

Price Setting Rules and Competition

by

Mia Katina Tam

Submitted to the University of Adelaide

in partial fulfilment of

the requirements for the degree of

Doctor of Philosophy

in

Economics

April 2023

Contents

List of Tables	iii
List of Figures	iv
Abstract	v
Declaration	vii
Acknowledgements	viii
1 Introduction	1
2 Infinitely repeated games	7
2.1 Introduction	7
2.2 Theoretical background	15
2.2.1 Stage game	15
2.2.2 Pricing frequency and collusion	18
2.3 Experimental design	22
2.3.1 Hypothesis	25
2.4 Results	29
2.4.1 Average payoffs	31
2.4.2 Dynamics	35
2.5 Conclusion	39
3 Continuous time	41
3.1 Introduction	41
3.2 Model and implementation	46

3.2.1	Model	47
3.2.2	Parameters and Static Equilibrium.	48
3.3	Hypothesis and experimental design	50
3.3.1	Experimental implementation and procedure	52
3.4	Results	55
3.5	Conclusion	62
4	Behavioural mechanism	65
4.1	Introduction	65
4.2	Experimental Design and Hypothesis	67
4.3	Results	70
4.4	Conclusion	75
5	Conclusion	77
	Appendices	80
A.1	Chapter 2 - Instructions	81
A.2	Chapter 2 - Individual markets	105
A.3	Chapter 3 - Instructions	115
A.4	Chapters 3 and 4 - Control questions	124
A.5	Chapter 4 - Instructions	125
	References	130

List of Tables

2.1	Treatments and summary of parameters.	23
2.2	Demographics across treatments	29
2.3	Average price of independent observations.	30
2.4	Price proportions in each treatment.	31
2.5	Average supergame payoffs for independent observations.	31
2.6	Pairwise rank-sum tests of equality.	33
2.7	OLS regression on market supergame payoffs (clustering on matching group).	33
2.8	Pairwise test on equality of coefficients.	34
2.9	Hypothesis and rank-sum test on supergame payoffs between treatments.	34
2.10	Mann-Whitney rank-sum test, p-value for difference between treatments for each supergame	38
3.1	Parameters and Nash, joint-profit maximising, and best shot deviation prices and profits.	50
3.2	Demographics across treatments	58
3.3	Average price of independent observations.	58
3.4	Summary statistics for moving average $\bar{p} - \underline{p}$	60
4.1	Difference in deviation profit, and loss from signalling the JPM price instead of Nash, for each treatment.	68
4.2	Demographics across treatments	70
4.3	Average price of independent observations.	71
4.4	Summary statistics for moving average $\bar{p} - \underline{p}$	72

List of Figures

2.1	Possible equilibria with two firms.	21
2.2	Average sub-period price in all 5 supergames in δ_h (top) and δ_l (bottom).	30
2.3	Distribution of the average payoff in each treatment for each supergame.	32
2.4	Real value of average payoffs.	36
2.5	Frequency of change in price rank in each supergame, divided by the number of price change opportunities.	37
3.1	Best response functions for Day and Night demand levels, duopoly.	49
3.2	Example screenshot for the continuous treatment.	54
3.3	Average price over the experiment. Average price per second (top). Average period price, Day and Night (bottom).	56
3.4	Prices in each market.	57
3.5	Fraction of time (seconds) played by each group in each pricing category.	59
3.6	Fraction of markets within each category in each period.	61
3.7	Average period price for Unstable/Low markets.	62
4.1	Screenshot of Short-Discrete treatment screen shown to participants.	69
4.2	The top figure shows the average period price for each treatment (dashed lines), and separately for day and night in <i>short-discrete</i> (solid green lines). The bottom figure shows the difference in average period price between treatments.	71
4.3	Average market period price in each category.	72
4.4	Prices in each market.	73
4.5	Fraction of markets in each category over time.	74
4.6	Distribution of average period prices in markets that are Unstable/Low.	75
4.7	Average period price for markets that are Unstable/Low, over time.	75

Abstract

This thesis examines how pricing frequency impacts the ability of firms to tacitly collude. Online platforms have reduced search costs for consumers by making it easier to compare across firms. However, they also make it easier for firms to use prices as a communication tool and keep track of competitors. Various jurisdictions have implemented transparency schemes in the retail petrol market, with the aim of promoting competition between firms and enhancing consumer welfare. However, it is unclear whether the increased competition from lower consumer search costs outweighs a potentially enhanced environment for tacit collusion. A sub-set of jurisdictions have implemented additional restrictions on how and when petrol stations can change their price. It is unclear whether these pricing-frequency restrictions lead to enhanced competition and consumer welfare. This thesis uses laboratory experiments to investigate how pricing frequency impacts the level of tacit collusion in markets.

Chapter 2 examines whether pricing frequency impacts tacit collusion in the context of infinitely repeated games. We use four treatments to distinguish whether behaviour is driven by collusion incentives, adaptive pricing, and/or satisficing. On one dimension we vary whether there one or four pricing periods per stage game. On the other, we vary the induced discount factor (firm patience). Results show that increased pricing frequency leads to lower collusion levels. This is due to adaptive pricing, whereby firms continuously undercut one another, leading to a decay of cooperation over time.

Chapter 3 investigates how pricing frequency impacts collusion levels in a 60 minute, real-time duopoly experiment. Unlike the experiment in Chapter 2, where each period is distinct, we implement pricing-frequency restrictions in an otherwise real-time environment. We compare an unregulated market, with no restrictions on how and when firms can change their price, to one where prices are only updated simultaneously at set 20 second intervals. Results show collusion levels are significantly higher without pricing-frequency restrictions. This seemingly contradicts the finding from Chapter 2 where higher pricing frequency reduces collusion levels.

To test the mechanism behind our results, in Chapter 4 we implement a third treatment where prices are updated every 2.5 seconds. If collusion incentives such as the cost of signalling and gain from deviating drive results, then average prices in the 2.5 second treatment should be closer to our unrestricted treatment. If a behavioural mechanism is responsible, prices should be closer to the 20 second treatment. Results show collusion levels are virtually identical between the 2.5 and 20 second treatments. Thus, firms price adaptively, with pricing periods acting as a behavioural prompt to start a price war. This explains the seemingly contradictory results from Chapters 2 and 3, and the distinction between decisions made in discrete and continuous time becomes salient.

This thesis contributes to the experimental literature on cooperation in repeated games and continuous time. Results show that imposing discrete decision periods in an otherwise continuous-time environment reduces cooperation levels. Future experiments implementing repeated games should therefore consider the impact of timing as part of their design.

Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship, and funding from the Australian Research Council (DP210103919).

Acknowledgements

I am deeply grateful for the exceptional mentorship, guidance, and encouragement provided by my supervisor, Ralph-C Bayer. Ralph is a wonderful role model, incredibly patient, generous, and always willing to share his vast knowledge and passion for economics – in his own research area and in economics more generally. I have learnt so much from him, and he never fails to light up the conversation with new insights during our discussions. I would also like to thank my co-supervisor Paul Pezanis-Christou for his support, encouragement, and assistance throughout the years. I feel very fortunate to have had Ralph and Paul's guidance and company throughout the many challenges and discoveries I've encountered.

I am also appreciative of guidance and support from HDR coordinators Stephanie McWhinnie, Giulio Zanella, and Raul Barreto, who have fostered a lively research environment for research candidates at the School of Economics. I am particularly grateful to Stef for her support over the years. She played an instrumental role in helping me get the job that led to the realisation I wanted to do research, and this is something I will always be grateful for. As a role model she demonstrates integrity, intelligence, and perspective, and her guidance has been invaluable. I would also like to extend thanks and appreciation to the wider faculty and support staff at the School, who I have enjoyed interacting with and learning from.

I am grateful to David Byrne and Tom Wilkening for their suggestions and feedback on Chapters 3 and 4. Their guidance was invaluable and helped in pushing my thinking forward on how to refine and strengthen the applications of my research.

I extend thanks and gratitude to Kim and Rui, who have helped me immensely with moral support and as lab mates. I also thank the friends I've made in the PhD room who have provided much needed laughter, good company, and interesting insights from their research.

