected number of encounters with cooperators by the projected number of periods:

$$\phi_k = (1-\beta)e_k^T(I-\beta Q)^{-1}\sigma \tag{2.3}$$

The significance of the contact rate given in equation (2.3) is that it captures the probability that subject i meets a cooperator in any given period of the indefinitely repeated interaction. This rate is calculated on the basis that there are currently $k \geq 1$ defectors in the economy. The contact rate can be interpreted as follows: if there are currently few defectors in the community, the probability of being matched with cooperators in the continuation game will be high. This results in ϕ_k being close to one. Similarly, once there are many defectors in the community, the probability of being matched with cooperators is low. This means ϕ_k is close to zero.

It is apparent that in order for cooperative behaviour to be sustained by a contagious punishment mechanism, it has to be the case that subject i has no incentive to deviate in an environment with no defectors. This means that none of the players have any reason to commit the initial deviation from the rule of conduct, which causes the ultimate collapse of cooperative behaviour:

$$\frac{\pi_{CC}}{1-\beta} \geq \frac{1}{1-\beta}[\pi_{DC}\phi_1 + \pi_{DD}(1-\phi_1)] \tag{2.4}$$

Equation (2.4) stipulates the condition required for the existence of a constructive equilibrium. The condition states that subject i's pay-off from continued cooperation must be greater than her expected earnings, should she choose to commit the first defection.

Finally, it is necessary to show that it is not optimal to slow down the community punishment, once a contagious response has been triggered. Therefore, let subject i be one of $k \geq 2$ defectors currently in the community. The following condition must be satisfied for subject i to have no incentive to slow down the punishment process once it has begun:

$$\sum_{k'=k}^{N} Q_{kk'}[\sigma_{k'}(\pi_{CC}+\beta v_{k'-1})+(1-\sigma_{k'})(\pi_{CD}+\beta v_{k'})] \leq \sum_{k'=k}^{N} Q_{kk'}[\sigma_{k'}\pi_{DC}+(1-\sigma_{k'})\pi_{DD}+\beta v_{k'}] \quad (2.5)$$

The right side of equation (2.5) gives the pay-off earned by subject i if she participates in the punishment action. Given the number of defectors in the economy increases from k to k' with probability $Q_{kk'}$, there are two possible cases to be considered. Firstly, subject i is involved in a mixed match with probability $\sigma_{k'}$. In this case subject i will earn π_{DC} by choosing to defect. Alternatively, subject i is paired with another defector with probability $(1-\sigma_{k'})$. If this outcome is realized, the pay-off she earns is π_{DD}. The variable $\beta v_{k'}$ gives subject i's discounted pay-offs in the continuation game, given that there are now k' defectors in the economy.

The left hand side of equation (2.5) illustrates the case of subject i attempting to slow down the contagious punishment process. In particular, subject i deviates from the punishment action by choosing to play C in the current period (instead of D, which is the punishment action). Again there are two possible cases: subject i is either involved in a match with a cooperator, or otherwise she is paired with a defector. If matched with a cooperator, subject i's play of C will earn her π_{CC} instead of π_{DC}. However, since subject i's partner is not turned into a defector at this stage, the future pay-offs in the continuation game are now greater than before. In particular, they are $\beta v_{k'-1}$ instead of $\beta v_{k'}$. On the other hand, if subject i is paired with a defector, her cooperation will earn her π_{CD} instead of π_{DD}. Yet, in this particular case there is no effect on the continuation pay-offs, as no new infection would have occurred in any case. The use of the outside option in the punishment protocol is not individually optimal. This is because $\pi_{OO} < \pi_{DD}$ in each stage game. Also, it is apparent that opting out after witnessing a defection violates the condition stipulated in equation (2.5).

The conditions specified above are sufficient to confirm the existence of a constructive equilibrium, which is supported by a contagious punishment regime. Equations (2.4) and (2.5) are satisfied by the parametrization used in this study.

Proposition 11. *All subjects cooperating in every period constitutes an equilibrium of the repeated game even in the absence of personal histories.*

2.2.3 Equilibria in the Stage Game (Costless Outside Option)

As before, the outside option produces an additional stage game Nash equilibrium. For the case of the costless outside option that equilibrium is not Pareto dominated, since now $\pi_{DD} = \pi_{OO}$. However, in a one shot game even the costless outside option should never be used. This is because by choosing to abstain a player is guaranteed to earn $\pi_{OO} = \pi_{DD}$. On the other hand, a defection in the stage game will achieve at least π_{DD} with a chance of attaining π_{DC}, should the opponent choose to cooperate.

Proposition 12. *The costless outside option should never be played in the stage game.*

2.2.4 Equilibria in the Repeated Game (No Reputation/Costless Outside Option)

The equilibria from the repeated game with the costly outside option also exist when the outside option is costless. This is due to the fact that subjects again only observe the action profiles for the pairings they are involved in. Hence, there is no possible value in choosing to opt out at any point, as within each stage game opting out is not individually optimal. Also, the use of the outside option as part of a community punishment protocol is not feasible, as this again violates equation (2.5).

Proposition 13. *Identical equilibria exist in both private monitoring treatments.*

2.2.5 Equilibria in the Repeated Game (Reputation/Costly Outside Option)

The discussion so far has shown how, with severe informational restrictions in place, a community can maintain cooperation. This requires strict rules of conduct and ruthless punishment action. However, such an equilibrium seems undesirable, as it is highly fragile and involves the mass punishment of innocent players ([13]). I now extend my analysis to the case where personal histories are observable. Unlike before, individual encounters between the players are not unrelated in this set up. In particular, subjects can now observe a complete decision history for the individual they are paired with in the current round of the interaction. My analysis follows from [2] and [53]. First, I consider the case where the outside option is costly.

Given that personal histories are now available for review, each subject is essentially assigned a label, which gives a true account of her previous actions. When two players meet, they first observe each other's label and then select their respective action choices. Each subject's label is updated every round, based solely on the decision she makes. Thus, a player's label accurately reflects her entire behavioural history. Formally, we have the following set up:

1. In period t, subject i holds a label $z_i(t)$, where $z_i(t) = (a_i(1),...,a_i(t-1))$.
2. When subjects i and j are matched in period t and select their actions $a_i(t)$ and $a_j(t)$, their respective labels will subsequently be revised as follows:

$$(z_i(t+1)) = f(z_i(t), a_i(t)) \tag{2.6}$$

$$(z_j(t+1)) = f(z_j(t), a_j(t)) \tag{2.7}$$

3. At the start of period t, each subject is aware of her own label and can observe that of her opponent.

When two players are matched, they review one another's history repository before selecting an action. Each subject's behavioural record is subsequently updated based only on her own decision. Given this transmission of information in the economy, an equilibrium exists as follows:

1. If two players are matched, who have defected in one or both of the two preceding rounds, they will mutually minimax each other, i.e. they will both abstain from the pairing, and thus, earn π_{OO}.
2. If two players are matched, who have cooperated in the two preceding rounds, they will both play C and therefore earn π_{CC}.
3. Finally, if a player, who cooperated in the previous two periods, is matched with a subject, who defected in either of the two preceding rounds, the co-operator will minimax the defector by choosing not to play. However, the other player will cooperate. Both players will earn the minimax pay-off π_{OO}.

Penalizing a defector is costly even to the punisher. This is because both the punisher and the deviator earn the pay-off π_{OO}. However, as the outside option is used to penalize opportunism, the punisher does not tarnish her own reputation by defecting on the deviator.

I now show that the strategy described above constitutes an equilibrium[4]. I begin by ensuring that it is individually optimal for subject i to cooperate in round t given that she defected in the previous period:

$$(1-\beta^2)\pi_{OO} + \beta^2\pi_{CC} \geq \pi_{OO} \tag{2.8}$$

The right hand side of equation (2.8) captures the pay-off for subject i if she continues to defect. In that case she will earn π_{OO} in each subsequent pairing, as her opponents choose to abstain. The left side of the expression gives the expected pay-off for subject i if she cooperates and her defection is forgiven in period $t+2$ (that is two periods from now). The equilibrium strategy requires that any deviation must be punished for the duration of two periods.

Next, I show that even for a player, who cooperated in her two previous matches, the actions prescribed by the equilibrium strategy are in her best interest:

$$(1-\beta)\min[\pi_{OO}, \pi_{CC}] + \beta\pi_{CC} \geq (1-\beta)\pi_{DC} + \beta[(1-\beta^2)\pi_{OO} + \beta^2\pi_{CC}] \tag{2.9}$$

In equation (2.9) the left side gives the lower bound of a subject's pay-off if she follows the equilibrium strategy. Given that the outside option is available, subject i will only ever choose to cooperate or opt out. It is evident that choosing not to play is costly even to the punisher. However, the option to

[4] For the general proof see [2]).

abstain allows for the punisher to preserve her label. Hence, she will not be penalized in the future. The right side of the expression gives the maximum pay-off subject i can attain if she chooses to deviate from the equilibrium strategy.

Proposition 14. *When behavioural histories are available the outside option is used to target deviators and enforce cooperation.*

2.2.6 Equilibria in the Repeated Game (Reputation/Costless Outside Option)

Finally, I will construct an equilibrium for the treatment, which combines the observability of personal histories with the costless outside option. The necessary conditions for a constructive equilibrium are identical to those outlined above. It follows therefore that any defection by one of the subjects will result in the deviator being penalized for two rounds. As before, no permanent damage will come to the community as a result of an individual defection.

Proposition 15. *Identical equilibria exist in both reputation treatments.*

2.3 The Experiment

2.3.1 Experimental Design

The stage game participants are faced with varies across treatment configurations, depending on whether the outside option is either costly or costless. Also, I manipulate whether or not personal histories are observable.

Table 2.2 Prisoner's Dilemma with Outside Option

	Left (Cooperate)	Right (Defect)	Opt Out
Left (Cooperate)	30,30	5,40	π_{OO}, π_{OO}
Right (Defect)	40,5	15,15	π_{OO}, π_{OO}
Opt Out	π_{OO}, π_{OO}	π_{OO}, π_{OO}	π_{OO}, π_{OO}

The Stage Game: Each of the stage games consists of two players independently and simultaneously selecting an action from the available set of choices. Whenever the outside option is costly, subjects are faced with the stage game in table 2.2, such that $\pi_{OO} = 10$. Otherwise, when the option not to play is costless, $\pi_{OO} = 15$. The numerical entries in each cell give the pay-offs for the row and column player, respectively. For the sake of neutrality the action choices Cooperate and Defect are labelled as Left and Right. Whenever personal histories are available for review, subjects can observe all previous decisions of the person they are currently matched with. No other information about an individual's identity is revealed at any time.

The Indefinitely Repeated Game: An indefinitely repeated game consists of an unknown number of stage games being played by the subjects. This infinite horizon set up is achieved through a probabilistic rule of continuation. Since the continuation probability is constant, the expected number of additional rounds is given by $(1-\beta)^{-1}$. As β equals 0.9, each repeated game is expected to last for eleven rounds in total. The continuation probability and the random matching regime are identical across all sessions and treatments. The exact number of rounds in each repeated game is determined by drawing random integers between 1 and 100 from a uniform distribution, like in ([5]). A draw of 91 or any higher number ends the game. The subjects are anonymous agents, who are randomly matched in every period of the interaction. There are ten subjects in the economy, which results in five pairings being realized every time. These pairings are for only one round, equally likely, and independent over time.

The Experimental Session: Each of the experimental sessions involves ten participants, who in total play six indefinitely repeated games (I refer to them as cycles). The subjects first undergo three cycles without personal histories. These are then followed by another three cycles where personal histories are observable. Within each session, subjects therefore encounter only one of the two outside options.

The Experimental Treatments: This experiment employs a 2x2 design and therefore consists of four different treatments. The individual treatments differ in terms of the guaranteed pay-off associated with the outside option, as well as the observability of personal histories. For each treatment configuration five sessions are conducted in the laboratory[5].

Table 2.3 Comparison of Experimental Treatments

	Costly Outside Option	Costless Outside Option
No Reputation	5 Sessions	5 Sessions
Reputation	5 Sessions	5 Sessions

Two of the experimental treatments are characterized by private monitoring, which means that subjects can merely observe the outcomes of their own matches. In one of those treatments, subjects play the stage game with the costly outside option ($\pi_{OO} = 10$). In the other private monitoring treatment subjects are faced with the costless outside option instead ($\pi_{OO} = 15$).

In the two remaining treatments personal histories are observable for all players. This means that every subject can review the previous decisions made by the opponent she is matched with in the current round of play. These treatments, which combine the observability of personal histories with the option to abstain, test how subjects respond to the possibility of punishing deviations whilst maintaining a favourable reputation.

In summary, this study varies whether personal histories are observable in the economy, as well as the guaranteed pay-off associated with the outside option. Firstly, subjects are either only aware of their own histories, or they can observe a complete decision history for every opponent they meet in the continuation game (No Reputation vs. Reputation). Secondly, participants can either abstain from a one shot interaction by accepting a lower stage game pay-off than π_{DD}, or they can opt out from the stage game and secure a pay-off equal to π_{DD} (Costly Outside Option vs. Costless Outside Option).

Experimental Procedures: A total of 100 people participated in this study. The subjects were undergraduate students at the University of Exeter, who were recruited from all disciplines other than Economics and Psychology. The students signed up online using the ORSEE recruitment system ([22]). All sessions for this study were conducted at Exeter University's own FEELE laboratory. Dividers were in place between individual desks so as to ensure that subjects could not observe one another or communicate with other participants in any way. The participants interacted with each other anonymously through computers. The experimental interface for this study was programmed using z-Tree ([16]).

[5]The data from chapter one was re-used for the framework with the costly outside option.

Experimental instructions were laid out on all desks, and they were also read out aloud[6]. After the instructions had been read by the experimenter, subjects were required to complete a quiz, which tested their understanding of the experiment. All answers were checked by the experimenter and additional questions were answered in private. The correct answers to the quiz were subsequently announced to the room and were therefore common knowledge.

Every experimental session consisted of six cycles. On average, each cycle consisted of 11.48 periods in total. The variance for the underlying distribution of randomly drawn numbers was 3.68. People's earnings were determined solely by their own actions, as well as the actions of their playing partners. Subjects were paid their cumulative earnings for the second and third cycles with and without reputation. The experimental earnings were converted to money using a fixed exchange rate, such that 60 tokens = £1. The payable amount included a £3 show-up fee, which was paid to all participants. People's earnings differed significantly across the treatment configurations. In particular, subjects made on average £3.26 from the treatments with private monitoring. Across the two reputation treatments their earnings differed significantly depending on the outside option. Participants made on average £6.21 when the outside option was costly. By contrast, average earnings were substantially higher for the configuration with the costless outside option, at £8.65. All subjects were paid individually and advised they were not obliged to disclose how much money they had earned.

[6]The experimental instructions are included in Appendix B.

2.3.2 Experimental Hypotheses

I first consider the outside option in the treatments with private monitoring where subjects only observe the outcomes of their own pairings. In this situation choosing not to play can never be individually optimal. This is due to the fact that opting out yields a guaranteed pay-off $\pi_{OO} \leq \pi_{DD}$ in every stage game. There are no reputational concerns in this framework, since behavioural histories cannot be accessed ([45] and [55]).

Hypothesis 10. *None of the subjects should choose to opt out in the private monitoring treatments.*

The outside option is not game theoretically meaningful when subjects merely observe the results of their own interactions. As a result, cooperation levels should not differ significantly between the two private monitoring treatments.

Hypothesis 11. *No significant difference in cooperation levels should be observed across the private monitoring treatments.*

Cooperation can be sustained as an equilibrium in the repeated game even when personal histories are not observable (Proposition 11). However, this constructive equilibrium is intuitively unappealing, as it is highly fragile and dictates the punishment of innocent players. On the other hand, when personal histories are available each agent is assigned a label, which gives an accurate account of her past behaviour. This aspect of observability is useful, as it allows cooperators to identify one another and target deviators efficiently. Hence, a single defection by any subject will not cause a collapse of cooperative behaviour in the economy. Instead, the community will punish the deviator (Propositions 14/15). This can only have the effect of increasing cooperation levels in the group.

Hypothesis 12. *Cooperation levels should be significantly higher in the reputation treatments.*

Whereas the outside option should never be utilized in the absence of personal histories, it is part of the punishment protocol when behavioural histories are observable. Even though choosing to opt out is costly to the punisher, it will prevent her personal record being tarnished. Therefore, she will not subsequently be subject to disciplinary action. As a result, no permanent damage will come to the community as a result of a deviation. This in turn prevents an erosion of the continuation pay-offs for all subjects.

Hypothesis 13. *The outside option should be utilized to punish deviators in the treatments with reputation.*

Since the outside option is used to punish defections, the costliness of administering the punishment depends on the stage game structure. Whenever the guaranteed pay-off is such that $\pi_{OO} < \pi_{DD}$, subjects incur a lower stage game pay-off when choosing not to play. It follows that there is a direct cost associated with preserving a favourable reputation. On the other hand, when the outside option yields a pay-off equal to π_{DD}, players are indifferent between opting out and achieving mutual defection in the prisoner's dilemma. Hence, subjects should be more prepared to opt out when the outside option is costless.