This thesis would not have been possible without the love and support from Michael, Beela, Oskar, and my parents. Although, I may have finished writing sooner had I not had paws all over my face every night. I am grateful to Aneta for her unwavering support and encouragement.

I dedicate this thesis to my Nana.

Chapter 1

Introduction

Many markets have become increasingly transparent through the emergence of accurate and up-to-date online platforms. Examples of such markets include retail petrol, business hotels, and online retailers more generally. Transparency reduces search costs for consumers by making it easier to compare prices across sellers, but has a secondary effect of facilitating the monitoring of competitors' prices. Competition laws prevent explicit agreements to collude being made between firms, but do not prevent tacit communication through prices. The question arises as to whether increases in competition from reduced consumer search costs outweigh the potential increase in tacit collusion from the reduction in communication costs. Answering this question empirically is not straight forward. Concerns about tacit collusion in retail petrol (gasoline) markets have resulted in the implementation and enforcement of competition regulations in various jurisdictions around the world. These generally take the form of transparency programs and price setting restrictions, such as price ceilings, limitations on the frequency of price changes, or only allowing for price reductions (Haucap and Müller, 2012). How can these pricing restrictions help or hinder the ability of firms to tacitly collude in otherwise high-information environments?

The motivating policy behind this thesis is the retail petrol regulations in West-

ern Australia (WA). These regulations impose 24-hour price commitment on petrol stations, and stations are required to submit their price for the next day by 2pm. The following day's prices are then made publicly available from 2:30pm through the FuelWatch program.¹ This price is then implemented at 6am the next day, and remains fixed for the next 24-hour period. It is unclear if pricing frequency regulations have actually led to welfare increases for motorists, due to the difficulty of controlling for multiple unobservables in these markets.

Our initial question can be restated to ask if simultaneous pricing-frequency restrictions can reduce tacit collusion in markets. This thesis comprises three related essays that seek to experimentally examine if changes to allowable pricing frequency, whether through regulations or technology, impacts the ability of firms to tacitly collude. Can theory help us answer this question? The modelling of strategic competition in quantities and prices by Cournot (1838) and Bertrand (1883) paved the way for rigorous theoretical investigation of market characteristics and welfare outcomes (see Stigler (1964)). In particular, theoretical work on infinitely repeated games and associated folk theorems for identifying equilibria (i.e., Friedman (1971); Fudenberg and Maskin (1986)) shed light on whether collusive equilibria in markets may be implementable. In complementary work, Maskin and Tirole (1988) provide the theoretical underpinning for the impact of price commitment in infinite discrete time duopolies, proving that Edgeworth cycles can be a Markov perfect equilibria. Thus, under certain assumptions such as sufficient firm patience and timing of actions, theory tells us that collusive outcomes are indeed implementable. However, it does not tell us how they are achieved, nor provide a sharp outcome on how collusive a market may be. The theory of infinitely repeated games is unable to answer our question.

The retail petrol market has received attention from empiricists, given the different policies implemented by regulators and the availability of data on petrol

¹The retail petrol market in WA is regulated under the Petroleum Products Pricing Act 1983 WA.

prices (Eckert, 2013). Byrne and de Roos (2019); Wang (2009) provide evidence of tacit collusion in the WA petrol market after the introduction of the 24-hour price commitment regulations. Despite evidence of tacit collusion, it remains unclear whether there is *less* collusion in WA with the 24-hour price commitment regulations. One might think of comparing prices in WA with those in New South Wales (NSW), which also has a pricing transparency program, FuelCheck, but no restrictions on the frequency of allowable price changes.² Any price differential for retail petrol in WA and NSW cannot be entirely attributed to the difference in regulatory environment. The complexities of markets are vast, including differing consumer preferences, services provided, capacity constraints and brand-group pricing issues. Comparing prices, even with the inclusion of controls, does not adequately solve the identification problem.

Dewenter et al. (2017) and Becker et al. (2021) have made progress towards identifying the causal impact of pricing regulations in the Austrian retail petrol market using difference-in-difference and synthetic control methods, respectively. Both studies find a pro-competitive effect of the Austrian rule, which allows for one price increase per day and unlimited decreases. Dewenter et al. (2017) do not find any significant effect of the WA regulations on fuel price levels. More evidence is needed, as both studies formulate their controls using prices from other countries, where the underlying data generating process may be different.

Laboratory experiments are well placed to cleanly test the impact of pricing frequency regulations on tacit collusion. This thesis experimentally investigates the impact of price setting rules that restrict interaction frequency on cooperation in Bertrand oligopolies. All experiments were conducted with approval from the Human Research Ethics Committee at the University of Adelaide (H-2019-179).

In Chapter 2 we investigate how price commitment impacts cooperation in the context of an infinitely repeated game. We bridge the gap between traditional

²The NSW industry is regulated under the Fair Trading Act 1987 No 68 NSW.

experiments on infinitely repeated games and the more recent development of experiments investigating how behavioural factors impact cooperation levels. In our theory-driven null hypothesis, we test whether increasing the attractiveness of collusion, as measured through collusion incentives,³ impacts the level of collusion in markets. Our first alternative hypothesis is derived from the observed decay in cooperation levels in experiments using a repeated prisoner's dilemma (Chaudhuri, 2011). Alternative hypothesis 1 states that average levels of collusion will decrease as the number of price changes increases. The second alternative hypothesis states that as discount factors decrease, the level of collusion increases. This is contrary to standard repeated-game logic on the impact of discount factors, as tested in our null hypothesis. Instead, alternative hypothesis 2 tests whether collusion is driven by satisficing behaviour (Simon, 1956). Participants target a real payoff level and thus compensate for the lower discount factor by increasing nominal prices and profits. The third alternative hypothesis is that both the number of price changes, and the discount factor, impact observed collusion levels.

Four treatments are used, varying the number of price changes per period to be one or 4, and the induced discount factor.⁴ Using theory and observations of behaviour in experiments, we derive a null and three alternative hypotheses. We run five supergames in each session for our aggregative triopolies. Results support alternative hypothesis 1, and collusion levels are significantly higher in treatments with one price change compared to treatments with 4 price changes per stage game. We do not find evidence to support our null hypothesis or alternative hypothesis 2. Thus we conclude that restricting the number of price changes results in higher levels of collusion. Adaptive pricing, where firms undercut one another, takes longer to achieve with fewer available price changes.

³We derive collusion incentives from the theory of infinitely repeated games. Based on our parameters we calculate the value of colluding at the joint profit maximising price, minus the value of one-shot deviation with Nash profits forever (grim trigger). As this difference decreases, we expect lower levels of collusion, as it becomes less attractive.

⁴We use a combination of a continuation probability and inflation rate. For our purposes, using a continuation probability only would not allow us to achieve the induced discount factor low enough for our purposes without severely restricting the length of supergames.

In Chapter 3, we test the impact of restricting pricing frequency in a 60 minute (90 period), continuous-time duopoly. Following the findings in Chapter 2, we are interested in how collusion dynamics evolve over time, without being restricted by the stopping probability or inflation rate required to implement an infinitely repeated game in the lab. Profits are gained in a flow over time, which more closely follows the flow of consumers stopping for petrol over each day. Another key difference with our experiment in Chapter 2 is that we remove the break between periods to ensure that any restart effects, or chosen focal points for coordination, are due to differences between treatments, and not due to the experimental design. Each cycle (or period) runs for 40 seconds, with a repeating pattern of 20 second days and 20 second nights. These shocks implement high and low demand levels at known intervals at the same time for all participants. Firms need to actively change prices to adapt to each demand level, ensuring participants are active in pricing for the entire 60 minute experiment.

We compare a treatment where firms are unrestricted in the number of price changes, to a second treatment that implements 20 second pricing periods for an otherwise real-time market. Our null hypothesis states that price levels are equal in both treatments. The alternative hypothesis depends on the inequality between treatments. If prices are lower in the unrestricted treatment, the introduction of discrete pricing periods every 20 seconds makes it easier for firms to collude (i.e., introduces a clear focal point and makes it easier to keep track of prices). If prices are lower in the 20 second treatment, higher cost of initiation and gain from deviating result in lower levels of collusion. Results show a clear treatment difference, with higher prices in the unrestricted treatment. This is seemingly contrary to the results in Chapter 2, where more available price changes led to lower levels of collusion.

In Chapter 4, we test for the mechanism driving the results in Chapter 3 by using an additional treatment with 2.5 second pricing periods. If prices in this treatment are closer to the unrestricted treatment in Chapter 3, then we confirm that collusion incentives lead to higher prices. This is because the cost of signalling

(initiation) and gain from deviating (incentive to cheat) are lower in the unrestricted treatment. Thus, by reducing the simultaneous price commitment period from 20 seconds down to 2.5 seconds, these collusion incentives are closer to those in the unrestricted treatment. However, a behavioural explanation could be responsible for the difference, thus the 2.5 second length is chosen to ensure the decision period is still discretised.

Prices in the 2.5 second treatment are similar to those in the 20 second treatment. Results show a higher proportion of fully collusive markets, and higher prices more generally, in the unrestricted treatment. Thus, we conclude that collusion incentives are not behind the higher prices in our unrestricted treatment, and instead it is the introduction of discrete pricing periods in an otherwise continuous time market. We hypothesise this is driven by a status quo effect (Samuelson and Zeckhauser, 1988) in the unrestricted treatment, whereas the introduction of structure into the market through a 20 or 2.5 second pricing period acts as a prompt for firms to re-strategise. This has a pro-competitive effect on markets.

The final chapter concludes by summarising the findings from this thesis in context of the literature and future research directions.

Chapter 2

Infinitely repeated games

2.1 Introduction

“The best known results in the theory of repeated games, the folk theorems, focus attention on the multiplicity of equilibria in such games, a source of consternation for some. We consider multiple equilibria a virtue - how else can one hope to explain the richness of behavior that we observe around us?” - Mailath and Samuelson (2006)

The development of repeated game theory opened the door to the analysis of how time impacts the nature and outcome of strategic interaction. The classical prisoner’s dilemma is one such example, where repeated interaction can lead to the possibility of a cooperative outcome being reached when it would otherwise be unattainable as a sub-game perfect equilibrium. Various Folk Theorems for infinitely repeated games show how punishment strategies can be used to support cooperative outcomes, providing a theoretical explanation for why cooperative outcomes are often observed in society (Friedman, 1971; Fudenberg and Maskin, 1986). For many real-world situations, achieving a cooperative outcome improves the overall welfare of society, e.g., action on climate change, students working on a group assignment, and attempts by nations to promote regional stability. However, there are circumstances in which the welfare implications of cooperation are negative from the policy maker’s point of view. Cooperation between firms in oligopoly markets can lead to

diminished consumer welfare. This chapter is concerned with how the latter is impacted by the frequency of interaction within a stage game. We focus on games of complete information to reflect real world markets where prices are readily available online. This allows us to focus on identifying the causal impact of interaction frequency, without confounds that may result from imperfect monitoring.

Friedman (1971) and Fudenberg and Maskin (1986) show that for a supergame, or an infinitely repeated stage game, all individually rational stage-game payoffs that are feasible can be implemented as an average period payoff - as long as players are sufficiently patient. Patience is measured against the threshold critical discount factor, which in turn is determined using the payoffs of the game and strategies used by players. Thus, there are infinitely many potential sub-game perfect equilibria that are Pareto dominant compared to playing Nash at every stage.¹ With a large set of potential equilibria in a supergame, including Nash, it is unclear on which equilibrium players will eventually land. If cooperation is indeed sustainable in equilibrium, identifying which is more likely to occur is not straight forward.

This theoretical gap creates challenges for applications where more refined predictions are required. Although, as Mailath and Samuelson (2006) state, multiplicity of equilibria provides an explanation for varied behaviour, it can also hinder analysis of policy. Theory does not predict which equilibrium will actually be played under different policy regimes. For instance, in situations where cooperation leads to concern over anti-competitive conduct in markets, policymakers are without theoretical guidance on how potential market regulations may affect equilibrium selection. Potential solutions to this coordination problem have been proposed, such as the identification of focal points (Schelling, 1960), measures of strategic uncertainty (Heinemann et al., 2009), and payoff and risk dominance (Harsanyi, 1995; Harsanyi and Selton, 1988). Unfortunately these approaches do not make predictions on how interaction frequency impacts the overall level of cooperation in markets. A secondary concern

¹This requires that there are feasible stage game payoffs that pareto-dominate stage-game Nash.

arises as to whether there are behavioural factors that impact cooperation levels, which theory is currently unable to capture.

The lack of predictive power of theory, and the potential impact of behavioural factors, has led to experimental work on cooperation in infinitely repeated games. Infinitely repeated games can be implemented in laboratory experiments by incorporating a probability that the game ends each period (stopping probability) and an inflation rate. Dal Bó and Fréchette (2018) provide an extensive meta analysis of how characteristics of repeated games impact the overall level of cooperation. The main vehicle to study infinitely repeated games and cooperation has been the prisoner's dilemma. Dal Bó and Fréchette (2018) use a meta analysis of experiments using an infinitely repeated prisoner's dilemma to show how cooperation levels are impacted by variables including the critical discount factor, game length and experience, matching protocol, and public monitoring. However, it is not clear that repeated-game logic is the only explanation for observed results on differences in cooperation levels. The strategy space in a prisoner's dilemma is not rich enough to conclusively prove the existence of repeated game play according to theory, and separate this from potential behavioural factors such as social preferences, reasoning ability, and the building of reputations. Any defection from observed cooperation is indistinguishable from the implementation of a punishment strategy. Thus it is unclear whether punishment, or another behavioural influence, is behind the observed change in cooperation.

A relatively new strand of literature has made progress in this area by showing differences in decisions made in discrete compared to continuous time. In the context of a 60 second continuous prisoner's dilemma, Friedman and Oprea (2012) find that continuous time helps stabilise cooperation, as do Bigoni et al. (2015). Leng et al. (2018) find that games in continuous time are behaviourally different to those in discrete time, in the context of a minimum effort coordination game.

Our contribution seeks to bridge the gap between the experimental literature on

infinitely repeated games and the literature on the timing of decisions and cooperation. We identify potential behavioural explanations for observed differences in cooperation levels, while embedding the decision environment within the context of an infinitely repeated game. This allows us to disentangle the impact of collusion incentives, driven by repeated-game logic, from potential behavioural factors. The context of Bertrand competition allows us to explore these impacts over a wider strategy space than if we used a 2 by 2 game. We focus on simultaneous pricing decisions in markets with complete information, capturing the essence of the retail petrol market regulations in WA. The classical Bertrand model of oligopoly (Bertrand, 1883) is appropriate to proxy the general retail petrol market, where firms sell otherwise homogeneous products that are differentiated by location and services provided. Although a Hotelling location model (Hotelling, 1929) could also capture these characteristics, we are more interested in the choice of price, and impact on price dynamics over time.