Hypothesis 14. *Subjects should opt out more readily in the reputation treatment with the costless outside option.*

Choosing to abstain means that a deviation will not cause a chain reaction in the community. Instead, once the defector has been punished for two periods, the economy will revert to a state of complete cooperation (Proposition 14/15). Given that not playing is a more attractive option when $\pi_{OO} = \pi_{DD}$, cooperation levels should be higher in the reputation treatment with the costless outside option.

Hypothesis 15. *Across the reputation treatments, cooperation levels should be higher when the outside option is costless.*

Whenever behavioural histories are available, a defection by any subject will result in the deviator being punished for two consecutive periods (Proposition 14). This equilibrium condition is the same in both reputation treatments, regardless of the outside option being costly or costless. It follows therefore that in both reputation treatments subjects should closely monitor their current partner's decisions in the two preceding periods. Both of those decisions should have a significant effect on their action choice.

Hypothesis 16. *The current partner's decisions in the two preceding matches should significantly influence a subject's behaviour in both reputation treatments.*

Finally, it follows from the theoretical analysis that a subject should cooperate in the aftermath of a defection, in order to be admitted back onto the equilibrium path.

Hypothesis 17. *Subjects should choose to cooperate following a defection when personal histories are observable.*

To sum up, cooperation should be maximized in the reputation treatment with the costless outside option. Also, more cooperation should be evident in both reputation treatments relative to the treatments with private monitoring. There should be no difference in cooperation levels between the two private monitoring treatments. Finally, both outside options should only be utilized when subjects have reputational concerns.

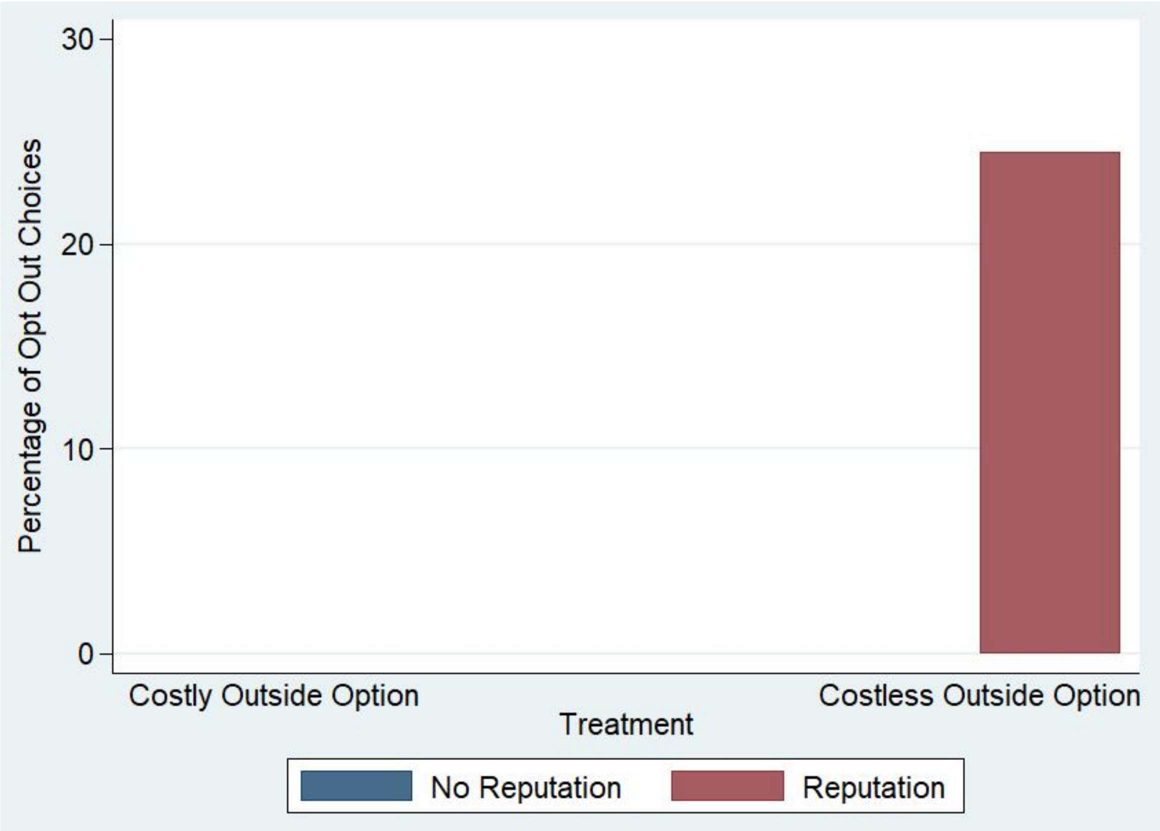

Figure 2.2 Use of the Outside Option by Treatment

2.4 The Results

2.4.1 Non-parametric Analysis

I begin the analysis of the experimental data by presenting results on the overall differences between individual treatments. I aggregate all cycles by treatment and analyse any differences in the individual choices made. To start with, I examine the use of the outside option across different treatment configurations (Figure 2.2).

It is evident that none of the subjects choose to opt out in the two private monitoring treatments. This is likely due to the fact that they consider the outside option as a dominated action choice. In addition, there is no reputational value in opting out, as players observe merely the strategy profiles of their own matches.

Result 13. *The subjects make no use of the outside option in either of the private monitoring treatments.*

However, even when behavioural histories are observable, subjects only utilize the outside option when it is costless to do so. Opting out accounts for 24.18% of all individual choices in the reputation

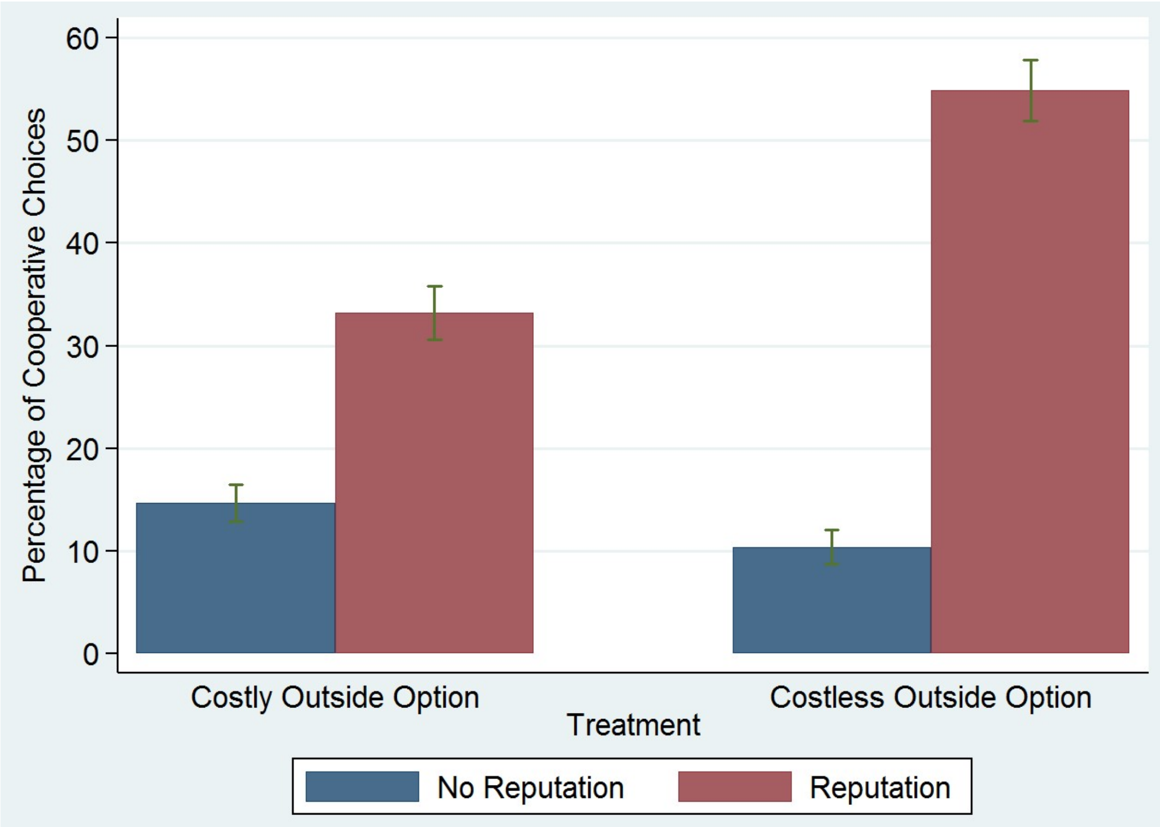

Figure 2.3 Average Cooperation by Treatment (Conditional on Entry)

treatment with the costless outside option. By contrast, none of the subjects choose to opt out when the outside option is costly. A Mann-Whitney test rejects the null hypothesis of equality at the 1% level (z=3.000, p=0.003).

Next, I analyse differences in the individual choice to cooperate across the four treatment specifications (Figure 2.3). I disregard all opt out choices, which I effectively treat as non-decisions. My approach is very much in line with the decision not to play, since the subject effectively chooses to abstain from the current pairing.

Cooperation levels in the two private monitoring treatments are nearly identical. Cooperative choices account for 10.41% of all decisions when the outside option is costless and 14.25% when it is costly (z=1.149, p=0.251, Mann-Whitney test). This is an immediate consequence of both outside options being ignored entirely.

Result 14. *There is no significant difference in cooperation levels across the two private monitoring treatments.*

I now analyse the differences in cooperation between the private monitoring treatments and those with reputation. The observability of personal histories is manipulated in a within-subjects design. Cooperation averages 32.6% in the reputation treatment with the costly outside option, which marks a significant increase in cooperative behaviour over the corresponding treatment with private monitoring (z=2.023, p=0.043, Wilcoxon Sign-Rank test). A similar trend is also evident for the framework with the costless outside option, where the reputation treatment produces a cooperation rate of 54.97%. This is again a significant increase in cooperation over the private monitoring environment (z=2.023, p=0.043, Wilcoxon Sign-Rank test).

Result 15. *Cooperation increases significantly when behavioural histories are observable.*

Finally, I turn my attention to the two reputation treatments, where I compare the impact of the two different outside options. A Mann-Whitney test shows that the difference in cooperation levels between the two reputation treatments is statistically significant at the 5% level. (z=2.193, p=0.028). This finding implies that conditional on entry cooperation is higher when it is costless not to play.

Result 16. *The reputation treatment with the costless outside option achieves a higher rate of cooperation than the corresponding treatment with the costly outside option.*

2.4.2 Treatment Specific Analysis

In order to uncover further differences across individual treatment configurations, I analyse the data for each of the treatments separately. First, I use a random effects probit model to analyse the dynamics of the individual decision to cooperate ([23])[7]. In this regression, the dependent variable is the decision taken by a subject in each of her stage games. The variable takes on the value of one when a subject chooses to cooperate, and is equal to zero otherwise[8]. Again, I exclude all opt out decisions for the reputation treatment with the costless outside option.

Table 2.4 Random Effects Probit Regression on the Choice to Cooperate

	$\pi_{OO} = 15$ (Cooperate = 1)	Reputation - $\pi_{OO} = 15$ (Cooperate = 1)	$\pi_{OO} = 10$ (Cooperate = 1)	Reputation - $\pi_{OO} = 10$ (Cooperate = 1)
2nd Cycle	-0.81^{*}	0.47	0.37	0.20
	(0.42)	(0.34)	(0.77)	(0.27)
3rd Cycle	-0.98^{**}	0.23	-0.63^{**}	0.65^{**}
	(0.32)	(0.23)	(0.30)	(0.32)
Period	-0.19^{**}	-0.03^{***}	-0.07^{*}	-0.13^{**}
	(0.09)	(0.01)	(0.04)	(0.05)
2nd Cycle x Period	0.14	-0.16^{***}	-0.02	0.07^{**}
	(0.11)	(0.06)	(0.04)	(0.03)
3rd Cycle x Period	0.11	0.07^{*}	-0.05	-0.01
	(0.07)	(0.04)	(0.06)	(0.09)
Length Previous Cycle	0.05	-0.02	-0.38	-0.03
	(0.07)	(0.04)	(0.44)	(0.06)
Restart	0.55^{***}	0.17	0.44	0.41^{**}
	(0.20)	(0.13)	(0.27)	(0.18)
Intercept	-1.11	0.67^{***}	2.70	-0.24
	(0.84)	(0.24)	(4.58)	(0.35)
lnsig2u				
_cons	-0.61	-0.59	0.29	0.41
	(0.31)	(0.22)	(0.23)	(0.29)
N	1,270	1,121	1,379	1,235

Robust Standard Errors in Parentheses $^{*}\ p < 0.1$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

[7] A Hausman test on the data set fails to reject the null hypothesis.
[8] A logit regression model produces nearly identical results.

Next, I run an additional regression for the reputation treatment with the costless outside option. I do this to incorporate the decision not to play in my econometric analysis. Since the decision variable can now take on three possible values (0 = Defect, 1 = Cooperate, 2 = Opt Out), I use a multinomial probit model[9].

Table 2.5 Multinomial Probit Regression for the Reputation - π_{00} = 15 Treatment

(Defect = Base Outcome)	Cooperate	Opt Out
2nd Cycle	0.36	0.47
	(0.33)	(0.57)
3rd Cycle	0.21	0.68^{**}
	(0.33)	(0.28)
Period	-0.04^{**}	0.09^{**}
	(0.02)	(0.04)
2nd Cycle x Period	-0.15^{**}	-0.09
	(0.06)	(0.08)
3rd Cycle x Period	0.08	-0.09
	(0.06)	(0.54)
Length Previous Cycle	-0.02	0.07^{***}
	(0.05)	(0.02)
Restart	-0.20	-2.38^{***}
	(0.13)	(0.54)
Intercept	0.77^{**}	0.22
	(0.37)	(0.27)
N	1,480	1,480

Robust Standard Errors in Parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

It is evident from table 2.4 that the dynamics of individual decision making vary significantly across the different treatment configurations. In the treatments with private monitoring subjects completely disregard the outside option and instead embrace the benefits of defection. On the other hand, it appears that cooperation is far more stable in those treatments, which feature personal histories. This behavioural evidence is in line with the theoretical result that reputational incentives promote

[9]The random effects assumption is not applicable to this regression model.

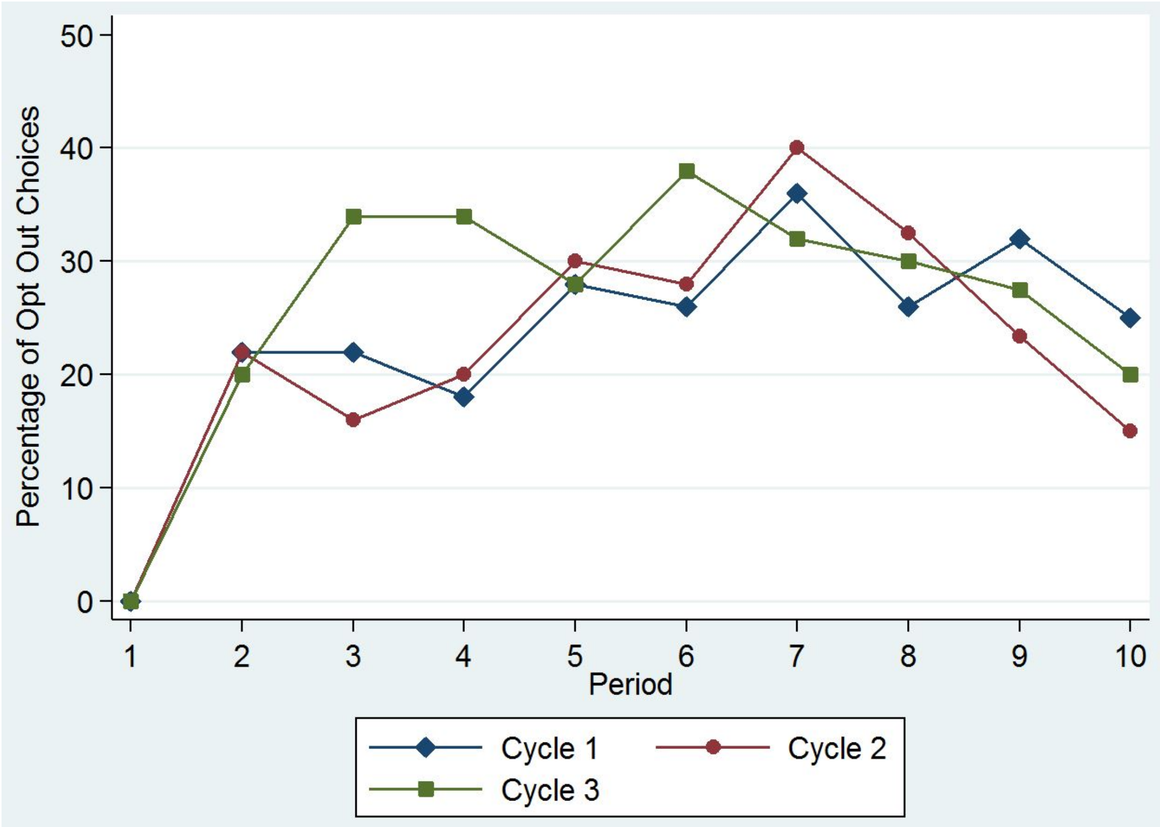

Figure 2.4 Use of the Costless Outside Option

cooperation within a community ([30]).