To understand how pricing frequency regulations impact collusion levels, we need to identify how they impact the stage game in an infinitely repeated sequence. Pricing frequency impacts the potential payoffs to players by dividing each period payoff by the number of sub-periods (or pricing periods). This necessarily impacts the payoffs used to calculate the critical discount factor. For example, with more pricing periods in a stage game, the payoff from deviating is lower as other firms can punish sooner, resulting in a lower critical discount factor. In other words, the threshold level of patience for cooperation to be sustained in equilibrium is reduced. We again run into the same problem, that in environments where collusion is sustainable, it is unclear how pricing frequency might impact overall collusion levels. Although theory tells us how pricing frequency impacts the critical discount factor, above this threshold it is unclear whether there is any effect. This further demonstrates why empirical investigation of potential behavioural factors is necessary. Regarding equilibrium selection, our contribution is on the question of how regulations impact

average collusion levels, rather than understanding how collusion is initiated.²

The level of control in laboratory experiments makes them an appropriate tool to help us understand how cooperation is impacted by different characteristics of repeated games (Abbink and Brandts, 2008). Within the experimental literature, multiple determinants of cooperation have been explored in the context of repeated games, and industrial organisation more generally (Bigoni et al., 2015; Friedman and Oprea, 2012; Dal Bó and Fréchette, 2011; Leng et al., 2018). Factors impacting cooperation levels include:

1. The number of firms in the market. As the number of firms in experimental markets increases above two, collusion levels decrease (Huck et al., 2004; Dufwenberg and Gneezy, 2000; Abbink and Brandts, 2008);
2. Use of a stochastic vs deterministic ending rule. Bigoni et al. (2015) find more cooperation with deterministic ending rules, rather than stochastic ones;
3. The sustainability of collusion as an equilibrium outcome for infinitely repeated games. Being supported in equilibrium is not a sufficient condition for cooperative outcomes to arise (Dal Bó and Fréchette, 2011; Dal Bó and Fréchette, 2018);
4. Strategic complements vs strategic substitutes. More cooperation is observed when strategies are complementary (i.e., price competition) compared to when they are substitutes (i.e., quantity competition) (Potters and Suetens, 2009; Anderson et al., 2010);
5. Information and experience. The availability, or lack thereof, of payoff information and ability of participants to learn can support cooperation in the long run (Huck et al., 2017; Friedman et al., 2015). Cooperation increases with experience (Dal Bó and Fréchette, 2011), as with explicit communication (Fonseca and Normann, 2012);

²Using the terminology of Green et al. (2014), we are able to provide insight on the implementation stage related to equilibrium selection, but not the initiation stage.

6. Alternate vs simultaneous move games (Leufkens and Peeters, 2011).

Thus, the literature has explored multiple determinants of cooperation in repeated games. However, the impact of pricing frequency has not been addressed. Our experiments do not seek to replicate these findings, rather complement them by providing evidence on how a potentially crucial element impacts cooperation. The experiments of (Leufkens and Peeters, 2011) are the closest to achieving this by examining short-run price commitment, however the focus is whether strategic uncertainty resulting from a simultaneous move game leads to more collusion than an alternate move structure. In addition, more work is required in the industrial organisation context of repeated games to understand the impact of decision timing, or frequency.

We experimentally test the impact of pricing frequency on collusion levels in the context of an infinitely repeated game. Referencing the characteristics from the literature, we implement a market with three firms simultaneously setting prices (strategic complements), and a stochastic ending rule. Our design is novel as we incorporate flow profits into the stage game. We use four treatments to differentiate between the impact of general collusion incentives (derived through the repeated game itself), and the potential behavioural impact of higher pricing frequency. Our experimental market comprises a symmetric aggregative triopoly, where firms compete in prices selling differentiated products. The choice of prices by firms are strategic complements, and thus prices up to the monopoly price lead to higher profits compared to Nash.

Our experiments are not designed to strictly test theory. Theoretically where collusion is not sustainable in equilibrium, the Nash outcome is expected. If collusion is sustainable, an average price above Nash, and the Nash outcome, is feasible. Thus, finding no difference between treatments is a valid theoretically. Instead we formulate a null hypothesis which states that collusion increases with classic collusion incentives. This refers to the difference in colluding forever, and defecting and being

punished indefinitely under a grim trigger strategy. Even if collusion is a feasible equilibrium action in two treatments, where the incentives are higher we may expect higher levels of collusion.

Our first alternative hypothesis is that more price changes per periods leads to lower levels of collusion. The intuition behind this hypothesis is from the observed decay in contributions over time in repeated public goods games, with repetition leading to lower levels of cooperation (Chaudhuri, 2011). This observation is in the spirit of Maskin and Tirole (1988) Edgeworth price cycles, which are commonly observed in retail petrol markets. If participants are pricing adaptively, cooperation levels will decrease with each new decision. This would appear as a price cycle in an industrial organisation context, with a restart effect at the start of the game. If participants are pricing adaptively, then the rigidity imposed in the marketplace slows the pace of price wars, leading to higher average prices under lower pricing frequency. Fewer available price changes may stifle competition and lead to higher prices on average. Thus, a longer period of time will be required to achieve the same number of price reductions.

An alternative behavioural hypothesis is that pricing frequency does not impact collusion levels, but that discount factors do in a non-standard way. Standard theory would suggest that lower discount factors would reduce cooperation rates.³ If instead participants aim for a target level of real profit, then over time to compensate for a lower discount factor, prices and therefore nominal profits will rise. Such behaviour could otherwise be described as satisficing (Simon, 1956), and has been identified as a factor driving collusive outcomes in markets (Dixon, 2000). Differences in collusion would then be observed between discount factors, but not across pricing frequencies. Finally, it could be that both of these behaviours impact collusion. This would lead to a distinct ranking of collusion levels for our four treatments.

³Compared to a critical discount factor, if the agent's patience level is below this, then collusion is not sustainable. Standard collusion incentives would suggest that the more patient the agent is, the higher the incentive to maintain cooperation.

On one dimension we vary the number of allowable price changes per period, to be either one or four. On the other, we vary the induced discount factor (patience level) using a continuation probability and inflation rate. The use of an inflation rate on top of the continuation probability is necessary to achieve discount factors low enough for our purpose, without making supergames too short. To test the impact of pricing frequency on cooperation we need to hold the incentives derived from repeated games constant, in case these incentives do play a role in determining cooperation levels. To achieve this we select one of the inflation rates such that the difference in continuation payoff from colluding and deviating is equalised between a treatment with 1 price change per period, and another with 4. The average supergame payoff is used as a measure of collusiveness as it allows us to easily compare across treatments with different discount factors.⁴

We find no evidence to support our null hypothesis that higher collusion incentives lead to higher levels of collusion. Results support our first alternative hypothesis - with more price changes per period collusion levels are lower due to within-period competition. Lower discount factors leading to higher collusion levels due to satisficing behaviour is not supported by the data. We are able to show, in the context of an infinitely repeated game, that higher action frequency impacts cooperation levels by promoting adaptive pricing.

The remainder of the chapter is structured as follows. First, we outline our market in the context of the stage game. We then show how pricing frequency within a stage game impacts the critical discount factor, and the theoretical uncertainty in how pricing frequency impacts tacit collusion. Second, we outline our experimental design, implementation and hypothesis. Third, we present results. Finally, we conclude.

⁴This is not the standard calculation of the average. The average period payoff is calculated as the stage game payoff that would yield the same supergame payoff derived from the realised payoffs of the supergame. As we don't have an actual infinite sequence, we only use realised payoffs to make this calculation. The calculation is explained in more detail later in this chapter.

2.2 Theoretical background

We are interested in how the number of allowable price changes in a market impacts average prices. In other words, how does the frequency of interaction impact the ability of firms to coordinate in an infinitely repeated game? repeated-game logic sheds light on what conditions are necessary for collusion to be a theoretical possibility. However, this does not guarantee that players will successfully coordinate, nor on which strategy (price).

First we outline our stage game model. We then extend this to an infinitely repeated game to examine how pricing frequency per period impacts the ability of firms to coordinate. Finally, we explain how this leads to the theoretical uncertainty that underpins our empirical investigation.

2.2.1 Stage game

In this section we describe our general model and the characteristics of our market. With experimental implementation in mind, our model seeks to simplify the game, while working within the constraints of implementation.⁵ We start with a simultaneous move, symmetric Bertrand oligopoly with N firms $I = \{1, 2, \dots, N\}$ selling differentiated goods. There is a discrete and finite strategy space, so firms are able to set prices at integers $p_i \in \{1, 2, \dots, p^{max}\} \forall i \in I$. Firms are given a payoff matrix which was designed around the implementation constraints. This was initially based on the standard linear differentiated Bertrand model. Payoffs are chosen to satisfy the properties usually assumed for symmetric differentiated Bertrand oligopolies, as follows.

⁵Constraints included the implementation of the discount factor (discussed further below), including a reasonable continuation probability, and avoiding extreme expected payment differences between treatments.

Symmetry

Profit for each firm depends on their own price p_i , and the other firms' prices p_{-i} , and can be written as $\pi_i(p_i, p_{-i})$. To simplify the model for experimental implementation, we assume that a firm's profit only depends on their own price, and the average of other firms' in the market. Profits are symmetric, such that:

$$\pi_i(p_i, p_{-i}) = \pi_j(p_j, p_{-j}) \text{ if } p_i = p_j; p_{-i} = p_{-j} \quad (2.1)$$

Existence of best responses

Given the discrete strategy space, this implies that a best response always exists such that:

$$\exists p_i^*(p_{-i}) = \arg \max_{p_i} \pi_i(p_i, p_{-i}) \quad \forall p_{-i}, i, -i \quad (2.2)$$

Aggregative

In a market with more than two firms, characterising the oligopoly game in two dimensions simplifies the game significantly, which is particularly important for the purpose of experimental implementation in the lab. Thus, we characterise the game as aggregative, such that each firm's profit function depends only on own price p_i , and the aggregate price of other firms in the market. Our aggregator function takes the average of other firm prices in the market, denoted \bar{p} , and assume a unique Nash equilibrium.⁶ We state this as follows:

$$\pi_i(p_i, p_{-i}) = \pi_i(p_i, \bar{p}_{-i}) \text{ if } \sum_{k \neq i} p_k = \sum_{k \neq i} \bar{p}_k \quad (2.3)$$

⁶The theoretical literature on aggregative games provides analysis well beyond our present use, including existence of equilibrium and welfare analysis in asymmetric oligopolies (Anderson et al., 2020; Jensen, 2010; Alós-Ferrer and Ania, 2005; Corchón, 1994; Caplin and Nalebuff, 1991).

In other words, firm i 's profit, for a given price, is the same for price vectors of the other firms in the market that yield the same average price. Therefore the firm's profit function can be written as a function of own price, and the aggregate (average) of the other firms in the market:

$$\pi_i(p_i, p_{-i}) \Rightarrow \pi(p_i, \bar{p}_{-i}) \quad (2.4)$$

Strategic complements

The products being sold by firms are imperfect substitutes, such that the cross-price elasticity of demand between them is positive. For a given p_i , if \bar{p}_{-i} increases, the quantity demanded for firm i increases too. With respect to the retail petrol market, this captures the differentiation between petrol stations due to location and/or customer loyalty programs, for products which are otherwise identical. The final important characteristic of our market is that firms' strategies are strategic complements, such that the best response is weakly increasing in response to an increase in \bar{p}_{-i} :

$$p_i^*(\bar{p}_{-i}) \geq p_i^*(\bar{p}'_{-i}) \text{ if } \bar{p}_{-i} > \bar{p}'_{-i} \quad (2.5)$$

We can now define the stage game prices and profits. The unique and symmetric Nash equilibrium price $p^N = p_i^*(\bar{p}_{-i}^N)$ leads to Nash profits:

$$\pi^N := \pi(p^N, p^N) \quad (2.6)$$

The joint profit maximising price, p^C , is the price that satisfies the following:

$$p^C = \arg \max_p \sum_{i=1}^N \pi_i(p, p) \quad (2.7)$$

Thus, the fully collusive profits of each firm are:

$$\pi^C := \pi(p^C, p^C) \quad (2.8)$$

Finally, the deviation profit from the stage game occurs with price:

$$p^D = \arg \max_p \pi_i(p, p^C) \quad (2.9)$$

Deviation profits are therefore:

$$\pi^D := \pi(p^D, p^C) \quad (2.10)$$

From our assumptions it follows that:

$$p^N < p^D < p^C \quad (2.11)$$

and

$$\pi^D > \pi^C > \pi^N \quad (2.12)$$

The above describes the general characteristics of our stage game. The following section extends this to an infinitely repeated game, showing the impact of the frequency of price changes.

2.2.2 Pricing frequency and collusion

By extending the stage game to an infinitely repeated game, or supergame, we can examine how the frequency of price changes impacts the ability for firms to co-

ordinate. The theoretical literature on infinitely repeated games is vast, but our research question is concerned with the implications of one particular folk theorem for equilibrium in infinitely repeated games, as articulated by Fudenberg and Maskin (1986) and Friedman (1971). To briefly summarise, this folk theorem tells us that as long as firms are sufficiently patient, it is possible for payoffs (averaged across the supergame) that Pareto dominate the stage game Nash equilibrium to be implemented as a subgame perfect Nash equilibrium (SPNE) of the supergame. What is considered sufficiently patient can be determined through the critical discount factor $\bar{\delta}$.

The stage game is now referred to as a period, and within-stage game price changes are referred to as sub-periods. To introduce time, we allow n price changes per stage game (number of sub-periods), discount factor (inflation) r , and continuation probability ρ . The choice to include both a discount factor and a continuation probability was made to facilitate experimental implementation. Firm i 's discount factor δ_i can be stated as:

$$\delta_i = \rho(1 - r) \quad \forall j \quad (2.13)$$

Grim trigger, reverting to playing Nash forever if anyone deviates from the collusive agreement, is the most extreme trigger/punishment strategy. Grim trigger may support the collusive outcome if firms are patient enough. Less extreme strategies, for example where punishment only lasts as long as necessary to eliminate deviation profits, are also possible. However, grim trigger gives us the upper threshold of required firm patience in the supergame, thus we focus on this.

We can determine the threshold patience level that allows for the collusive outcome. A firm will not deviate if the present value of profit from cooperating is greater than the profit from deviating today, plus the present value of Nash profits forever. This can be written as:

$$\frac{\pi^C}{1 - \delta_i} \geq \pi^D + \frac{\delta_i \pi^N}{1 - \delta_i} \quad (2.14)$$

In each period, if a firm deviates from the collusive agreement, the grim trigger strategy would require firms to respond by playing p^N forever. Considering the impact of the frequency of price changes, it can be observed that for $n > 1$, if a competitor deviates from the collusive agreement, the punishment will take effect in the subsequent sub-period. Therefore, a firm that wishes to deviate will still want to maximise its period profit, and play the collusive price for $n - 1$ sub-periods, then deviate in the last sub-period. The period profit from this deviation strategy is defined as:

$$\pi^D = \frac{(n - 1)\pi^C + \pi^D}{n} \quad (2.15)$$

Knowing this, we can then determine the threshold firm patience, $\bar{\delta}_n$ in order for collusion under grim trigger strategy to be sustainable. It is defined as below:

$$\bar{\delta}_n = \frac{\pi^D - \frac{\pi^C}{(1+r)}}{\pi^D - \frac{\pi^N}{(1+r)}} \quad (2.16)$$

Notice that as n increases, π^D decreases, therefore $\bar{\delta}_n$ decreases.⁷ It is theoretically more conducive to collusion (required firm patience lower) as the frequency of price changes increases. The sub-period payoff from deviating decreases as n increases. The present value of cooperating, Π^C for a given level of firm patience δ_i for the entire supergame can be written as:

$$\Pi^C = \frac{\pi^C}{(1 - \delta_i)} \quad (2.17)$$

⁷Differentiating (2.16) with respect to n recalling the π^D depends on n and that $\partial\pi^D/\partial n < 0$ immediately reveals this.




Figure 2.1: Possible equilibria with two firms.

Similarly, the present value of deviating today, plus the discounted value of Nash profits for each subsequent period is:

$$\Pi^D = \pi^D + \frac{\delta_i \pi^N}{(1 - \delta_i)} \quad (2.18)$$

Thus, firms will cooperate so long as $\Pi^C \geq \Pi^D$.

A folk theorem tells us that any individually rational equilibrium is implementable as a SPNE (Fudenberg and Maskin, 1986). For the two firm example, all possible equilibria are shown in the shaded portion of Figure 2.1. In other words, any present value of payoffs, π_i and π_j is a potential SPNE of the supergame. It is possible for the threat of grim trigger to support the collusive outcome if firms are sufficiently patient, such that the firm's discount factor greater than the critical discount factor $\delta_i \geq \bar{\delta}$. This leads to our empirical question of how interaction frequency impacts the ability of firms to coordinate on these equilibrium. The nature of firm interaction in our model, strategies being symmetric strategic complements, means that symmetric profits, starting from Nash profits, are a Pareto improvement on the last. Thus firms have an incentive to coordinate on prices as close to the joint-profit maximising price as they can.