It is interesting that even in the presence of personal histories subjects make no use of the costly outside option. This clearly suggests that they are not willing to accept a reduced stage game pay-off in order to punish deviations whilst retaining a positive reputation. Instead, they defect in order to attain the equilibrium pay-off in the current round of play. Thus, subjects deliberately choose the implicit cost of having a defection added to their behavioural history over the direct cost of incurring a reduced pay-off ([14]).

Result 17. *Subjects refuse to bear a direct cost in their current match to enforce cooperation.*

By contrast, when the outside option is costless, and therefore equivalent to the one shot Nash equilibrium, subjects embrace the benefits of having the freedom not to play. The usage of the costless outside option is largely in line with the theoretical predictions, as there is no substantial variation across consecutive cycles in terms of people opting out (Figure 2.4). Instead, subjects never opt out in the first round of a cycle, whereas the outside option is readily used in the subsequent periods. It is therefore apparent that the players comprehend the importance of having the option to abstain.

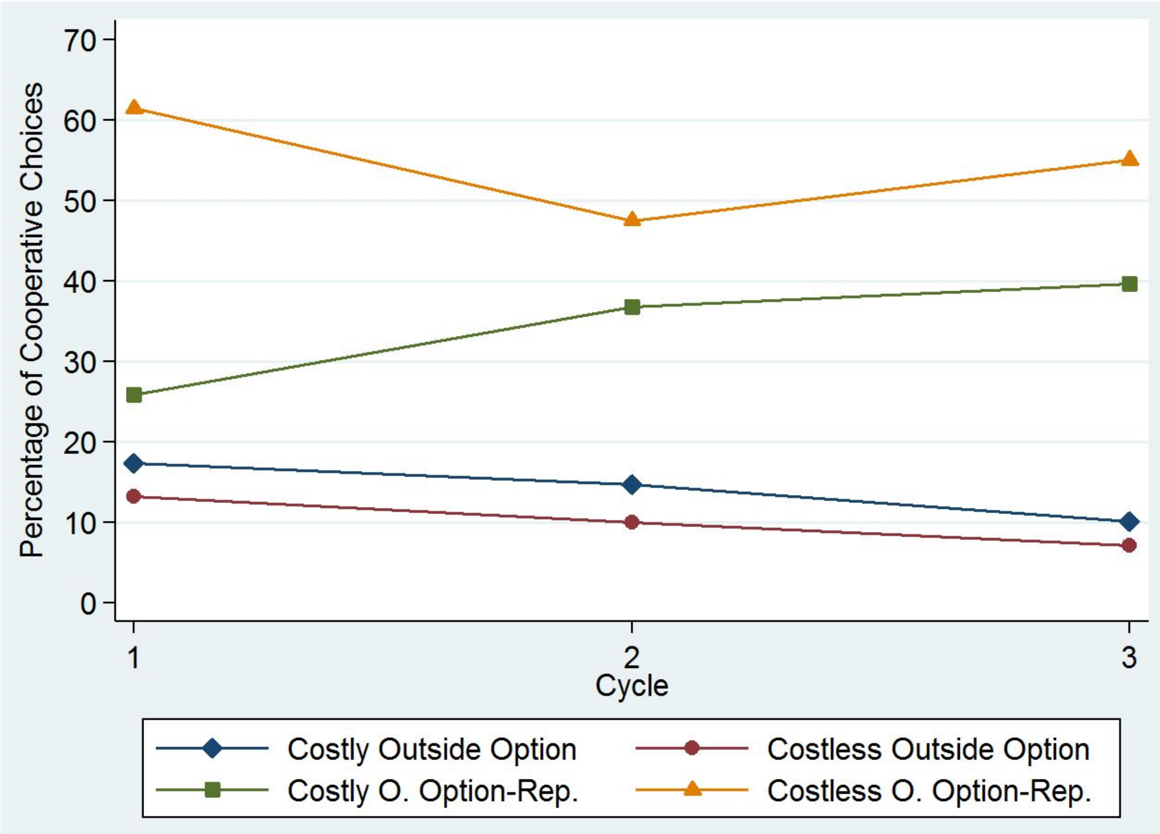

Figure 2.5 Average Cooperation across Cycles (Conditonal on Entry)

Result 18. *None of the participants opt out at the start of a cycle, whereas the outside option is used substantially in the later rounds of the game.*

The costless outside option leads to significantly higher cooperation levels relative to the reputation treatment with the costly outside option (Figure 2.5). Therefore, it can be concluded that the costless outside option is effectively used by the subjects to penalize deviations without committing a defection themselves. Thus, the spread of opportunism within the economy is contained, and a single defection does not trigger a chain reaction. The costless outside option serves to enforce cooperation, rather than punish deviators.

Even for the reputation treatment with the costly outside option cooperation rates increase across consecutive cycles. However, this cannot be attributed to the outside option itself, which the subjects disregard completely. Instead, it seems that over the course of the three continuation games, the subjects begin to grasp the interdependence of individual encounters. Therefore, in the later cycles people attempt to embrace the benefits of sustained cooperation in the community ([10]). However, as subjects are not prepared to use the costly outside option, a single deviation inevitably triggers a wave of defections, and therefore a collapse of sociable behaviour.

Result 19. *People's refusal to use the costly outside option means that the spread of deviations cannot be stopped.*

Finally, it is apparent that in the private monitoring treatments cooperation declines across consecutive cycles. This is likely due to the fact that subjects do not consider themselves to be part of a group when they can only observe the results of their own pairings, and thus, have no accurate information about the state of the economy. Instead, they consider the repeated game as a dynamic sequence of one shot games.

2.4.3 Individual Reputation Effects

The Costly Outside Option

I now further analyse the two reputation treatments. In particular, I study the effect of behavioural histories on an individual's decision to either cooperate, defect or abstain. I seek to quantify the effects associated with reputational concerns and the extent to which the outside options affect people's behaviour. I start by analysing the costly opt out treatment.

The variable *Partner's Reputation* is a ratio of the number of cooperative choices relative to the total number of decisions made. The variable therefore takes on values between zero and one, depending on the relative number of cooperative choices in a person's behavioural history. I analyse the effect of the two most recent decisions in a separate regression model, so as to avoid multicollinearity. *Partner Lag 1* measures the effect of the current partner's decision in the preceding match, whereas *Partner Lag 2* corresponds to her decision two periods ago.

I include an explanatory variable in the model to control for a subject's own decision in the preceding match. This is because a deviation in the previous pairing is likely to influence her decision behaviour in the present round of play. It follows from the theoretical discussion that a subject should repent following a deviation, so as to be admitted back into the community. I also capture a player's decision in the first period of the current cycle. This approach seeks to address the underlying issue of endogeneity, which stems from the regressor *Having Defected in Previous Match* being correlated with the first lag of the dependent variable.

It is clear from table 2.6 that even in the presence of a costly outside option there is a significant relationship between the current partner's behavioural history and a person's decision to cooperate. However, the magnitude of this effect does not significantly vary over time. It would seem intuitive that the partner's behavioural history should be a more significant factor for decisions made late in the cycle. As subjects choose an action in each round, there is a richer set of choices to be observed in the later stages.

Result 20. *Personal histories have a significant effect on the individual decision to cooperate. However, the magnitude of this relationship does not evolve over time.*

The effect of personal histories does not evolve over the course of consecutive periods. This clearly indicates that subjects only take into consideration the most recent decisions made by the current partner. This result is very much in line with theoretical predictions. In order to determine the precise effect of being able to observe people's past behaviour, it is necessary to regress the decision to cooperate on the partner's most recent choices. It follows from hypothesis 16 that the current partner's two most recent decisions should have a significant impact on an individual's action choice.

Table 2.6 Random Effects Probit Regression on the Choice to Cooperate

	Coefficient (Cooperate = 1)	dy/dx (Cooperate = 1)	Coefficient (Cooperate = 1)	dy/dx (Cooperate = 1)
2nd Cycle	0.02	0.06**	0.26	0.08**
	(0.19)	(0.03)	(0.30)	(0.04)
3rd Cycle	0.49	0.04*	0.65	0.05*
	(0.54)	(0.02)	(0.62)	(0.03)
Period	−0.12*	−0.03***	−0.02	−0.01***
	(0.06)	(0.01)	(0.02)	(0.00)
2nd Cycle x Period	0.05**		0.02	
	(0.02)		(0.05)	
3rd Cycle x Period	−0.06		−0.07	
	(0.10)		(0.10)	
Length Previous Cycle	−0.06	−0.01*	−0.06	−0.01*
	(0.04)	(0.01)	(0.05)	(0.00)
Partner's Reputation	0.80***	0.17***		
	(0.12)	(0.03)		
Partner's Reputation x Period	0.10*	0.02		
	(0.06)	(0.01)		
Having Defected in Previous Match	−0.46***	−0.10***	−0.69**	−0.15**
	(0.17)	(0.04)	(0.25)	(0.05)
Cooperation in Period One	0.59***	0.13***	0.69**	0.15**
	(0.18)	(0.04)	(0.25)	(0.04)
Partner Lag 1			0.78***	0.16***
			(0.24)	(0.04)
Partner Lag 2			0.48	0.10
			(0.34)	(0.07)
Intercept	−0.48		−0.86**	
	(0.42)		(0.28)	
lnsig2u				
_cons	−0.25		−0.54	
	(0.21)		(0.41)	
N	1,051	1,051	841	841

Robust Standard Errors in Parentheses   * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

It is apparent from table 2.6 that subjects have a very short memory when considering their current partner's behavioural record. Only the opponent's choice in her previous match has a significant effect on the individual decision to cooperate. This observation is clearly at odds with the theoretical result that a defection by any player should be punished for the duration of two consecutive periods.

Overall, an analysis of the marginal effects shows that the benefit from maintaining a positive reputation is much smaller than is commonly believed. Cooperation by the current partner in the previous round merely increases the probability of cooperation by 16%.

Result 21. *Reputational concerns play a minor role when it is costly to abstain from an interaction.*

The Costless Outside Option

I now analyse the effect of personal histories in the presence of a costless outside option. I again use a multinomial probit model in order to incorporate all opt out choices[10]. A multinomial logit model produces nearly identical marginal effect estimates.

Table 2.7 Multinomial Probit Regression on the Individual Reputation Effects

(Defect = Base Outcome)	Cooperate	Predict Outcome Cooperate	Opt Out	Predict Outcome Opt Out
2nd Cycle	0.07	-0.12^{**}	0.54	0.03
	(0.43)	(0.04)	(0.89)	(0.07)
3rd Cycle	-0.45	-0.13^{***}	0.31	0.10^{**}
	(0.42)	(0.02)	(0.39)	(0.04)
Period	-0.05	-0.02^{*}	0.03	0.01
	(0.04)	(0.01)	(0.10)	(0.02)
2nd Cycle x Period	-0.11^{*}		-0.11	
	(0.05)		(0.11)	
3rd Cycle x Period	0.01		-0.01	
	(0.07)		(0.07)	
Length Previous Cycle	-0.03	0.00	-0.09^{*}	-0.02^{**}
	(0.04)	(0.08)	(0.05)	(0.01)
Partner's Reputation	1.62^{***}	0.49^{***}	-0.47^{*}	-0.27^{**}
	(0.59)	(0.12)	(0.25)	(0.11)
Partner's Reputation x Period	-0.07	-0.01	-0.01	0.01
	(0.07)	(0.01)	(0.10)	(0.02)
Having Defected in Previous Match	-1.03^{***}	-0.18^{**}	-0.60^{*}	-0.02
	(0.23)	(0.02)	(0.36)	(0.06)
Cooperation in Period One	1.12^{***}	0.21^{***}	0.45^{**}	-0.02
	(0.11)	(0.03)	(0.18)	(0.04)
Intercept	-0.15		-0.57^{**}	
	(0.48)		(0.70)	
N	1,156	1,156	1,156	1,156

Robust Standard Errors in Parentheses — * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

[10] As before, the random effects assumption does not apply to this model.

For the regression model in table 2.7, the variable *Partner's Reputation* indicates the ratio of cooperative and opt out choices relative to the total number of decisions made. This approach implies that the decision not to play is perceived in a positive way and therefore boosts an individual's reputation. However, the assumption seems justified based on the previous analysis of the data, which suggests that the outside option is used to enforce cooperative behaviour.

It is apparent from the data that a costless outside option significantly changes the impact of behavioural histories on people's decision making. The partner's reputation has a marginal effect of 49% on a person's decision to cooperate, which is significantly greater than in the treatment with the costly outside option.

Result 22. *The marginal effect of behavioural histories on individual decisions is considerably greater when the option not to play is costless.*

In addition, the regression clearly shows that the partner's reputation has a significantly negative effect on the decision to abstain from a stage game pairing. This finding clearly supports the hypothesis that subjects consider the outside option as a means of penalizing deviations without damaging their own reputation.

Once again, the impact of personal histories does not fluctuate over time. This suggests that subjects only take into account the more recent choices made by their current partner. An analysis of the two most recent decisions illustrates that people's responses to personal histories are very much in line with the theoretical predictions (Table 2.8). In particular, when making a decision subjects monitor the two most recent choices made by their current playing partner[11]. It is therefore apparent that players review a partner's behavioural history more carefully when it is costless to abstain from the interaction.

Table 2.8 shows that a previous decision to cooperate or abstain by the current partner increases the probability of cooperation. This confirms that the decision not to play is generally viewed favourably by others and therefore benefits a person's reputation. However, a decision to abstain has a significant effect at lag one, whereas it is only marginally significant at lag two. By contrast, cooperative choices have a significant effect for both preceding matches. Cooperation by the partner at lag one increases the probability of a cooperative decision by 26%. Similarly, cooperation by the partner at lag two makes cooperative behaviour 19% more likely. This shows that people condition their decision making on the partner's behaviour in the two previous matches when it is costless to abstain from the pairing.

Result 23. *Behavioural histories are significantly more important in the presence of a costless outside option.*

[11] *Partner Lag 3* has no significant effect on an individual's decision making.

Table 2.8 Multinomial Probit Regression on the Individual Reputation Effects

		Predict Outcome		Predict Outcome
(Defect = Base Outcome)	Cooperate	Cooperate	Opt Out	Opt Out
2nd Cycle	−0.07	−0.10**	0.52	0.00
	(0.50)	(0.05)	(0.73)	(0.05)
3rd Cycle	−0.15	−0.13***	0.79***	0.10**
	(0.54)	(0.03)	(0.28)	(0.04)
Period	−0.11***	−0.03***	0.02	0.01
	(0.02)	(0.01)	(0.06)	(0.01)
2nd Cycle x Period	−0.08*		−0.12	
	(0.04)		(0.09)	
3rd Cycle x Period	−0.05		−0.09**	
	(0.07)		(0.04)	
Length Previous Cycle	−0.01	0.01	−0.08	−0.02**
	(0.05)	(0.01)	(0.05)	(0.01)
Having Defected in Previous Match	−1.10***	−0.18***	−0.68*	−0.03
	(0.19)	(0.05)	(0.37)	(0.07)
Cooperation in Period One	1.15***	0.21***	0.48**	−0.02
	(0.08)	(0.02)	(0.20)	(0.04)
Partner Lag 1				
Cooperate	0.86**	0.26***	−0.14	−0.13*
	(0.37)	(0.09)	(0.20)	(0.07)
Opt Out	0.52**	0.13**	−0.20	−0.09
	(0.25)	(0.06)	(0.19)	(0.07)
Partner Lag 2				
Cooperate	0.53**	0.19***	−0.39***	−0.14***
	(0.23)	(0.05)	(0.08)	(0.02)
Opt Out	0.23	0.06*	0.02	-0.02
	(0.19)	(0.03)	(0.14)	(0.02)
Intercept	0.05		0.49	
	(0.30)		(0.54)	
N	998	998	998	998

Robust Standard Errors in Parentheses — * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

2.5 Concluding Remarks

This chapter seeks to explore the significance of outside options in economies made up of strangers. These economies are indefinite in duration and the subjects are randomly and anonymously rematched in every period. The players have no means of communicating and they are matched with a sequence of unknown opponents. Attaining cooperation in such communities is inherently difficult, since a single deviation in the economy will result in a total collapse of sociable behaviour ([29] and [6]). Making first order histories observable for all players has been shown to be somewhat useful in maintaining cooperation ([53]). However, it still remains difficult to achieve high social welfare in such a market, since there is no mechanism to differentiate deviation from punishment. The option to abstain from any given stage game can help address this problem, because it allows for defections to be contained and cooperation to be enforced ([25]).

In my study, I test experimentally whether subjects are prepared to use the option not to play in order to enforce cooperation rather than penalize defections by others directly. I employ a fixed probability of continuation to simulate an indefinite interaction ([4]). The subjects are repeatedly rematched to play the prisoner's dilemma. However, I modify the conventional form of the game to include an option not to play. I vary whether the outside option either yields a pay-off smaller than or equal to the Nash equilibrium in the prisoner's dilemma. Thus, all interactions between the subjects are entirely optional and people are able to unilaterally abstain from any given pairing. In order to further analyse the implications of being able not to play, I also include a reputation parameter in my experiment. Specifically, I vary whether subjects can either only observe the outcomes for the stage games they play, or they can review the behavioural history of every person they meet. It is apparent that there can be no strategic value in opting out from an interaction when people only observe the results of the pairings they are directly involved in ([55]). However, when personal histories are observable the outside option can be used to penalize deviations without committing a defection.

The results of the experiment clearly indicate that making personal histories observable has a significant effect on social welfare. In particular, both reputation treatments achieve significantly higher levels of cooperation than the treatments with private monitoring ([5]). However, it is evident from the data that cooperation is maximized in the economy, which combines personal histories with the costless outside option. Subjects are clearly reluctant to abstain from an interaction when the resulting reward is below the Nash equilibrium pay-off. On the other hand, people are keen to embrace the option not to play when the pay-off is equivalent to the Nash equilibrium. This is reflected in the fact that the costly outside option is never played, whereas the costless outside option is used substantially. By contrast, there is no significant difference between the two treatments with private monitoring. There both the costly and the costless outside option are ignored completely. This finding is very much in line with the result that subjects, in the absence of personal histories, either defect from the beginning or quickly learn to defect in the continuation game. It is evident that private monitoring means players see the interaction as a mere series of one shot games.

I further develop this result by closely examining the individual reputation effects for those treatments that feature personal histories. I find for the treatment with the costly outside option that the importance of behavioural histories is rather limited. It is apparent from the data that subjects only take into consideration their current partner's decision in the preceding match. This indicates people have a very short memory when it comes to penalizing opportunism. The result is complemented by the finding that cooperation by the current partner in her preceding match only increases the probability of cooperation by 16%. In the treatment with the costless outside option, the effect of being able to observe personal histories is more complex. Subjects review the history of a playing partner much more thoroughly when it is costless to abstain from the pairing. In particular, it seems that people take into consideration the two most recent choices under this configuration. This constitutes a significant change from the costly outside option environment. Also, it is apparent that a previous decision to abstain by the current partner is perceived in a positive way, and therefore makes cooperation more likely. This finding implies that subjects fully understand the significance and the value of being able to walk away.

All in all, this experiment confirms the theoretical result that giving subjects the option not to interact with others can increase social welfare and produce considerable gains for cooperators [24]. However, it is evident that the beneficial effects associated with outside options are maximized when people have reputational incentives. Also, it is apparent that subjects respond far more positively to the option not to play when it is costless for them to do so. This suggests that individuals are reluctant to make sacrifices at the current time in order to retain a positive reputation moving forward. Therefore, reputation clearly plays a less crucial role when it is costly to avoid interactions with others.

Chapter 3

Reputation in Patient Hospital Selection: Empirical Evidence from Australia

3.1 Introduction

In recent years extensive reforms to the healthcare sector in many countries have largely enabled people to choose their preferred hospital. Such changes to the mechanism, which allocates patients to healthcare providers, give rise to strategic considerations and more direct competition among medical institutions ([40]). As a result, reputation effects increasingly play an important role in the individual decision to attend a particular hospital. The existing literature focuses largely on how patients perceive and rate the quality of medical care. However, little research has thus far been dedicated to the importance of reputation effects in choosing a hospital from a wide range of alternatives.

Hospital selection has traditionally been studied using survey data. [21] identifies the location of a clinic as a crucial factor in choosing a medical institution. The study finds that patients in rural areas typically consider their local institution to be inferior to metropolitan alternatives. This preconceived bias is evident in both users and non-users of rural hospitals. [54] finds that the decision to avoid the nearest hospital also depends significantly on patient specific characteristics, such as age, employment status and treatment specific requirements.

[15] investigates the determinants of patient satisfaction. Survey data shows that patient evaluations depend most critically on the overall outcome of a hospital stay. The paper therefore concludes that a clinic's reputation is heavily correlated with the perceived quality of medical care. [7] links patient satisfaction with the decision to recommend a clinic. The study finds that patient ratings hinge primarily on hospital staff. Good scores for the interpersonal skills of doctors and nurses are heavily correlated with the decision to recommend an institution. Associations are evident between perceived

staff competence and a hospital's reputation.

[48] studies how patients decide which hospital to attend for major surgical procedures. Most respondents state that they would rely on advice from their GP and opinions of others rather than performance data when selecting a medical institution. Other significant decision factors include prior experience with a hospital and recommendations from family and friends. Hence, the study concludes that directing patients towards high quality institutions by publishing performance data is likely to have limited success.

Empirical evidence on hospital selection is both scant and inconclusive. [27] conducts an analysis into the effect of hospital performance reports. The study shows that the publication of a performance report generally leads to quality improvements. However, the market share for individual hospitals appears unaffected by these reviews. There is no shift away from low quality institutions and there is equally no evidence of a surge towards clinics performing above average. By contrast, [44] finds that for top hospitals in the US an improvement in the rankings has a positive effect on the volume of non-emergency patients. In particular, an improvement by merely one position leads to an increase in patient numbers by approximately one per cent. The study therefore concludes that there is a strong relationship between a clinic's position in the rankings and the decision to attend that institution.

Choice behaviour in the healthcare sector has become the subject of economic research in modern times. However, a review of the existing literature shows that little attempt has to date been made to quantify the importance of hospital reputation. The research so far has mainly identified observable factors, which impact on hospital choice. Yet, the question of how much reputation effects influence the decision to attend a particular clinic remains largely unanswered. This gap in the literature is due to very limited survey data on patient satisfaction, as well as conflicting empirical evidence. This study seeks to complement the existing research by proposing new estimation methods to gauge the magnitude of reputation effects using hospital administrative data.

This chapter analyses hospital reputation in the context of birth episodes. Expectant mothers generally dedicate substantial effort to selecting the hospital best suited to their needs and preferences. This market therefore lends itself to the modelling and estimation of reputation effects. I focus my analysis on childbirths in the Melbourne metropolitan area. I consider women with private health insurance plans, as the Australian healthcare system permits these patients to freely select their preferred hospital from a wide range of options. Previous research suggests that distance is an important influence on hospital choice ([3]). Therefore, I treat the distance patients are willing to travel to a hospital as a valid indicator of its reputation. By controlling for patient specific complexities and observable clinic features, I numerically capture the importance of reputational considerations in hospital selection.

The Australian healthcare sector is well suited to my analysis of hospital reputation. According to official data, 45.8% of Australians hold private health insurance policies in 2011 ([35]). The high percentage of private patients in that country is due to existing legislation. Whereas Australia has a universal healthcare system named Medicare, penalties are imposed for high income earners, who do not hold private health insurance. For example, the Medicare Levy Surcharge (MLS) is payable on incomes above $90,000 for singles and $180,000 for families ([38]). Incentives and benefits are available to those, who take out private health insurance, such as rebates on insurance premiums. It follows that for many Australians it is financially sensible to purchase policies, in order not to pay the MLS. Many others opt for private health insurance to avoid long waiting times for elective procedures and to freely choose their healthcare provider, specialist doctor or hospital.

My analysis consists of two parts. Firstly, I seek to explain the distance a patient travels to the hospital of her choosing. As such, I treat the distance travelled as the variable of interest. My approach stems from the notion that patients are likely to travel longer distances to attend a more reputable hospital. This allows me to quantify how strongly people value an institution's reputation and how such considerations affect their decision making. I compute hospital fixed effects and reveal the institution specific differences using a random effects model. In the second part of the study, I control for the distance a patient is likely to travel ([20]). Hence, I consider all hospitals within a specified travel distance of a person's residence. I then estimate the probabilities of individual hospitals being selected from the set of options. To confirm the robustness of my results, I vary the maximum distance a patient can travel.

The results clearly show that there are substantial hospital dependent differences in the willingness to travel. These differences are robust to a wide range of patient characteristics. My analysis reveals that the hospital fixed effects consist mainly of clinical performance indicators and other demographic factors. These account for approximately half of the variation in the distance patients travel. In the second part of my study, I find that even when controlling for the number of accessible hospitals, there are substantial differences in the probability of individual institutions being chosen. These differences between clinics do not depend on the willingness to travel and are robust to changes in the clustering of the standard errors. Therefore, my findings provide a reliable measure of reputation effects in the hospital selection of pregnant women.

The remainder of the chapter is structured as follows: The dataset used in this study is introduced in section two. Section three outlines the econometric model, with the two strands of my analysis being explained in detail. In section four, I present the results of my research, including an overview of regression estimates and a breakdown of hospital specific differences. Section five summarizes my key findings and concludes the chapter.

3.2 The Data

I use the Victorian Admitted Episodes Dataset (VAED) from the year 2011/12. The VAED is an administrative dataset, which contains detailed information on all treatment episodes across private and public hospitals in the state of Victoria. The dataset holds extensive information on medical care episodes, treatment outcomes and other patient specific details. These include gender, age, marital status, statistical local area (SLA), ethnicity, diagnoses, treatments and surgical procedures. In my data, 41.33% of those treated hold private health insurance, which is representative of the Australian population.

I analyse the hospital choices of pregnant women with private health insurance, who reside in the Melbourne metropolitan area. My study examines the decision making of private patients, as they are completely unrestricted in their selection of a hospital. By contrast, patients relying on Australia's universal healthcare system do not enjoy the same freedom when choosing a medical institution. My dataset consists of 11,967 inpatient childbirth episodes across 18 Melbourne hospitals. Out of the 18 clinics in my sample, ten are public hospitals and eight are private. There are in total 26 medical centres in Melbourne with obstetric care facilities. However, I exclude eight of those, as they are not visited by any of the women in my sample. I complement my dataset with information on the demographics of the hospital catchment areas and common indicators of clinical performance. To estimate the distance a patient travels, I find the center of her SLA and measure the travel distance to her chosen hospital. For women, who attend a medical centre within their own SLA, I fix the travel distance at 1 km. This is a good estimate given the average size of an SLA. I cannot obtain the exact distance travelled, since addresses are not included in the VAED. This is due to privacy concerns.

In the first strand of the analysis I treat the distance travelled as the dependent variable. I estimate a linear model featuring hospital dummies as regressors. Thus, I obtain institution specific differences in the willingness to travel, whilst controlling for other factors. I subsequently examine the patient specific characteristics and hospital performance criteria, which impact on the distance travelled. Hence, I identify the clinic specific factors, which cause fluctuations in the willingness to travel.

In the second part of my study I compute the probability of individual hospitals being chosen. To facilitate this analysis, I define a maximum distance the patient is willing to travel. Thus, I determine the set of hospitals, which the individual chooses from. The outcome variable takes on the value of one for the clinic chosen and is equal to zero for all other alternatives. To verify the robustness of my results, I vary the maximum travel distance, which changes the composition of the choice set. To achieve sufficient variation in the available options, I fix the travel distance at four different levels. I choose the 50th, 90th, 95th and 99th percentiles of the distance variable in the dataset.

The distribution of distances travelled, illustrates great variation in the hospital selection of pregnant women. Patients travel on average 20.55 km to reach their preferred clinic. The standard deviation shows that the majority of women in my sample travel between 5.61 km and 35.49 km from their home to the hospital. This finding is consistent with the relatively high density of medical institutions in the Melbourne metropolitan area. The average patient in my sample has access to 11.27 hospitals within 41 km of her home. Table 3.1 gives a detailed summary of the distances travelled by the women in my sample.

Table 3.1 Distance Travelled (km)

Mean	**Std. Dev.**	**Variance**	**p10**	**p25**	**p50**	**p75**	**p90**	**p95**	**p99**	**Min**	**Max**
20.55	14.94	223.17	5.31	8.53	16.22	28.38	40.96	49.81	63.27	1	123.77

The explanatory variables in my study are mostly patient specific characteristics, as well as hospital performance indicators. The variable *w18wies* represents the continuous cost weight, which determines the reimbursement a hospital receives for treating a patient. A higher cost weight indicates a larger payment resulting from a more complex care episode. The case specific cost weight is calculated by comparing the medical needs of individual patients to those of other people with similar conditions ([57]). In my dataset the variable does not indicate the treatment cost to the patient herself. This is because the insurance provider bears the full cost of the care episode. However, the cost weight serves as a measure for the relative complexity of a childbirth.

In this study of hospital choices by pregnant women, it is essential to account for circumstantial differences between consecutive birth episodes. For instance, a woman giving birth for the first time may apply different criteria in her hospital selection than for any subsequent childbirth. In addition, experiences from the first birth will likely influence hospital choices for future birth episodes. To take such differences into consideration, I control for whether or not a woman is giving birth for the first time. I obtain the relevant information by using VAED datasets from previous years. I review childbirth records going as far back as 2001/02 and match corresponding obstetric care episodes through the unique Patient-ID [1].

In 2011 Melbourne has a population of approximately 3.85 million people, which accounts for 19% of Australia's national population. The metropolitan area measures 9,990 km^2 in size and is home to approximately 70% of people living in the state of Victoria. Therefore, Melbourne is the regional capital of south-eastern Australia. The city itself exhibits great ethnic diversity with 38% of its inhabitants born abroad ([39]). To reflect these factors in my analysis, I control for the birthplace of a patient by indicating whether or not she was born in Australia. To further account for diverse ethnic backgrounds, which may influence hospital selection, I also consider the indigenous status of some

[1] Ten years is a suitable time frame given the average age of women in my sample.

patients in my sample. Finally, I capture differences in the neighbourhood or the social environment of a patient by dividing the metropolitan area into three regions: southern, eastern and north-western. Implementing these controls is straightforward, as the necessary information is included in the dataset.

Hospital performance indicators are easily accessible in Australia from the Institute of Health and Welfare ([34]). On their website prospective patients can actively compare the recent performance of individual medical centres. Some existing research indicates that publicly available performance indicators influence the decision of which hospital to attend ([44]). As a result, I consider these factors in my analysis. The performance indicators used in my study include the number of childbirths at a hospital in the previous year, as well as the reported number of *Staphylococcus Aureus* bloodstream infections. Both of these statistics are given on the website and are therefore common knowledge. I also control for the number of stillborn babies from the previous year, although this information is not available on the internet. I obtain the number of stillbirths from the VAED for 2010/11.

Apart from common clinical performance measures, I also control for the demographics of the hospital catchment area. It is crucial to account for socio-economic factors, as they directly affect the number of people likely to hold private health insurance. Also, I must consider the population of fertile women, as it inevitably influences the number of childbirths at the local clinic. To control for the demographics of the various hospital catchment areas, I use official census data from the Australian Bureau of Statistics ([1]). The data used comes from the 2011 census, which conveniently captures the same year as my dataset. I include variables to control for the number of women between the ages of 17 and 49 in each hospital catchment area, as well as household income levels.

A detailed overview of all control variables used in this study, including descriptions, means and standard deviations, is given in table 3.2.