2.3 Experimental design

In this section we explain experimental design. First we outline the intuition behind the experimental design, and then we detail the parameters and implementation. We implement a triopoly (three firms) to balance the overall competitiveness of an experimental market, yet still capture the nature of the real-world markets where pricing frequency regulations have been implemented (more than two firms competing).

We conduct four treatments, varying induced firm patience δ_r (above or below the threshold $\bar{\delta}_n$) and the number of allowable price changes per period (n). This allows us to isolate the impact of allowable pricing frequency on average price.

For a given n , varying whether collusion is sustainable or not will allow us to identify a baseline average price level and behaviour of firms. Our baseline treatment sets $n = 1$, with induced firm patience below $\bar{\delta}_1$, such that collusion is not sustainable. We label this low firm patience treatment $n1\delta_l$, which, as explained below, is implemented using a high inflation rate.

In the remaining three treatments we induce firm patience above $\bar{\delta}_n$ so that collusion is sustainable in each supergame. This allows us to distinguish between differences in the incentive to deviate, and the number of price changes, between treatments. Our second treatment, $n1\delta_h$, holds n constant and induces firm patience above the threshold required to sustain collusion by lowering the inflation rate. We set $n=4$ for our third treatment $n4\delta_h$. Increasing n by necessity decreases the continuation payoff from deviating. This means that for a given level of firm patience, changing n will result in two differences between treatments - the number of allowable price changes; and the payoff from deviating. We therefore run our fourth treatment, $n4\delta_l$, where induced firm patience is selected such that the continuation payoff from deviating is the same as in $n1\delta_h$. This will allow us to identify the impact of altering the number of price changes from one to four.

Table 2.1: Treatments and summary of parameters.

		$n = 1$	$n = 4$
	Inflation rate	$(\bar{\delta}_1 = 0.63)$	$(\bar{\delta}_4 = 0.30)$
$\delta_h = 0.72$	0.1	$n1\delta_h^*$	$n4\delta_h^*$
$\delta_l = 0.56$	0.29	$n1\delta_l$	$n4\delta_l^*$

$$p \in [1, 7] \quad p^N = 2, \quad p^C = 6$$

Continuation probability = 0.8;

*Collusion sustainable in equilibrium.

As above, increasing the number of price changes decreases the required firm patience, $\bar{\delta}$ for grim trigger to sustain collusion. We vary whether grim trigger is able to theoretically sustain collusion. This is achieved using a combination of a continuation probability and inflation rate (implementing the discount factor) that is varied to induced firm patience. The continuation probability for all treatments is 0.8. In other words, the probability that the game will end after each period is 20%. We vary the induced firm patience δ_i by altering the inflation rate which reduces the value of profits earned in future periods. Using an inflation rate of 10%, induced firm patience under the low inflation treatment is $\delta_h = 0.72$. The high rate of inflation is selected to make the payoff from deviating when $n = 4$ equal to the continuation payoff from deviating in $n = 1$ with low inflation.⁸ Thus we get an inflation rate of 29% and $\delta_l = 0.56$. The four treatments are summarised in Table 2.1.

The experiment was conducted at the Adelaide Laboratory for Economics Experiments using z-Tree (Fischbacher, 2007) in 2020 and 2021. Participants were recruited using ORSEE (Greiner, 2015) and were mainly students (undergraduate and postgraduate) of the University of Adelaide. For each treatment we ran four sessions with between 12 and 18 participants in each. Overall, 279 participants participated in our experiment. Once the instructions were handed out to participants,

⁸The payoff from deviating is calculated as the difference between the present value of cooperating, minus the present value of deviating (assuming grim trigger is implemented). Thus, given the low inflation rate chosen, we can select a high rate such that the difference is the same for both treatments, despite changing n . For our parameterisation, the present value of cooperating is 153.6 ECU for $n1$, while the deviating yields 144.9 ECU (difference = 8.71). Thus, we select the inflation rate that yields the same difference in payoffs for $n4$, resulting in the present value of cooperating forever being 98.7 ECU and deviating 90.0 ECU for $n4\delta_l$.

the experimenter (the same person for every session) read the first page of the instructions which provided general information about the task. Participants were then allowed to read the remainder at their own pace, and answer the set of control questions that tested their understanding. The control questions were designed to ensure participants understood the game's timing, how to use the profit table, the impact of inflation, and the method of payment.⁹ A practice round (with computers) was run prior to the commencement of the experiment. Instructions, including the profit table, can be found in Appendix A.1. Participants played five supergames in each session. We pre-drew the number of periods in each supergame to facilitate comparability between treatments. We wanted to avoid any bias caused by changing the order of supergames with differing lengths. Each supergame lasted for 4, 8, 5, 10 and 3 periods, respectively. We used matching groups of six, and participants were randomly matched into groups of three at the start of each supergame.

A simplified profit table was used, and can be seen in the example instructions in the appendix, with possible prices the integers $\in [1, 7]$. Integers between 1 and 7 were chosen to make it easier for participants to calculate the average price of the other two firms in the market. In addition, given the profit function, the difference in profit (ECU) for decimal places between these integers is minimal, as can be seen in the profit table found in the instructions. The Nash price is 2, the joint-profit maximising price is 6, and the deviation price is 3. With groups of three, this meant that each participant's profit depended upon their own price, and the average price of the other two members of their group. Participants were paid for the last period of a randomly selected supergame (Chandrasekhar and Xandri, 2014). On average, subjects earned AUD \$26.3 , including a AUD \$10 show-up payment. The

⁹All participants needed to successfully answer the following questions before the practice round began. Participants in the $n1$ treatments were asked "How many seconds does each day last?", while those in the $n4$ treatments were asked how many periods were in a day, and how long each period lasted for. Next, participants were required to identify profit in ECU for themselves, Seller 2 and Seller 3, given a hypothetical set of prices (this was the same for all treatments). Next, they were told they had profit of 20 ECU on Day 3, and needed to determine how much this was worth in AUD, which differed across δ . Finally, participants were given a table which identified the profit (in AUD) made each day for a hypothetical game. Participants needed to determine the final payment if Game 3 was chosen for payment.

experiment, including the reading of instructions, control questions, and payment, took one and a half hours to complete.

2.3.1 Hypothesis

We explain how our experimental design allows us to test hypotheses on how the frequency of price changes impacts tacit collusion in the context of infinitely repeated games. The latter requires deeper examination of repeated game characteristics, such as the discount factor and continuation profit, in order to identify the causal relationship between the former.

In Section 2.2 we outlined the stage game and how the critical discount factor is determined. To test our hypotheses it is necessary to have a payoff measure that is comparable across treatments, given the different induced values for δ_i using a stopping probability and an inflation rate i . Thus, rather than using the average price in each treatment, we instead test for differences using the average supergame payoffs. The average supergame payoff for an infinitely sequence of profits $\bar{\pi}_g$ is calculated as:

$$(1 - \delta_i) \sum_{t=1}^{\infty} \delta^{t-1} \pi_t \quad (2.19)$$

Importantly, the impact of inflation suggests that to compare across treatments, real (inflation adjusted) prices should be considered instead of nominal prices averaged over a supergame. Note that we are interested in calculating the average supergame payoff, rather than average prices. Average supergame payoffs are necessary to identify overall equilibrium selection of an infinite game. Equation 2.19, which calculates this, is valid for an infinite sequence. As we do not have realised profits for an infinite sequence in our experiment, we need to alter this slightly in order to re-scale the stream of realised supergame profits. This will provide us with the average payoff, which if received in every stage-game of an infinite sequence,

yields the same present value as the realised sequence of payoffs in the supergame. With discounting, for each supergame g we are able to calculate the average payoff $\bar{\pi}$, given the termination period T , as:

$$\bar{\pi}_g = \frac{\sum_{t=1}^T \delta^t \pi_t}{\sum_{t=1}^T \delta^{t-1}} \quad (2.20)$$

This gives us an average stage game payoff, which if received in every period of the supergame, would yield the same average payoff as the actual stream of payoffs incurred. This is a convenient measure of collusion, given that period profits may vary, and also provides a discount rate weighted payoff. This will allow us to make comparisons between treatments.