Table 3.2 Summary Statistics of all Control Variables

Variable	Description	Mean	Std. Dev.
Age	Patient's Age in Years	33.53	4.23
w18wies	Cost Weight for Reimbursement	1.09	0.66
Born Australia	=1 if Born in Australia	0.74	0.44
Indigenous Status	=1 if Patient has Indigenous Status	0.01	0.08
Length of Stay	Length of Stay in Days	4.61	2.95
Married	=1 if Married/De Facto	0.98	0.16
Eastern Region	=1 if Resides in Eastern Region	0.21	0.41
North Western Region	=1 if Resides in North Western Region	0.37	0.48
Southern Region	=1 if Resides in Southern Region	0.41	0.49
First Birth	=1 if First Birth	0.49	0.46
C-Section	=1 if C-Section	0.40	0.49
Number of Hospitals within 41 km	Hospitals within 41 km	11.27	2.53
Number of Childbirths 2010/11	Number of Births in 2010/11	2,967.42	1,962.68
Bloodstream Infection Rate 2010/11	Infection Rate in 2010/11	0.51	0.52
Number of Stillbirths 2010/11	Number of Stillbirths in 2010/11	0.96	0.20
Percentage of Women in Hospital Area	Percentage of Women (17-49)	27.45	3.47
Number of Women in Hospital Area	Number of Women (17-49)	34,195.42	10,789.76
Number of Households Q1 in Hospital Area	Income $0 - $624 per week	9,419.70	3,048.18
Number of Households Q2 in Hospital Area	Income $625 - $1,213 per week	9,640.28	4,011.40
Number of Households Q3 in Hospital Area	Income $ 1,214 - $2,148 per week	10,723.46	4,282.50
Number of Households Q4 in Hospital Area	Income above $2,148 per week	13,465.18	2,189.34
Percentage of Q3 and Q4 in Hospital Area	Households in Q3 and Q4	56.94	5.10
Distance to Nearest Hospital	Distance to Nearest Hospital (km)	7.57	5.86
Excess Travel	= Distance Travelled/Distance to Nearest Hospital	8.71	11.66
I51	=1 if Hospital Chosen	0.01	0.06
I67	=1 if Hospital Chosen	0.02	0.13
I76	=1 if Hospital Chosen	0.01	0.08
J00	=1 if Hospital Chosen	0.01	0.03
J49	=1 if Hospital Chosen	0.01	0.04
J55	=1 if Hospital Chosen	0.01	0.02
J62	=1 if Hospital Chosen	0.01	0.02
J70	=1 if Hospital Chosen	0.01	0.05
J73	=1 if Hospital Chosen	0.07	0.08
J81	=1 if Hospital Chosen	0.27	0.44
J95	=1 if Hospital Chosen	0.15	0.35
M66	=1 if Hospital Chosen	0.01	0.07
O23	=1 if Hospital Chosen	0.07	0.26
X04	=1 if Hospital Chosen	0.07	0.25
X06	=1 if Hospital Chosen	0.23	0.42
X13	=1 if Hospital Chosen	0.05	0.22
X73	=1 if Hospital Chosen	0.06	0.23
Z42	=1 if Hospital Chosen	0.06	0.24

3.3 The Model

I measure the importance of reputation effects through two separate strands of analysis. In part one, I study the distance a patient travels to her chosen hospital. As the women in my sample have the ability to freely select any medical institution, I express hospital reputation as differences in the willingness to travel. In the second part of the study, I compute the probabilities of individual hospitals being selected by prospective patients. Hence, I identify those institutions, which are most likely to be visited.

3.3.1 The Travel Distance

I specify a regression model, which exposes the determinants of the distance a patient travels. The basis of the analysis is that people typically attend a nearby clinic to avoid unnecessary travel. Therefore, if choosing between identical hospitals, a patient will always select the nearest available option ([36]).

In order to measure the importance of reputation effects in hospital selection, I express differences between clinics as fluctuations in the willingness to travel. In particular, I identify those institutions in the sample, which stand out by attracting patients from further away. I define the distance travelled as a function of patient characteristics x_i and a vector of hospital specific dummy variables d_h. Thus, I demonstrate how the selection of a particular clinic changes the distance a patient is prepared to travel. As there are 18 institutions in the sample, I include 17 hospital dummies in the model. I estimate the coefficients β_{h1} through to β_{h17}, which represent hospital specific differences in the willingness to travel. The error term u_{ih} represents unobserved individual random effects, which are assumed to be i.i.d.

$$Dist_{ih} = x_i\beta + \beta_{h1}d_1 + ... + \beta_{h17}d_{17} + u_{ih} \tag{3.1}$$

Next, I decompose the systematic differences between clinics into observable and unobservable factors. I aim to assess the relative significance and magnitude of reputational considerations in choosing a hospital. To achieve this, I specify a panel model that defines the distance travelled in terms of patient attributes, hospital performance indicators and a hospital random effect e_h:

$$Dist_{ih} = x_i\beta + w_h\gamma + e_h + u_{ih} \tag{3.2}$$

3.3.2 The Hospital Selection

The second part of my analysis is concerned with computing the probabilities of individual hospitals being chosen. I specify the maximum distance a patient can travel to effectively restrict the number of medical institutions she has access to. I then study each patient's hospital selection based on the options available to her.

An individual's willingness to travel, determines the number of hospitals she can choose from. The set of options increases in size, as the patient travels further and additional clinics therefore become accessible. I identify those medical institutions each patient in the sample has access to. Therefore, I increase the dataset to a multiple of its original size. I analyse the explicit and implicit decisions a person makes. In particular, I consider not only the hospital a patient attends, but also capture those medical centres she chooses to avoid. The exact number of observations pertaining to each woman's hospital selection depends on how many clinics are accessible to her.

I manipulate the group of hospitals a patient chooses from by varying the maximum distance she can travel. This is done, based on the distribution of the distance variable in the dataset. I choose the 50th, 90th, 95th and 99th percentiles. The specified distances are 16.22 km, 40.96 km, 49.81 km and 63.27 km, respectively.

I construct a probit model that links a binary dependent variable with observable patient factors and the set of all hospitals in the sample. Y^*_{ih} is a latent variable, which captures the probability of patient i selecting hospital h from the set of medical institutions. Formally, the decision is modelled as follows:

$$Y^*_{ih} = x_i\beta + \beta_{h1}d_1 + ... + \beta_{h17}d_{17} + u_{ih} \tag{3.3}$$

In the model, Y^*_{ih} is proportional to an observable binary outcome variable Y_{ih}. Therefore, if patient i chooses hospital h from the set of options, $Y_{ih} = 1$. Otherwise, if hospital h is not selected, the binary variable $Y_{ih} = 0$.

The $1 \times (k+1)$ vector x_i contains observable patient characteristics, which influence the probability of patient i choosing to attend hospital h. The parameter β is a vector of size $(k+1) \times 1$ that captures the effect of individual patient factors on hospital selection. The inclusion of x_i in the model is critical, as patient specific features are likely to affect the individual decision of which medical institution to attend.

Each patient in the sample selects a hospital from N_i clinics accessible to her. However, the vector d_h in the model represents all hospitals in my dataset. As a result, I compute selection probabilities for individual clinics, which are not patient specific. Instead, I derive the probability of a hospital being visited, which depends on the maximum distance a pregnant woman is willing to travel.

The main parameter of interest to be estimated in the model is the vector β_h. It contains the probability of each hospital being visited, conditional on the assumed maximum travel distance. The disturbance term u_{ih} is a patient-specific hospital effect, which is i.i.d normal across all observations in the dataset.

3.4 The Results

3.4.1 The Willingness to Travel

I seek to evaluate the significance of reputation effects in individual hospital selection. I therefore begin my analysis by computing the differences in patients' willingness to travel, which are specific to individual institutions (Table 3.3).

I estimate a linear model, which links the distance travelled to patient characteristics and a set of hospital indicators. I use the dependent variable in its logarithmic form to obtain relative changes in the distance travelled. The hospital dummies in the regression model constitute a binary choice. Each of those variables takes on the value of one if the relevant hospital is chosen, and it is equal to zero otherwise. As there are in total 18 hospitals in my sample, the regression model features 17 dummy variables. The 18th clinic serves as a reference. It follows that variation in the willingness to travel is measured against the benchmark institution[2]. The variable *Distance to Nearest Hospital* captures the distance an individual travels, in order to reach the medical centre closest to her.

To verify the robustness of the estimates, I vary the clustering of the standard errors. Initially, I assume that all hospital choices are purely patient specific. As a result, I do not apply any clustering to the regression model. This approach rules out any tangible link between patient specific traits and the chosen medical institution. Next, I impose the assumption that patients, who attend the same clinic, share many common treatment characteristics and are exposed to the same hospital environment. For instance, women giving birth at the same maternity ward are likely treated by the same doctors and looked after by the same nursing staff. Therefore, I re-run my regression model and cluster the standard errors at the hospital level.

Table 3.3 shows that there is substantial variation in the willingness to travel, which is specific to individual hospitals. It appears that some medical institutions in my sample attract patients from further away, relative to other clinics. I control for the number of hospitals each pregnant woman has access to within 41 km of her home. The vast majority of my estimates are robust to changes in the clustering of the standard errors. This implies that differences in the willingness to travel are deeply ingrained in the decision making of individual patients. Also, it is interesting that coefficients controlling for personal attributes are mostly insignificant in both regression models. Hence, it appears that fluctuations in the willingness to travel are mainly hospital specific, and largely independent of patient characteristics. The only exception is the birth mode, which is statistically significant.

[2]The choice of the reference hospital does not affect the results.

Table 3.3 Linear Regression on the Distance Travelled

VARIABLE	log(Distance Travelled) (Cluster Variable: None)	log(Distance Travelled) (Cluster Variable: Hospital)
w18wies	−0.01	−0.01
	(0.03)	(0.02)
Age	−0.00	−0.00
	(0.00)	(0.00)
Born Australia	−0.02	−0.02
	(0.01)	(0.03)
Indigenous Status	0.02	0.02
	(0.07)	(0.05)
Length of Stay	0.00	0.00
	(0.01)	(0.00)
Married	0.02	0.02
	(0.04)	(0.04)
Eastern Region	−0.27***	−0.27
	(0.02)	(0.24)
North Western Region	−0.43***	−0.43
	(0.02)	(0.26)
C-Section	0.05**	0.05*
	(0.02)	(0.03)
First Birth	−0.01	−0.01
	(0.01)	(0.02)
Number of Hospitals within 41 km	−0.11***	−0.11***
	(0.00)	(0.04)
Distance to Nearest Hospital	0.04***	0.04**
	(0.00)	(0.02)
I51	1.65***	1.65***
	(0.09)	(0.11)
I67	2.52***	2.52***
	(0.05)	(0.24)
I76	2.66***	2.66***
	(0.08)	(0.23)
J00	0.42**	0.42***
	(0.20)	(0.08)
J49	0.66*	0.66***
	(0.35)	(0.19)
J55	2.92***	2.92***
	(0.16)	(0.16)
J62	3.29***	3.29***
	(0.31)	(0.13)
J70	1.22***	1.22***
	(0.11)	(0.13)
J73	2.34***	2.34***
	(0.07)	(0.23)
J81	1.61***	1.61***
	(0.03)	(0.18)
J95	1.43***	1.43***
	(0.03)	(0.14)
M66	1.37***	1.37***
	(0.09)	(0.12)
O23	1.94***	1.94***
	(0.03)	(0.10)
X04	1.31***	1.31***
	(0.03)	(0.13)
X06	2.28***	2.28***
	(0.03)	(0.20)
X13	1.15***	1.15***
	(0.03)	(0.09)
X73	1.20***	1.20***
	(0.04)	(0.22)
_cons	2.08***	2.08***
	(0.07)	(0.43)
N	11,967	
$Adjusted\ R^2$	0.55	

Robust Standard Errors in Parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

3.4.2 The Determinants of the Travel Distance

I now examine the variation in travel behaviour, which is linked to hospital characteristics. I separately explore the effect of hospital criteria and patient specific factors on the distance an individual travels, using a random effects model.

The regression output in table 3.4 suggests that the distance a person travels is determined by a variety of hospital criteria and patient specific factors. Interestingly, the relative complexity of a birth episode, which is indicated by the cost weight, does not affect the distance travelled. This suggests that women with elaborate treatment requirements apply the same criteria in their hospital selection as those with more straightforward medical needs. However, patients, who give birth by c-section, travel on average three per cent further compared to those giving birth naturally. This is an intuitive result, as people undergoing surgical procedures would likely have more treatment specific concerns when choosing a medical institution. Furthermore, it is evident that the distance travelled is only marginally shorter for first birth episodes. Therefore, the hospital choice does not critically depend on whether or not a woman is pregnant for the first time.

An expectant mother's hospital selection is affected significantly by her age, whereas her ethnic background and marital status do not determine the willingness to travel. An increase in the patient's age by just one year reduces the distance travelled by two per cent. Older women are therefore more inclined to attend a nearby hospital. This seems logical, as the medical risks related to childbirth increase with age. Hence, older patients are likely to choose local hospitals, which are easily accessible at all times ([54]). The patient's marital status does not impact on the distance she is prepared to travel. This finding is unexpected, as women in domestic relationships would likely consider additional factors, such as overnight accommodation for spouses in their decision making. The regression also reveals that there is no difference between women born in Australia and those born overseas. As a result, the patient's cultural background does not significantly influence her hospital choice. This is in line with the fact that an individual's indigenous status does not affect the distance she travels, either.

The number of hospitals in the local area is a key determinant of the distance patients travel to their chosen clinic ([32]). For the women in my sample, every additional medical centre within 41 km of their residence, reduces the distance travelled by eight per cent. Similarly, the travel behaviour of individual patients also depends on the location of the nearest hospital. In particular, the distance travelled increases by five percent for a one kilometer increase in the distance to the closest institution. This finding is intuitive, and it lends robustness to other estimates in the model. These results also illustrate the validity of my approach, which is based on the notion that patients are intuitively averse to unnecessary travel. However, despite the relatively even distribution of medical institutions, there are noticeable differences in the hospital selection of women from various parts of Melbourne. In particular, patients from southern and north western districts travel around five per cent further

Table 3.4 Linear Regression on the Distance Travelled

VARIABLE	log(Distance Travelled)
w18wies	0.04
	(0.03)
Age	-0.02^{**}
	(0.00)
Born Australia	-0.05
	(0.03)
Indigenous Status	0.10
	(0.08)
Length of Stay	-0.00
	(0.01)
Married	-0.02
	(0.04)
Eastern Region	-0.05^{***}
	(0.02)
North Western Region	-0.03
	(0.02)
C-Section	0.03^{**}
	(0.01)
First Birth	-0.02^{*}
	(0.01)
Number of Hospitals within 41 km	-0.08^{***}
	(0.00)
Distance to Nearest Hospital	0.05^{***}
	(0.00)
Bloodstream Infection Rate 2010/11	-0.35^{***}
	(0.02)
Number of Childbirths 2010/11	0.03^{***}
	(0.00)
Number of Stillbirths 2010/11	-0.02^{***}
	(0.00)
log(Households Income Q3)	3.35^{***}
	(0.09)
log(Households Income Q4)	1.60^{***}
	(0.05)
log(Number of Women in Hospital Area)	-5.26^{***}
	(0.11)
_cons	11.28^{***}
	(0.47)
N	11,967
$Adjusted\ R^2$	0.45

Standard Errors in Parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

compared to those from eastern parts of the city. This suggests that individual hospital choices are significantly influenced by the social environment a person is exposed to. Previous research shows that recommendations from family and friends affect hospital selection to a significant degree ([48]).

It is evident that clinic characteristics and performance criteria are key factors in the choice of a hospital ([44]). The patients in my sample are clearly reluctant to attend medical institutions with high numbers of healthcare related bloodstream infections. This effect is numerically large, as a one per cent increase in the infection rate reduces the distance travelled to that clinic by as much as 35%. The rate of bloodstream infections at individual hospitals is publicly available from the Institute of Health and Welfare. Therefore, the significant impact of this variable suggests women closely review and utilize the information available on the departmental website. This finding is further supported by the impact of childbirths in the previous year. For every additional birth at a clinic in 2010/11, the distance travelled increases by three per cent. It is interesting that the number of stillbirths from the previous year is also a significant determinant of the distance a patient is willing to travel. This is because the number of stillborn babies at individual clinics is not publicly available. For every additional stillbirth the distance travelled decreases by as much as two per cent. This relationship is strong, considering information regarding stillbirths is spread solely through unofficial channels.

To ensure the validity of my estimated coefficients, I take into consideration demographic differences across the various hospital catchment areas. I control for such factors in my model, using data from the 2011 Australian census. The number of women in the hospital catchment area controls for the population of fertile females in the neighbourhoods surrounding each hospital. The variable has a significant effect on the hospital selection of patients in the sample. For a one per cent increase in the number of women aged 17 to 49, the distance travelled to the local clinic decreases by 5.26%. This trend reiterates the notion that pregnant women will intuitively select a nearby hospital for obstetric care and childbirth episodes. Finally, it is apparent from the regression that the economic demographics of the hospital catchment area also influence the distance a patient travels. For a one per cent increase in the number of households in income quartile three, the distance travelled increases by 3.35%. Similarly, a one per cent increase in households in income quartile four increases the distance by 1.60%. This indicates that hospitals in relatively wealthy neighbourhoods attract patients from further away.

The analysis so far has revealed the factors, which influence the distance a patient travels to the hospital of her choice. Next, I further develop my results by computing hospital specific selection probabilities.

3.4.3 The Selection Mechanism

I now extend the analysis of hospital selection by studying how patients choose a medical institution from a given set of accessible alternatives. I define the maximum distance a person can travel and analyse her decision making based on the options available to her. In order to ensure the robustness of my results, I vary the maximum distance people can travel. As a result, I effectively alter the set of clinics a prospective patient is able to choose from.

For each woman in my data sample, I construct the set of hospitals within her range. The dependent variable in the model is binary, based on a patient's decision. It is equal to one for the hospital visited, and zero for all other foregone alternatives. I therefore consider the hospital a patient attends, as well as the institutions she avoids. Initially, I link the decision to visit a hospital with patient characteristics, as well as clinical performance criteria (Table 3.5).

I include patient specifics in the model to control for any individual effects, which may be correlated with the hospital a person attends. The regressor *Excess Travel* serves as a control variable for the location of the nearest hospital, in relation to a patient's home. The variable captures the distance to a hospital relative to the nearest alternative. The resulting distance ratio equals one for the hospital closest to a patient's home, and it is strictly larger than one for all other clinics. The results are largely consistent with my previous findings. The selection of a hospital from a set of accessible options is influenced by essentially the same factors as the willingness to travel. This is an intuitive result. As such, patients choose medical institutions mainly based on their recent performance.

A one percent increase in the rate of bloodstream infections decreases the probability of the hospital being selected by 14%. This estimate is in line with the first strand of my analysis, where a high rate of bloodstream infections adversely affects the distance a patient is prepared to travel. Similarly, the number of stillbirths also influences the probability of a clinic being chosen in the way one would expect. For every additional stillborn baby, the selection probability decreases by two per cent. The impact of stillbirths on hospital selection is therefore evident throughout the analysis. Finally, it is apparent that other patient specific factors, such as the location of the nearest hospital and the number of accessible alternatives also influence the probability of a medical institution being selected.

Next, I estimate the probabilities of individual clinics being selected from a set of available options (Table 3.6). I do this by including hospital dummies, rather than clinic specific performance measures, in my model. This allows me to relate selection criteria specifically to the individual hospitals in the dataset. The estimated coefficients for the hospital dummies can be interpreted as selection probabilities. These probabilities are conditional on the assumed willingness to travel. An important advantage of this method is that it incorporates both explicit and implicit decisions.

Table 3.5 Probit Regression on the Hospital Selection (Maximum Distance: 40.96 km)

VARIABLE	Hospital Chosen=1 (Cluster Variable: Hospital)	Hospital Chosen=1 dy/dx
log(Distance Travelled)	−0.42***	−0.06***
	(0.10)	(0.02)
w18wies	0.06	0.01
	(0.10)	(0.01)
Age	−0.00	−0.00
	(0.01)	(0.01)
Born Australia	−0.01	−0.01
	(0.04)	(0.01)
Indigenous Status	−0.08	−0.01
	(0.44)	(0.06)
Length of Stay	−0.01	−0.01
	(0.02)	(0.00)
Married	−0.06	−0.01
	(0.10)	(0.02)
Eastern Region	−0.07	−0.01
	(0.13)	(0.02)
North Western Region	−0.22	−0.03
	(0.24)	(0.03)
C-Section	0.02	0.01
	(0.07)	(0.01)
First Birth	−0.01	−0.02
	(0.03)	(0.02)
Number of Hospitals within 41 km	−0.11***	−0.02***
	(0.03)	(0.00)
Excess Travel	−0.03*	−0.03**
	(0.02)	(0.01)
Bloodstream Infection Rate 2010/11	−1.08***	−0.14***
	(0.24)	(0.03)
Number of Childbirths 2010/11	0.02**	0.01**
	(0.01)	(0.00)
Number of Stillbirths 2010/11	−0.05**	−0.02**
	(0.02)	(0.00)
log(Households Income Q3)	−2.15	−0.28*
	(1.40)	(0.17)
log(Households Income Q4)	0.88	0.12
	(0.61)	(0.08)
log(Number of Women in Hospital Area)	−0.97	−0.13
	(1.80)	(0.23)
_cons	3.01	
	(7.94)	
N	134,318	
Pseudo R^2	0.20	

Robust Standard Errors in Parentheses — * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3.6 Probit Regression on the Hospital Selection (Maximum Distance: 40.96 km)

VARIABLE	Hospital Chosen=1 (Cluster Variable: Hospital)	Hospital Chosen=1 dy/dx
log(Distance Travelled)	-0.34^{***}	-0.04^{***}
	(0.08)	(0.01)
w18wies	0.13	0.02
	(0.13)	(0.02)
Age	-0.01	-0.01
	(0.01)	(0.01)
Born Australia	0.02	0.02
	(0.04)	(0.05)
Indigenous Status	-0.12	-0.01
	(0.50)	(0.06)
Length of Stay	-0.03	-0.03
	(0.03)	(0.03)
Married	-0.07	-0.01
	(0.11)	(0.01)
Eastern Region	-0.20	-0.02
	(0.17)	(0.02)
North Western Region	-0.13	-0.02
	(0.32)	(0.04)
Excess Travel	-0.04^{*}	-0.03^{**}
	(0.03)	(0.01)
C-Section	0.03	0.04
	(0.09)	(0.03)
I51	-1.56^{***}	-0.19^{***}
	(0.07)	(0.01)
I67	-0.59^{**}	-0.07^{**}
	(0.23)	(0.03)
I76	-1.12^{***}	-0.14^{***}
	(0.12)	(0.02)
J00	-2.06^{***}	-0.25^{***}
	(0.15)	(0.02)
J49	-1.61^{***}	-0.20^{***}
	(0.05)	(0.01)
J55	-2.24^{***}	-0.27^{***}
	(0.25)	(0.03)
J62	-2.27^{***}	-0.28^{***}
	(0.27)	(0.03)
J70	-1.74^{***}	-0.21^{***}
	(0.06)	(0.01)
J73	-1.36^{***}	-0.16^{***}
	(0.13)	(0.02)
J81	0.52^{***}	0.06^{***}
	(0.14)	(0.02)
J95	0.03	0.00
	(0.10)	(0.01)
M66	-1.34^{***}	-0.16^{***}
	(0.02)	(0.02)
O23	-0.11	-0.01
	(0.08)	(0.01)
X04	-0.36^{***}	-0.04^{***}
	(0.08)	(0.01)
X06	0.56^{***}	0.07^{***}
	(0.15)	(0.02)
X13	1.04^{***}	0.13^{***}
	(0.04)	(0.01)
X73	-0.40^{**}	-0.05^{**}
	(0.19)	(0.02)
_cons	0.37	
	(0.49)	
N	134,318	
Pseudo R^2	0.27	

Robust Standard Errors in Parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

To determine the extent to which hospital selection is conditional on the willingness to travel, I manipulate the distance parameter based on the distribution of the underlying variable. I change the maximum distance a patient can travel to 16.22 km, 49.81 km and eventually 63.27 km. These respective values correspond to the 50th, 95th and 99th percentiles of the variable *Distance Travelled* in the data sample. All selection probabilities are computed against a reference institution, which is omitted from the model[3]. To start with, I fix the maximum distance a patient can travel at 40.96 km. The corresponding regression model is shown in table 3.6.

It is apparent from the estimated coefficients that the selection probabilities for individual hospitals vary significantly[4]. On the other hand, it is evident that all coefficients associated with treatment specific requirements are statistically insignificant. This result complements my previous findings, such that hospital selection is determined mostly by clinic related factors.

The model features the logarithm of the variable *Distance Travelled* as a regressor. This variable shows the effect of an increase in the distance travelled on the probability of a hospital being selected. However, it only captures an increase in the distance within the specified bracket. Thus, the estimated value conveys the effect of additional travel only within the predetermined set of accessible clinics. The estimated impact on the probability of a hospital being chosen is minimal. However, the high significance of the coefficient indicates that patients are generally averse to additional travel. Similarly, the significant coefficient on the variable *Excess Travel* reiterates the natural tendency among people to attend the nearest hospital available.

I now seek to verify the robustness of my results by controlling for additional patient characteristics, which likely affect the selection of an obstetric care facility (Table 3.7). I take into account the distribution of hospitals across Melbourne and the number of clinics individual patients have access to. Thus, I control for differences in the decision behaviour, which are due to some patients living in more secluded parts of the city. I include a variable that captures the number of clinics within 41 km of a patient's home. This is because the maximum distance a patient can travel is fixed at 40.96 km.

In addition, I control for those observations in my dataset, which represent first birth episodes. The reason is that the circumstances of a first birth may differ significantly from any subsequent childbirth. Women, who are pregnant for the first time, may be disposed to choose a hospital close to their home, so as to ensure easy access at all times. Furthermore, experience gained during the first birth may influence a patient's perception of hospitals in the future. As a result, the criteria applied in the selection of medical institutions might differ significantly for those women, who give birth for the first time.

---

[3]As before, the choice of the reference hospital has no bearing on the results.

[4]These systemic differences between individual institutions are robust to changes in the clustering of the standard errors.

Table 3.7 Probit Regression on the Hospital Selection (Robustness Check 1)

VARIABLE	Hospital Chosen=1 (Cluster Variable: Hospital)	Hospital Chosen=1 dy/dx	Hospital Chosen=1 (Cluster Variable: Hospital)	Hospital Chosen=1 dy/dx
log(Distance Travelled)	−0.45***	−0.05***	−0.34***	−0.04***
	(0.11)	(0.01)	(0.08)	(0.01)
w18wies	0.09	0.01	0.13	0.02
	(0.13)	(0.01)	(0.13)	(0.02)
Age	−0.01	−0.01	−0.01	−0.01
	(0.06)	(0.07)	(0.01)	(0.00)
Born Australia	0.02	0.01	0.02	0.02
	(0.05)	(0.05)	(0.04)	(0.05)
Indigenous Status	−0.11	−0.01	−0.12	−0.01
	(0.50)	(0.05)	(0.50)	(0.06)
Length of Stay	−0.02	−0.02	−0.03	−0.03
	(0.03)	(0.03)	(0.03)	(0.03)
Married	−0.05	−0.06	−0.07	−0.01
	(0.10)	(0.10)	(0.11)	(0.01)
Eastern Region	−0.13	-0.02	−0.20	−0.02
	(0.15)	(0.02)	(0.17)	(0.02)
North Western Region	−0.28	−0.03	−0.13	−0.02
	(0.34)	(0.03)	(0.32)	(0.04)
Excess Travel	−0.00	−0.04	−0.04	−0.04
	(0.01)	(0.03)	(0.05)	(0.03)
C-Section	0.04	0.01	0.03	0.04
	(0.08)	(0.09)	(0.09)	(0.03)
Number of Hospitals within 41 km	−0.14***	−0.02***		
	(0.04)	(0.00)		
First Birth	−0.09*	−0.02*		
	(0.05)	(0.00)		
First Birth x log(Distance Travelled)	−0.04**	−0.01		
	(0.02)	(0.00)		
I51	−1.28***	−0.15***	−1.56***	−0.19***
	(0.10)	(0.02)	(0.07)	(0.01)
I67	−0.54**	−0.06**	−0.59**	−0.07**
	(0.26)	(0.03)	(0.23)	(0.03)
I76	−0.96***	−0.11***	−1.12***	−0.14***
	(0.12)	(0.02)	(0.12)	(0.02)
J00	−2.25***	−0.26***	−2.06***	−0.25***
	(0.21)	(0.02)	(0.15)	(0.02)
J49	−1.88***	−0.22**	−1.61***	−0.20***
	(0.18)	(0.02)	(0.05)	(0.01)
J55	−2.32***	−0.27***	−2.24***	−0.27***
	(0.30)	(0.03)	(0.25)	(0.03)
J62	−2.41***	−0.18***	−2.27***	−0.28***
	(0.35)	(0.01)	(0.27)	(0.03)
J70	−1.54***	−0.13***	−1.74***	−0.21***
	(0.04)	(0.02)	(0.06)	(0.01)
J73	−1.14***	−0.09***	−1.36***	−0.16***
	(0.13)	(0.02)	(0.13)	(0.02)
J81	0.74***	0.03**	0.52***	0.06***
	(0.17)	(0.01)	(0.14)	(0.02)
J95	0.28**	0.03**	0.03	0.00
	(0.12)	(0.01)	(0.10)	(0.01)
M66	−1.39***	−0.16***	−1.34***	−0.16***
	(0.06)	(0.00)	(0.02)	(0.02)
O23	−0.18	−0.02	−0.11	−0.01
	(0.14)	(0.02)	(0.08)	(0.01)
X04	−0.09	−0.01	−0.36***	−0.04***
	(0.11)	(0.01)	(0.08)	(0.01)
X06	0.86***	0.10***	0.56***	0.07***
	(0.19)	(0.02)	(0.15)	(0.02)
X13	0.77***	0.09***	1.04***	0.13***
	(0.07)	(0.01)	(0.04)	(0.01)
X73	−0.22	−0.03	−0.40**	−0.05**
	(0.21)	(0.02)	(0.19)	(0.02)
_cons	1.77**		0.37	
	(0.74)		(0.49)	
N	134,318		134,318	
Pseudo R^2	0.30		0.27	

Robust Standard Errors in Parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

The hospital specific differences in table 3.7 are not affected by the number of clinics a patient has access to. Whereas it is evident that the number of hospitals in the local area influences her decision making, this effect does not explain the substantial variation in the selection probabilities. Instead, the estimated marginal effects are robust to the number of clinics a patient chooses from. This result underlines the importance of institution specific criteria in patient-hospital matching.

Hospital selection for first births differs from subsequent birth episodes. Specifically, women giving birth for the first time are particularly reluctant to travel further than necessary. However, the relative magnitude of the estimated relationship shows that such considerations are of minor importance. It therefore appears that even when giving birth for the first time, the characteristics of individual clinics remain the key determinant of hospital choice. This result indicates that information regarding the perceived quality of medical institutions is communicated through unofficial channels ([56]).

The examination of the selection probabilities yields the same key results as the first strand of my analysis. It seems that hospital traits, which influence the distance a patient is prepared to travel, also drive the decision of which hospital to select from a limited range of alternatives. These factors are sufficiently strong to render case specific requirements secondary in the decision making of prospective patients. This highlights the strategic importance of hospital performance criteria in a market, which allows patients to freely choose their preferred medical care provider. The option to review and compare the performance of hospitals greatly increases direct competition between institutions in the healthcare sector ([40]).

Finally, I examine whether variation in the selection probabilities is dependent on the willingness to travel (Table 3.8). Hence, I vary the maximum distance to determine whether the existent differences between hospitals are robust to such changes. By manipulating the travel parameter, I effectively alter the set of hospitals a patient chooses from. As the maximum travel distance decreases, the set of available options shrinks in size. By contrast, an increase in the willingness to travel means that additional clinics enter into the choice set. The analysis of the marginal effects shows that the relative ordering of the hospitals is essentially robust to changes in the willingness to travel. Although the estimated choice probabilities vary depending on the specified distance, the ranking of the hospitals remains nearly unchanged.

Table 3.8 Probit Regression on the Hospital Selection (Robustness Check 2)

VARIABLE	(Distance: 16.22 km) dy/dx	(Distance: 40.96 km) dy/dx	(Distance: 49.81 km) dy/dx	(Distance: 63.27 km) dy/dx
log(Distance Travelled)	0.08**	−0.04***	−0.05***	−0.05***
	(0.05)	(0.01)	(0.01)	(0.01)
w18wies	0.02	0.02	0.01	0.01
	(0.02)	(0.02)	(0.01)	(0.01)
Age	−0.00	−0.00	−0.00	−0.00
	(0.00)	(0.00)	(0.00)	(0.00)
Born Australia	0.00	0.00	0.00	0.00
	(0.01)	(0.01)	(0.00)	(0.00)
Indigenous Status	−0.10	−0.01	−0.01	−0.01
	(0.10)	(0.05)	(0.05)	(0.04)
Length of Stay	−0.00	−0.00	−0.00	−0.00
	(0.00)	(0.00)	(0.00)	(0.00)
Married	−0.03	−0.01	−0.01	−0.01
	(0.02)	(0.01)	(0.01)	(0.01)
Eastern Region	−0.11*	−0.03	−0.02	−0.02
	(0.06)	(0.02)	(0.02)	(0.02)
North Western Region	−0.03	−0.02	−0.02	−0.01
	(0.08)	(0.04)	(0.03)	(0.03)
Excess Travel	0.02**	−0.04	−0.06	−0.05
	(0.00)	(0.04)	(0.04)	(0.03)
C-Section	0.00	0.01	0.02	0.00
	(0.02)	(0.01)	(0.08)	(0.01)
First Birth	−0.02	−0.01	−0.01	−0.01
	(0.02)	(0.01)	(0.01)	(0.01)
First Birth x log(Distance Travelled)	0.01	0.01	0.00	0.00
	(0.01)	(0.00)	(0.01)	(0.00)
I51	−0.56***	−0.19***	−0.15***	−0.13***
	(0.03)	(0.01)	(0.01)	(0.01)
I67	−0.19**	−0.07***	−0.06***	−0.05***
	(0.07)	(0.02)	(0.02)	(0.01)
I76	−0.42***	−0.14***	−0.10***	−0.08***
	(0.04)	(0.02)	(0.01)	(0.01)
J00	−0.55***	−0.25***	−0.22***	−0.20***
	(0.02)	(0.02)	(0.02)	(0.02)
J49	−0.52***	−0.20***	−0.16***	−0.14***
	(0.02)	(0.01)	(0.01)	(0.01)
J55	−0.70***	−0.27***	−0.22***	−0.19***
	(0.07)	(0.03)	(0.02)	(0.02)
J62	−0.71***	−0.28***	−0.23***	−0.19***
	(0.07)	(0.03)	(0.03)	(0.02)
J70	−0.55***	−0.21***	−0.18***	−0.15***
	(0.02)	(0.01)	(0.01)	(0.01)
J73	−0.52***	−0.17***	−0.13***	−0.11***
	(0.02)	(0.02)	(0.01)	(0.01)
J81	−0.05	0.06***	0.07***	0.06***
	(0.05)	(0.01)	(0.01)	(0.00)
J95	0.15***	0.00	0.01	0.02**
	(0.03)	(0.01)	(0.01)	(0.00)
M66	−0.42***	−0.16***	−0.13***	−0.11***
	(0.02)	(0.01)	(0.00)	(0.00)
O23	−0.13***	−0.01	−0.00	0.01*
	(0.05)	(0.01)	(0.01)	(0.00)
X04	−0.18***	−0.04***	−0.03***	−0.02***
	(0.04)	(0.00)	(0.01)	(0.00)
X06	0.00	0.07***	0.07***	0.07***
	(0.01)	(0.02)	(0.01)	(0.01)
X13	0.04	0.13***	0.07***	0.04***
	(0.03)	(0.01)	(0.00)	(0.01)
X73	−0.23***	−0.05**	−0.03**	−0.02**
	(0.07)	(0.02)	(0.01)	(0.00)
N	44,230	134,318	162,417	188,138
Pseudo R^2	0.35	0.27	0.28	0.29

Robust Standard Errors in Parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

3.5 Concluding Remarks

This chapter contributes to the existing literature on hospital selection by studying the significance of reputation effects in patient-hospital matching. I use an administrative dataset from Australia, which contains detailed information on the hospital choices of pregnant women from the Melbourne metropolitan area. The basis of my analysis is the structure of the Australian healthcare sector, which allows private patients to freely select their preferred hospital from a wide range of options. I conduct this research through two separate strands of analysis. Firstly, I treat the distance patients travel to be indicative of a hospital's reputation. I therefore identify the factors, which influence the distance between a patient's residence and her preferred clinic. In the second part of my study, I examine the hospital selection of patients by controlling for the number of institutions they have access to. I specify the maximum distance a patient can travel and restrict the analysis only to hospitals within her range.