Repeated-game logic suggests that treatment $n4\delta_h$ is likely to have the highest level of cooperation, as it has the highest difference between the continuation profit from cooperating forever, and the profit from deviating today and receiving Nash forever. Given our parameters, the difference between cooperating forever (153.66) and deviating and reverting to Nash forever (132.1) is 21.5 ECU in $n4\delta_h$. Thus we expect this treatment to have the highest average payoff. By design we have equalised this difference for $n1\delta_h$ and $n4\delta_l$ at 8.71. Therefore, repeated-game logic would not indicate a difference in average payoff between the two. As the difference in continuation payoff is lower than $n4\delta_h$, the average payoff should also be lower. Our final treatment $n1\delta_l$ is the only treatment where collusion is not theoretically sustainable as an equilibrium, as the value is negative at -4.04. Therefore, it should have the lowest average supergame payoff. The greater the relative benefits from colluding compared to deviating, the higher the expected average payoff $\bar{\pi}$. If participants behave according to standard repeated-game logic, they correctly take into account these collusion incentives through the number of available price changes per period, either one or 4, and the resulting discount factor. This can be stated as follows in our null hypothesis:

H₀: *The average period payoff $\bar{\pi}$ increases with the difference in continuation payoff from colluding and deviating $\Rightarrow \bar{\pi}_{n4\delta_h} > \bar{\pi}_{n1\delta_h} = \bar{\pi}_{n4\delta_l} > \bar{\pi}_{n1\delta_l}$.*

Alternatively, if it is pricing frequency within each period that leads to differences in $\bar{\pi}$, then it would be expected that average payoffs would be lower as n increases, irrespective of the discount factor. The rigidity imposed by only allowing one price change per day means that within-period competition is stifled because price changes are delayed. After observing prices and profits in each pricing period, firms may decide to adapt their price in response, leading to undercutting. This is similar to the undercutting observed in Maskin and Tirole (1988)'s alternate move model, with price cycles characterised by a period of undercutting before a relenting and a price jump occurring. Although the assumptions of their model are not entirely met in our simultaneous move game, it provides an interesting description of firm behaviour, and resulting cycles that are frequently observed in retail petrol markets around the world. In this case, our alternative hypothesis can be stated as:

H_{A1}: *Higher pricing frequency leads to lower average payoffs $\Rightarrow \bar{\pi}_{n1\delta_h} > \bar{\pi}_{n4\delta_h}$ and $\bar{\pi}_{n1\delta_l} > \bar{\pi}_{n4\delta_l}$.*

The second alternative hypothesis is that only discount factors are relevant, and not the frequency of price changes. For repeated-game logic, the number of price changes impacts how quickly firms can punish, and therefore the threshold level of patience - as measured through the critical discount factor. If the frequency of price changes does not matter, then punishment strategies would not be used *within* a period. Punishment strategies would only occur across periods. Thus, if pricing frequency does not impact collusion, then there should be no difference in $\bar{\pi}$ across n , for a given induced discount factor δ . The critical discount factor $\bar{\delta}$ would be the same for a given stopping probability and inflation rate if n does not impact behaviour.¹⁰

¹⁰See equations 2.15 and 2.16, where if $n > 1$ is not used as part of a punishment strategy such as grim trigger, then effectively $n = 1$ for both regimes.

In response to a lower discount factor, participants may compensate for earning lower real profits by raising prices in an attempt to raise real profits. This behaviour may occur in the context of experiments where the reality of implementation weakens the logic of an infinitely repeated game and participants target an aspirational real payoff. This impacts $\bar{\delta}$ in the following way. In high inflation environments, a higher profit level is required to achieve the same real payoff in environments with lower inflation. This would result in higher period profits, meaning that the average supergame payoff would need to be higher. We would therefore expect differences between discount factors, within n :

H_{A2} : *The lower the discount rate, the higher the average payoff $\Rightarrow \bar{\pi}_{n1\delta_h} < \bar{\pi}_{n1\delta_l}$ and $\bar{\pi}_{n4\delta_h} < \bar{\pi}_{n4\delta_l}$.*

Our final alternative hypothesis is that both the discount factor and pricing frequency impact average payoffs - a combination of the above two alternative hypotheses. For $H_{A1,2}$ to be supported, we would need to have evidence in favour of both H_{A1} and H_{A2} . If a higher discount factor results in less collusive outcomes due to participants targeting an aspirational real payoff, but pricing frequency reduces collusive outcomes through adaptive behaviour, then we would expect the following ordering of average supergame payoffs. The highest average payoff would be observed in $n1\delta_l$, where a low discount factor leads to higher prices as participants chase higher real profits. There are no within-period price changes to promote a decrease in prices. The lowest supergame payoff would be expected in $n4\delta_h$ where within-period competition is possible, and the discount factor is relatively high. It is unclear which effect dominates, thus the ordering between $n1\delta_h$ and $n4\delta_l$ is ambiguous. This can be stated as follows:

$H_{A1,2}$: *Higher discount factors and fewer price changes support collusive outcomes $\Rightarrow \bar{\pi}_{n1\delta_l} > \bar{\pi}_{n1\delta_h} \sim \bar{\pi}_{n4\delta_l} > \bar{\pi}_{n4\delta_h}$.*

Given theory on infinitely repeated games does not provide testable hypothesis (i.e., the multiplicity of equilibria), we are not testing this with our experiments.

The implication is that we cannot sharply define what is and isn't equilibrium play. Hence, it is not possible to identify outcomes driven by confusion. The control questions and practice round before the start of the experiment are designed to mitigate any concern regarding the potential impact of confusion.

2.4 Results

First, we present an overview of price results for each treatment to provide insight into the choices made by participants. Next, we test our hypotheses and conduct a brief exploration of dynamics. Overall, our results show that increased pricing frequency reduces collusion levels, mainly due to within-period price decreases. We find no evidence to suggest that the discount factor is a sufficient condition for collusive outcomes to arise.

Table 2.2: Demographics across treatments

		$n1\delta_h$	$n1\delta_l$	$n4\delta_h$	$n4\delta_l$
Completed high-school math %		78.33	74.44	74.07	65.15
Gender %	Female	46.67	45.56	42.59	45.45
	Male	53.33	52.22	55.56	54.55
	Other	0.00	2.22	1.85	0.00
Age group %	18-25 years	71.67	75.56	68.52	75.76
	26-30 years	13.33	14.44	14.81	7.58
	31+	15.00	10.00	16.67	16.67
Home country %	Asia	53.33	40	59.26	51.52
	Australia	35.00	52.22	29.63	37.88
	Other	11.67	7.77	11.10	10.61

Table 2.3 presents the summary statistics for the average price of independent observations in each treatment.¹¹ All treatments are highly competitive, with average prices being close to the Nash price of 2. On average, prices are higher for $n1$ treatments, while $n4\delta_h$ has the lowest mean price and standard deviation.