I compute clinic specific differences in the willingness to travel. Thus, I express reputation effects in hospital selection as differences in the distance a patient is prepared to travel. I complement that result by studying the determinants of the distance a patient travels to her preferred hospital. The results indicate that the travel decision is mainly influenced by a variety of clinic dependent characteristics. By contrast, it appears that patients' treatment requirements are comparatively insignificant in explaining the travel behaviour of expectant mothers.

I complement the above analysis by studying the hospital selection of women in the sample. In particular, I control the distance a patient can travel to construct a personalised set of accessible institutions. I then study individual decision making based on the available options. This approach allows me to capture not only the clinic a patient attends, but also the institutions she avoids. I compute marginal effects to obtain the probabilities of individual hospitals being selected. In line with my previous findings, it is apparent that there are significant differences between the medical centres in the sample. These differences are robust to a variety of factors, such as the number of hospitals a woman has access to within 41 km of her residence. Also, it is evident that variation in the selection probabilities is mainly due to hospital performance criteria. Finally, I find that the observed differences between individual clinics are robust to changes in the willingness to travel. I vary the maximum travel distance to manipulate the set of hospitals a person chooses from. However, it is apparent that the relative differences between medical centres remain almost entirely unaffected.

The key result of this study is that reputation effects are critical in the hospital selection of prospective patients. The main contribution of the chapter lies in the approach to modelling and estimating the importance of hospital reputation. In particular, my analysis complements the existing literature by illustrating the significance of reputation in the individual decision to attend a particular clinic. The analysis conducted in this chapter highlights the direct competition between hospitals in a market where patients can freely choose their preferred healthcare provider. I show that clinic attributes are the main determinant of which institution a person visits. In addition, it is apparent that the women in the sample conduct detailed research and compare the performance of various hospitals. As a result, it seems that patient factors and treatment requirements are merely of secondary importance in their decision making.

One limitation of this study is that the exact distance a patient travels is not known. This is due to patients' addresses not being included in the dataset for the sake of anonymity. However, my approach to estimating the distance travelled is consistent with the information provided and therefore does not affect the validity of the results. A further limitation of this chapter is that I do not control for the location of individual hospitals relative to one another. This lack of a spatial dimension in the analysis could influence the magnitude of the estimated hospital effects.

This research draws attention to the channels through which hospital performance information is disseminated. In particular, it is evident from my results that information about individual hospitals, which is available on the internet, greatly affects the behaviour of women in my dataset. Therefore, it appears that the publication of clinical performance indicators is helpful in directing patients towards medical institutions, which perform above average. This in turn provides an incentive to under-performing institutions to improve their care standards. Therefore, reputation can help eliminate inefficiencies and deliver productivity gains to the healthcare sector.

Chapter 4

Summary and Future Research

This thesis studies the importance of reputation effects in indefinite interactions. In the first two chapters I study decision making in a random matching economy of indefinite duration. Chapter three provides empirical evidence of reputational considerations based on an administrative dataset from the Australian healthcare sector.

Chapter one investigates the incentives generated by contagious punishment schemes and to what extent personal histories can help overcome extreme informational restrictions. By contrast, chapter two uses behavioural histories as an instrumental parameter to examine how subjects respond to having the option not to play. I find that reputational concerns influence individual choices in the way economic theory predicts. However, a key result of my research is that behavioural histories alone are not sufficient to promote substantial cooperation between anonymous strangers. Also, agents generally have a short memory when punishing opportunism by others. I find that subjects are averse to bearing a direct cost in order to maintain a favourable reputation in the future.

In chapter two, I identify a costless option to abstain as effective in enforcing cooperative behaviour between the subjects. When the option to opt out of an interaction is perceived to incur no direct cost, participants embrace the prospect of sustained cooperation in the economy. This is also reflected in the fact that people monitor their partner's decision history more carefully when the option to walk away yields a pay-off equivalent to the one shot Nash equilibrium.

An interesting research question for the future is the impact of an attractive outside option on cooperation rates in random matching communities. Based on my own results, it would seem likely that such an option not to play would be readily utilized in the continuation game and may therefore lead to even higher efficiency levels in the market. However, in practice an attractive option to abstain could promote defection rather than sociable behaviour, with the potential punishment being ineffective. Alternatively, it would be worthwhile to change the pay-off structure associated with the outside option, such that it is asymmetric. Thus, players can impose a more direct penalty on deviators by choosing not to play. A drawback of this extension is that it significantly complicates the theoretical analysis and makes the derivation of equilibria cumbersome.

In chapter three of my thesis, I add to the existing literature on patient-hospital matching. I analyse the importance of hospital reputation in the decision making of pregnant women. I propose econometric modelling techniques to gauge the importance of reputation effects using patient administrative data. My contribution is novel in its approach and effectively complements the existent research on hospital reputation. Empirical evidence on the impact of reputation effects in the healthcare sector has to date produced conflicting and inconclusive results. I measure the reputation of medical institutions using two separate research methods. Initially, I compute hospital specific differences in the distance patients are willing to travel. My approach is based on the notion that individuals will generally attend nearby clinics so as to avoid unnecessary travel. Subsequently, I derive selection probabilities for the institutions in my sample, by constructing a set of accessible hospitals for each patient. My results clearly indicate that substantial differences exist among individual hospitals, which are apparent throughout my analysis.

One possible addition to the model I propose is to incorporate a spatial dimension. This would further improve the accuracy of the estimated hospital effects, as it allows for the location of hospitals in relation to one another to be considered. The extension requires the use of an alternative-specific multinomial probit model, which permits the use of case-specific regressors.

The importance of reputation has increasingly become the subject of economic research over the last 40 years. This thesis contributes to that literature by examining the preference structures of long lived agents. I extend the conventional economic modelling of indefinite relationships and study how reputation paired with the option not to play can help improve market efficiency.

Bibliography

[1] *2011 Census Overview* [2018], https://www.abs.gov.au. 2018-07-02.

[2] Bó, P. D. [2007], 'Social norms, cooperation and inequality', *Economic Theory* **30**(1), 89–105.

[3] Burns, L. R. and Wholey, D. R. [1992], 'The impact of physician characteristics in conditional choice models for hospital care', *Journal of Health Economics* **11**(1), 43–62.

[4] Bó, P. D. [2005], 'Cooperation under the shadow of the future: Experimental evidence from infinitely repeated games', *American Economic Review* **95**(5), 1591–1604.

[5] Camera, G. and Casari, M. [2009], 'Cooperation among strangers under the shadow of the future', *American Economic Review* **99**(3), 979–1005.

[6] Camera, G. and Gioffré, A. [2013], 'A tractable analysis of contagious equilibria'.

[7] Cheng, S.-H., Yang, M.-c. and Chiang, T.-l. [2003], 'Patient satisfaction with and recommendation of a hospital: effects of interpersonal and technical aspects of hospital care', *International Journal for Quality in Health Care* **15**(4), 345–355.

[8] Cooper, R., V. DeJong, D., Forsythe, R. and Ross, T. [1996], 'Cooperation without reputation: Experimental evidence from prisoner's dilemma games', **12**, 187–218.

[9] Dal Bó, P. and Fréchette, G. R. [2011], 'The evolution of cooperation in infinitely repeated games: Experimental evidence', *American Economic Review* **101**(1), 411–29.

[10] Dal Bó, P. and Fréchette, G. R. [2018], 'On the determinants of cooperation in infinitely repeated games: A survey', *Journal of Economic Literature* **56**(1), 60–114.

[11] Duffy, J. and Ochs, J. [2009], 'Cooperative behavior and the frequency of social interaction', *Games and Economic Behavior* **66**(2), 785–812.

[12] Ehrhart, K.-M. and Keser, C. [1999], 'Mobility and cooperation: On the run'.

[13] Ellison, G. [1994], 'Cooperation in the prisoner's dilemma with anonymous random matching', *Review of Economic Studies* **61**(3), 567–588.

[14] Evans, R. and Thomas, J. P. [1997], 'Reputation and experimentation in repeated games with two long-run players', *Econometrica* **65**(5), 1153–1173.

[15] Finkelstein, B. S., Singh, J., Silvers, J. B., Neuhauser, D. and Rosenthal, G. E. [1998], 'Patient and hospital characteristics associated with patient assessments of hospital obstetrical care', *Medical Care* **36**(8), AS68–AS78.

[16] Fischbacher, U. [2007], 'z-tree: Zurich toolbox for ready-made economic experiments', *Experimental Economics* **10**(2), 171–178.

[17] Friedman, J. W. [1971], 'A non-cooperative equilibrium for supergames', *Review of Economic Studies* **38**(1), 1–12.

[18] Fudenberg, D. and Levine, D. K. [1989], 'Reputation and equilibrium selection in games with a patient player', *Econometrica* **57**(4), 759–778.

[19] Fujiwara-Greve, T. and Yasuda, Y. [2009], Cooperation in repeated prisoner's dilemma with outside options, GRIPS Discussion Papers 09-10, National Graduate Institute for Policy Studies.

[20] Geweke, J., Gowrisankaran, G. and Town, R. J. [2003], 'Bayesian inference for hospital quality in a selection model', *Econometrica* **71**(4), 1215–1238.

[21] Gooding, S. K. S. [1994], 'Hospital outshopping and perceptions of quality: Implications for public policy', *Journal of Public Policy Marketing* **13**(2), 271–280.

[22] Greiner, B. [2004], The online recruitment system orsee - a guide for the organization of experiments in economics, Papers on strategic interaction, Max Planck Institute of Economics, Strategic Interaction Group.

[23] Grosskopf, B. and Sarin, R. [2010], 'Is reputation good or bad? an experiment', *American Economic Review* **100**(5), 2187–2204.

[24] Hauk, E. [2003], 'Multiple prisoner's dilemma games with(out) an outside option: an experimental study', *Theory and Decision* **54**(3), 207–229.

[25] Hauk, E. and Nagel, R. [2001], 'Choice of partners in multiple two-person prisoner's dilemma games: An experimental study', *The Journal of Conflict Resolution* **45**(6), 770–793.

[26] Henrich, J., McElreath, R., Barr, A., Ensminger, J., Barrett, C., Bolyanatz, A., Cardenas, J. C., Gurven, M., Gwako, E., Henrich, N., Lesorogol, C., Marlowe, F., Tracer, D. and Ziker, J. [2006], 'Costly punishment across human societies', *Science* **312**(5781), 1767–1770.

[27] Hibbard, J. H., Stockard, J. and Tusler, M. [2005], 'Hospital performance reports: impact on quality, market share, and reputation', *Health affairs* **24**(4), 1150–1160.

[28] Kamei, K. [2017], 'Cooperation and endogenous repetition in an indefinitely repeated social dilemma', *Durham Business School, Working Paper No. 08* .

[29] Kandori, M. [1992], 'Social norms and community enforcement', *The Review of Economic Studies* **59**(1), 63–80.

[30] Kreps, D. M., Milgrom, P., Roberts, J. and Wilson, R. [1982], 'Rational cooperation in the finitely repeated prisoners' dilemma', *Journal of Economic Theory* **27**(2), 245 – 252.

[31] Kreps, D. M. and Wilson, R. [1982], 'Reputation and imperfect information', *Journal of economic theory* **27**(2), 253–279.

[32] Lane, P. and Lindquist, J. [1988], 'Hospital choice: A summary of the key empirical and hypothetical findings of the 1980s', *Journal of Health Care Marketing* **8**(4), 5–19.

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

[34] *Learn about your local Hospital* [2018], https://www.hospitals.gov.au. 2018-06-28.

[35] Lewis, S. and Willis, K. [2018], 'Do you really need private health insurance? here's what you need to know before deciding', https://www.theconversation.com.au. 2018-06-07.

[36] Luft, H., W. Garnick, D., H. Mark, D., J. Peltzman, D., S. Phibbs, C., Lichtenberg, E. and J. McPhee, S. [1990], 'Does quality influence choice of hospital?', *JAMA : the journal of the American Medical Association* .

[37] MacLeod, W. B. and Malcomson, J. M. [1988], 'Reputation and hierarchy in dynamic models of employment', *Journal of Political Economy* **96**(4), 832–854.

[38] *Medicare Levy Surcharge* [2018], https://www.ato.gov.au. 2018-08-21.

[39] *Melbourne Population* [2018], https://www.population.net.au/melbourne-population. 2018-06-21.

[40] Montefiori, M. [2008], 'Information vs advertising in the market for hospital care', *International journal of health care finance and economics* **8**(3), 145–162.

[41] Nosenzo, D. and Tufano, F. [2017], The Effect of Voluntary Participation on Cooperation, Discussion Papers 2017-12, The Centre for Decision Research and Experimental Economics, School of Economics, University of Nottingham.

[42] Orbell, J. M. and Dawes, R. M. [1993], 'Social welfare, cooperators' advantage, and the option of not playing the game', *American Sociological Review* **58**(6), 787–800.

[43] Ostrom, E. [2015], *Governing the commons*, Cambridge university press.

[44] Pope, D. [2009], 'Reacting to rankings: evidence from america's "best hospitals"', *Journal of Health Economics* **28**, 1154–1165.

[45] Riedl, A. and Ule, A. [2002], 'Exclusion and cooperation in social network experiments', *Unpublished Paper, CREED, University* .

[46] Rosenthal, R. W. [1979], 'Sequences of games with varying opponents', *Econometrica: Journal of the Econometric Society* pp. 1353–1366.

[47] Roth, A. and Murnighan, J. [1978], 'Equilibrium behavior and repeated play of the prisoner's dilemma', *Journal of Mathematical Psychology* **17**(2), 189–198.

[48] Schwartz, L. M., Woloshin, S. and Birkmeyer, J. D. [2005], 'How do elderly patients decide where to go for major surgery? telephone interview survey', *Bmj* **331**(7520), 821.

[49] Schwartz, S. T., Young, R. A. and Zvinakis, K. [2000], 'Reputation without repeated interaction: A role for public disclosures', *Review of Accounting Studies* **5**(4), 351–375.

[50] Selten, R. [1978], 'The chain store paradox', *Theory and Decision* **9**(2), 127–159.

[51] Semmann, D., Krambeck, H.-J. and Milinski, M. [2004], 'Strategic investment in reputation', *Behavioral Ecology and Sociobiology* **56**(3), 248–252.

[52] Stahl II, D. O. [1991], 'The graph of prisoners' dilemma supergame payoffs as a function of the discount factor', *Games and Economic Behavior* **3**(3), 368–384.

[53] Takahashi, S. [2010], 'Community enforcement when players observe partners' past play', *Journal of Economic Theory* **145**(1), 42 – 62.

[54] Varkevisser, M. and van der Geest, S. A. [2007], 'Why do patients bypass the nearest hospital? an empirical analysis for orthopaedic care and neurosurgery in the netherlands', *The European Journal of Health Economics* **8**(3), 287–295.

[55] Wilson, A. J. and Wu, H. [2017], 'At-will relationships: How an option to walk away affects cooperation and efficiency', *Games and Economic Behavior* **102**, 487 – 507.

[56] Wolinsky, F. and Kurz, R. [1984], 'How the public chooses and views hospitals', *Hospital and Health Services Administration* **29**, 58–67.

[57] Zhang, X., Hauck, K. and Zhao, X. [2013], 'Patient safety in hospitals – a bayesian analysis of unobservable hospital and specialty level risk factors', *Health economics* **22**(7).

Appendix A

Appendices For Chapter One

The Experimental Instructions

A.0.1 Instructions Control Treatment

Welcome to the FEELE lab,

You are about to participate in a session on decision making. You will be paid for taking part in cash. Your cash earnings depend partly on your own decisions, the decisions of other participants, and partly on chance. Throughout this session you will not be deceived or lied to in any way. Please turn off your mobile phones now before we begin.

Throughout the entire session you will be interacting with other participants through the computer terminals. It is essential that you do not talk or try to communicate with other participants in any way. Any attempt to communicate with other people will result in you being asked to leave without payment.

You should, in front of you, find a pay-off matrix, a pen and a blank sheet of paper. If you do not have all these items, please raise your hand now.

In this session there are two groups of ten people. It follows that you will be grouped with nine other people. You will be randomly matched with one of them at a time to interact with them. The pay-off matrix in front of you accurately describes the nature of that interaction. As you can see, each of the two players can select one of two possible action choices: Left or Right. If both players choose Left, they will each receive a payoff of 30 tokens. On the other hand, if one player chooses Left and the other chooses Right, the player who chose Right will receive a payoff of 40 tokens, whereas the player who chose Left will earn a payoff of 5 tokens. Finally, if both players choose Right, they will receive a payoff of 15 tokens each.

For each round of the interaction you will not know the identity of the participant you are paired with and you will have to select one of the action choices given on the payoff matrix (Left or Right). You can enter your choice into the computer by choosing either Left or Right. You confirm your decision by clicking the ‘Confirm’ button. At the end of each round you will be informed of the outcome of your interaction with the participant you are currently paired with. In particular, you will be informed of the other person’s action choice and the payoff you earned.

After any given round there is a chance that the interaction will not continue. In particular, there is a 1 in 10 chance after each round that the interaction will end. This means that in 1 case out of 10 the interaction will come to an end. However, if the interaction does continue, you will be matched with a new participant. You will not know the identity of that person and neither will they know yours. This process will be repeated until the interaction ends. There will be three such interactions.

At the end of the session some demographic data will be collected from you and afterwards you will be paid. This is done by converting your cumulative earnings for the second and third interactions into money. The fixed exchange rate is such that 60 tokens = £1. In addition, you will also receive a £3 show up fee. Everybody will be paid in private and you are under no obligation to tell other participants how much money you have earned.

Any Questions? Please raise your hand and the experimenter will attend to you individually.

A.0.2 Instructions Reputation Treatment

Welcome to the FEELE lab,

You are about to participate in a session on decision making. You will be paid for taking part in cash. Your cash earnings depend partly on your own decisions, the decisions of other participants, and partly on chance. Throughout this session you will not be deceived or lied to in any way. Please turn off your mobile phones now before we begin.

Throughout the entire session you will be interacting with other participants through the computer terminals. It is essential that you do not talk or try to communicate with other participants in any way. Any attempt to communicate with other people will result in you being asked to leave without payment.

You should, in front of you, find a pay-off matrix, a pen and a blank sheet of paper. If you do not have all these items, please raise your hand now.

In this session there are two groups of ten people. It follows that you will be grouped with nine other people. You will be randomly matched with one of them at a time to interact with them. The pay-off matrix in front of you accurately describes the nature of that interaction. As you can see, each of the two players can select one of two possible action choices: Left or Right. If both players choose Left, they will each receive a payoff of 30 tokens. On the other hand, if one player chooses Left and the other chooses Right, the player who chose Right will receive a payoff of 40 tokens, whereas the player who chose Left will earn a payoff of 5 tokens. Finally, if both players choose Right, they will receive a payoff of 15 tokens each.

For each round of the interaction you will not know the identity of the participant you are paired with and you will have to select one of the action choices given on the payoff matrix (Left or Right). You can enter your choice into the computer by choosing either Left or Right. You confirm your decision by clicking the ‘Confirm’ button. At the end of each round you will be informed of the outcome of your interaction with the participant you are currently paired with. In particular, you will be informed of the other person’s action choice and the payoff you earned.

After any given round there is a chance that the interaction will not continue. In particular, there is a 1 in 10 chance after each round that the interaction will end. This means that in 1 case out of 10 the interaction will come to an end. However, if the interaction does continue, you will be matched with a new participant. You will not know the identity of that person and neither will they know yours. However, from the second period of the interaction onwards you will be able to see for the person you are paired with in the current round, the decisions they have made in each of the previous rounds of the interaction. Similarly, your partner will also be able to review your choices in all previous rounds.

For example, if the interaction continues to the third round, you will at the beginning of that round be able to review the previous decisions of the participant you are matched with in the third round. You can do this before selecting an action choice (Left or Right) for the current round of the interaction. You will only see a person's respective action choices in each of the previous rounds of the interaction and receive no other information. This process will be repeated until the interaction ends. There will be three such interactions.

At the end of the session some demographic data will be collected from you and afterwards you will be paid. This is done by converting your cumulative earnings for the second and third interactions into money. The fixed exchange rate is such that 60 tokens = £1. In addition, you will also receive a £3 show up fee. Everybody will be paid in private and you are under no obligation to tell other candidates how much money you have earned.

Any Questions? Please raise your hand and the experimenter will attend to you individually.

A.0.3 Instructions Opt Out Treatment

Welcome to the FEELE lab,

You are about to participate in a session on decision making. You will be paid for taking part in cash. Your cash earnings depend partly on your own decisions, the decisions of other participants, and partly on chance. Throughout this session you will not be deceived or lied to in any way. Please turn off your mobile phones now before we begin.

Throughout the entire session you will be interacting with other participants through the computer terminals. It is essential that you do not talk or try to communicate with other participants in any way. Any attempt to communicate with other people will result in you being asked to leave without payment.

You should, in front of you, find a pay-off matrix, a pen and a blank sheet of paper. If you do not have all these items, please raise your hand now.

In this session there are two groups of ten people. It follows that you will be grouped with nine other people. You will be randomly matched with one of them at a time to interact with them. The pay-off matrix in front of you accurately describes the nature of that interaction. As you can see, each of the two players can select one of three possible action choices: Left, Right or Opt Out. If both players choose Left, they will each receive a payoff of 30 tokens. On the other hand, if one player chooses Left and the other chooses Right, the player who chose Right will receive a payoff of 40 tokens, whereas the player who chose Left will earn a payoff of 5 tokens. Finally, if both players choose Right, they will receive a payoff of 15 tokens each.

There is also an option to Opt Out, that is, you can choose not to play. In that case, regardless of the other person's choice you will both receive 10 tokens. Similarly, if your partner chooses to opt out, you will both receive 10 tokens as well, regardless of your choice.

For each round of the interaction you will not know the identity of the participant you are paired with and you will have to select one of the action choices given on the payoff matrix (Left, Right or Opt Out). You can enter your choice into the computer by choosing either Left, Right or Opt Out. You confirm your decision by clicking the 'Confirm' button. At the end of each round you will be informed of the outcome of your interaction with the participant you are currently paired with. In particular, you will be informed of the other person's action choice and the payoff you earned.

After any given round there is a chance that the interaction will not continue. In particular, there is a 1 in 10 chance after each round that the interaction will end. This means that in 1 case out of 10 the interaction will come to an end. However, if the interaction does continue, you will be matched with a new participant. You will not know the identity of that person and neither will they know yours. This process will be repeated until the interaction ends. There will be three such interactions.

At the end of the session some demographic data will be collected from you and afterwards you will be paid. This is done by converting your cumulative earnings for the second and third interactions into money. The fixed exchange rate is such that 60 tokens = £1. In addition, you will also receive a £3 show up fee. Everybody will be paid in private and you are under no obligation to tell other participants how much money you have earned.

Any Questions? Please raise your hand and the experimenter will attend to you individually.

A.0.4 Instructions Opt Out-Reputation Treatment

Welcome to the FEELE lab,

You are about to participate in a session on decision making. You will be paid for taking part in cash. Your cash earnings depend partly on your own decisions, the decisions of other participants, and partly on chance. Throughout this session you will not be deceived or lied to in any way. Please turn off your mobile phones now before we begin.

Throughout the entire session you will be interacting with other participants through the computer terminals. It is essential that you do not talk or try to communicate with other participants in any way. Any attempt to communicate with other people will result in you being asked to leave without payment.

You should, in front of you, find a pay-off matrix, a pen and a blank sheet of paper. If you do not have all these items, please raise your hand now.

In this session there are two groups of ten people. It follows that you will be grouped with nine other people. You will be randomly matched with one of them at a time to interact with them. The pay-off matrix in front of you accurately describes the nature of that interaction. As you can see, each of the two players can select one of three possible action choices: Left, Right or Opt Out. If both players choose Left, they will each receive a payoff of 30 tokens. On the other hand, if one player chooses Left and the other chooses Right, the player who chose Right will receive a payoff of 40 tokens, whereas the player who chose Left will earn a payoff of 5 tokens. Finally, if both players choose Right, they will receive a payoff of 15 tokens each.

There is also an option to Opt Out, that is, you can choose not to play. In that case, regardless of the other person's choice you will both receive 10 tokens. Similarly, if your partner chooses to opt out, you will both receive 10 tokens as well, regardless of your choice.

For each round of the interaction you will not know the identity of the participant you are paired with and you will have to select one of the action choices given on the payoff matrix (Left, Right or Opt Out). You can enter your choice into the computer by choosing either Left, Right or Opt Out. You confirm your decision by clicking the 'Confirm' button. At the end of each round you will be informed of the outcome of your interaction with the participant you are currently paired with. In particular, you will be informed of the other person's action choice and the payoff you earned.

After any given round there is a chance that the interaction will not continue. In particular, there is a 1 in 10 chance after each round that the interaction will end. This means that in 1 case out of 10 the interaction will come to an end. However, if the interaction does continue, you will be matched with a new participant. You will not know the identity of that person and neither will they know yours. However, from the second period of the interaction onwards you will be able to see for the person you are paired with in the current round, the decisions they have made in each of the previous rounds of the interaction. Similarly, your partner will also be able to review your choices in all previous rounds. For example, if the interaction continues to the third round, you will at the beginning of that round be able to review the previous decisions of the participant you are matched with in the third round. You can do this before selecting an action choice (Left, Right or Opt Out) for the current round of the interaction. You will only see a person's respective action choices in each of the previous rounds of the interaction and receive no other information. This process will be repeated until the interaction ends. There will be three such interactions.

At the end of the session some demographic data will be collected from you and afterwards you will be paid. This is done by converting your cumulative earnings for the second and third interactions into money. The fixed exchange rate is such that 60 tokens = £1. In addition, you will also receive a £3 show up fee. Everybody will be paid in private and you are under no obligation to tell other candidates how much money you have earned.

Any Questions? Please raise your hand and the experimenter will attend to you individually.

Appendix B

Appendices For Chapter Two

Experimental Instructions

B.0.1 Instructions Costless Outside Option Treatment

Welcome to the FEELE lab,

You are about to participate in a session on decision making. You will be paid for taking part in cash. Your cash earnings depend partly on your own decisions, the decisions of other participants, and partly on chance. Throughout this session you will not be deceived or lied to in any way. Please turn off your mobile phones now before we begin.

Throughout the entire session you will be interacting with other participants through the computer terminals. It is essential that you do not talk or try to communicate with other participants in any way. Any attempt to communicate with other people will result in you being asked to leave without payment.

You should, in front of you, find a pay-off matrix, a pen and a blank sheet of paper. If you do not have all these items, please raise your hand now.

In this session there is a group of ten people. It follows that you are grouped with nine other people. You will be randomly matched with one of them at a time to interact with them. The pay-off matrix in front of you accurately describes the nature of that interaction. As you can see, each of the two players can select one of three possible action choices: Left, Right or Opt Out. If both players choose Left, they will each receive a payoff of 30 tokens. On the other hand, if one player chooses Left and the other chooses Right, the player who chose Right will receive a payoff of 40 tokens, whereas the player who chose Left will earn a payoff of 5 tokens. Finally, if both players choose Right, they will receive a payoff of 15 tokens each.

There is also an option to Opt Out, that is, you can choose not to play. In that case, regardless of the other person's choice you will both receive 15 tokens. Similarly, if your partner chooses to Opt Out, you will both receive 15 tokens as well, regardless of your choice.

For each round of the interaction you will not know the identity of the participant you are paired with and you will have to select one of the action choices given on the payoff matrix (Left, Right or Opt Out). You can enter your choice into the computer by choosing either Left, Right or Opt Out. You confirm your decision by clicking the 'Confirm' button. At the end of each round you will be informed of the outcome of your interaction with the participant you are currently paired with. In particular, you will be informed of the other person's action choice and the payoff you earned.

After any given round there is a chance that the interaction will not continue. In particular, there is a 1 in 10 chance after each round that the interaction will end. This means that in 1 case out of 10 the interaction will come to an end. However, if the interaction does continue, you will be matched with a new participant. You will not know the identity of that person and neither will they know yours. This process will be repeated until the interaction ends. There will be three such interactions.

At the end of the session some demographic data will be collected from you and afterwards you will be paid. This is done by converting your cumulative earnings for the second and third interactions into money. The fixed exchange rate is such that 60 tokens = £1. In addition, you will also receive a £3 show up fee. Everybody will be paid in private and you are under no obligation to tell other participants how much money you have earned.

Any Questions? Please raise your hand and the experimenter will attend to you individually.

B.0.2 Instructions Costless Outside Option-Reputation Treatment

Welcome to the FEELE lab,

You are about to participate in a session on decision making. You will be paid for taking part in cash. Your cash earnings depend partly on your own decisions, the decisions of other participants, and partly on chance. Throughout this session you will not be deceived or lied to in any way. Please turn off your mobile phones now before we begin.

Throughout the entire session you will be interacting with other participants through the computer terminals. It is essential that you do not talk or try to communicate with other participants in any way. Any attempt to communicate with other people will result in you being asked to leave without payment.

You should, in front of you, find a pay-off matrix, a pen and a blank sheet of paper. If you do not have all these items, please raise your hand now.

In this session there is a group of ten people. It follows that you are grouped with nine other people. You will be randomly matched with one of them at a time to interact with them. The pay-off matrix in front of you accurately describes the nature of that interaction. As you can see, each of the two players can select one of three possible action choices: Left, Right or Opt Out. If both players choose Left, they will each receive a payoff of 30 tokens. On the other hand, if one player chooses Left and the other chooses Right, the player who chose Right will receive a payoff of 40 tokens, whereas the player who chose Left will earn a payoff of 5 tokens. Finally, if both players choose Right, they will receive a payoff of 15 tokens each.

There is also an option to Opt Out, that is, you can choose not to play. In that case, regardless of the other person's choice you will both receive 15 tokens. Similarly, if your partner chooses to Opt Out, you will both receive 15 tokens as well, regardless of your choice.

For each round of the interaction you will not know the identity of the participant you are paired with and you will have to select one of the action choices given on the payoff matrix (Left, Right or Opt Out). You can enter your choice into the computer by choosing either Left, Right or Opt Out. You confirm your decision by clicking the 'Confirm' button. At the end of each round you will be informed of the outcome of your interaction with the participant you are currently paired with. In particular, you will be informed of the other person's action choice and the payoff you earned.

After any given round there is a chance that the interaction will not continue. In particular, there is a 1 in 10 chance after each round that the interaction will end. This means that in 1 case out of 10 the interaction will come to an end. However, if the interaction does continue, you will be matched with a new participant. You will not know the identity of that person and neither will they know yours. However, from the second period of the interaction onwards you will be able to see for the person you are paired with in the current round, the decisions they have made in each of the previous rounds of the interaction. Similarly, your partner will also be able to review your choices in all previous rounds. For example, if the interaction continues to the third round, you will at the beginning of that round be able to review the previous decisions of the participant you are matched with in the third round. You can do this before selecting an action choice (Left, Right or Opt Out) for the current round of the interaction. You will only see a person's respective action choices in each of the previous rounds of the interaction and receive no other information. This process will be repeated until the interaction ends. There will be three such interactions.

At the end of the session some demographic data will be collected from you and afterwards you will be paid. This is done by converting your cumulative earnings for the second and third interactions into money. The fixed exchange rate is such that 60 tokens = £1. In addition, you will also receive a £3 show up fee. Everybody will be paid in private and you are under no obligation to tell other candidates how much money you have earned.

Any Questions? Please raise your hand and the experimenter will attend to you individually.