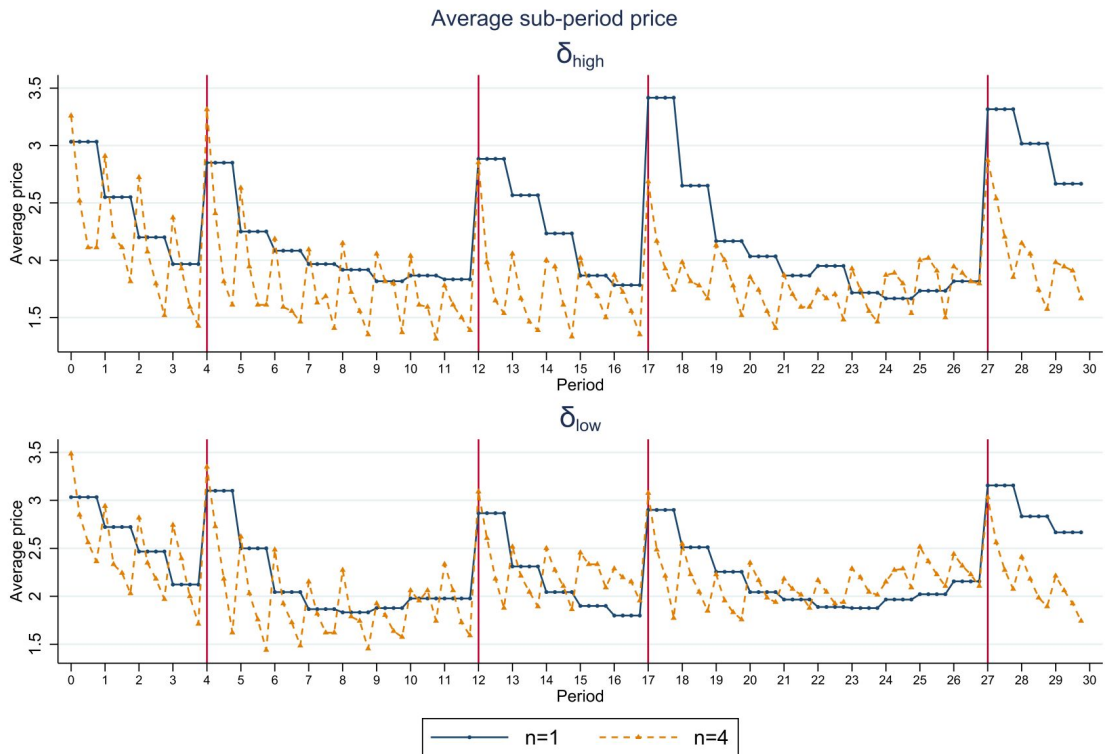
Figure 2.2 plots the average sub-period price in each treatment over the course of the experiment (for $n1$, this is the period price). The vertical lines indicate the start of a new supergame. Simple observation reveals price cycles in each treatment.

¹¹The average price over all supergames, for each matching group of six, in each treatment.

Table 2.3: Average price of independent observations.

Treatment	Mean	Std. dev.	Min	Max	Obs
$n1\delta_h$	2.26	0.33	1.65	2.82	10
$n1\delta_l$	2.29	0.52	1.64	3.72	15
$n4\delta_h$	1.86	0.20	1.66	2.18	9
$n4\delta_l$	2.16	0.53	1.60	3.40	11

For $n1$ we observe a jump at the start of each supergame, decreasing over each new price setting period. For $n4$, we observe period-length cycles characterised by within-period price decreases. This suggests that pricing frequency may impact the focal point chosen by participants, being the start of a supergame or new period. Starting prices in $n4$ treatments align with the downward trend of prices in the $n1$ treatments, suggesting that within-period price competition is driving the lower average price.

Figure 2.2: Average sub-period price in all 5 supergames in δ_h (top) and δ_l (bottom).

We now look closer at how frequently each price was played by participants. The Nash price, or lower (price ≤ 2), is played with much higher frequency for the $n4$ treatments, as shown in Table 2.4. Although the percentage of prices played above

Nash in $n4\delta_l$ are comparable to that in the $n1$ treatments, there is a higher proportion of prices below Nash at 41%, compared to approximately 26%. Considering the first sub-period only, the proportion of prices above Nash in our $n4$ treatments increases, in line with price cycle jumps before inter-period competition.

Table 2.4: Price proportions in each treatment.

	Treatment					
	$n1\delta_h$	$n1\delta_l$	$n4\delta_h$	$n4\delta_l$	Sub-period 1 only	
					$n4\delta_h$	$n4\delta_l$
Below Nash %	26.7	25.5	45.1	41.3	30.9	30.3
Nash (p=2) %	46.1	46.1	36.4	31.9	35.9	30.7
Above Nash %	27.3	28.0	18.5	26.8	33.2	39.1

2.4.1 Average payoffs

Accurate comparison across treatments requires a standardised measure of profit levels. We have provided an overview of the price choices made by participants, and in this section we determine whether this had a differential impact on overall supergame outcomes.

Table 2.5: Average supergame payoffs for independent observations.

Treatment	Mean	Std. dev.	Min	Max	Obs
$n1\delta_h$	32.52	1.35	29.87	34.49	10
$n1\delta_l$	33.00	1.78	29.50	36.46	15
$n4\delta_h$	30.56	1.15	28.85	32.22	9
$n4\delta_l$	31.35	2.16	27.76	35.57	11

Although our markets are highly competitive in all treatments, $\bar{\pi}$ remains a standardised measure of market outcomes. This is despite many markets not necessarily landing in the “hull” of Figure 2.1. Summary statistics for each independent observation on realised average supergame payoffs are shown in Table 2.5. Recall that the stage game Nash profit is 33. On average $\bar{\pi}$ is Nash or below in all treatments, further evidence of markets being over-competitive. Thus, prices in Table 2.3 appear to reflect the level of average supergame payoffs.

Figure 2.3 shows the distribution of average supergame payoffs in each treatment

for each supergame. Variation in $n1$ appears greater than the variation in average supergame payoffs in the $n4$ treatments. It is unsurprising that the median is slightly higher in $n1$ treatments for supergame 5, as it was the shortest in length and average prices began higher than in $n4$ for period 1.




Figure 2.3: Distribution of the average payoff in each treatment for each supergame.

We start by testing our null hypotheses relating to average supergame payoffs $\bar{\pi}$ in each treatment. According to supergame logic, we expect the highest $\bar{\pi}$ would be observed in treatment $n4\delta_h$. However, this treatment has the lowest $\bar{\pi}$ in Table 2.5. The results from conducting a Mann-Whitney rank-sum test on independent observations for treatment differences are summarised in Table 2.6. There are three statistically significant differences. For a given δ , prices are lower in $n4$ treatments compared to $n1$. The third significant difference is that there are higher prices in $n1\delta_l$ compared to $n4\delta_h$. For robustness, we run a simple OLS model on average market payoffs in each treatment, for each supergame, clustering standard errors by matching group. The regression estimates, and post-estimation tests of equality, are reported in Tables 2.7 and 2.8. Results from rank-sum tests on independent observations, and tests on the equality of regression coefficients, provide the same qualitative outcome.

Table 2.6: Pairwise rank-sum tests of equality.

Pairwise n	p-value	Pairwise δ	p-value	Pairwise $n \times \delta$	p-value
$n1\delta_h = n1\delta_l$	0.7	$n1\delta_h = n4\delta_h$	0.007***	$n1\delta_h = n4\delta_l$	0.11
$n4\delta_h = n4\delta_l$	0.34	$n1\delta_l = n4\delta_l$	0.036**	$n1\delta_l = n4\delta_h$	0.0012**

*p < 0.1, ** p < 0.05, ***p < 0.01.

Table 2.7: OLS regression on market supergame payoffs (clustering on matching group).

	Average payoff $\bar{\pi}$ (std. error)
Treatment (base $n1\delta_h$)	
$n1\delta_l$	0.48 (0.61)
$n4\delta_h$	-1.96 (0.55)
$n4\delta_l$	-1.17 (0.76)
Supergame (base 1)	
2	-1.7 (0.2)
3	-1.86 (0.3)
4	-2.00 (0.31)
5	-1.33 (0.35)
_cons	33.89 (0.41)

Result 1 *Collusion incentives do not lead to higher average supergame payoffs.*

For each treatment comparison for each hypothesis, Table 2.9 states the expected ordering of supergame payoffs, and the results of each rank-sum test compared to the hypothesised direction. We reject our null hypothesis as the hypothesised inequalities do not hold. This leads to the acceptance of our first alternative hypothesis H_{A1} that pricing frequency impacts average supergame payoffs. Both comparisons are statistically significant at the 5% level in the correct direction. We find no significant evidence in support of H_{A2} , thus we do not accept this, nor $H_{A1,2}$.

The null hypothesis states that higher $\bar{\pi}$ would be observed if collusion incentives, according to repeated-game logic, are higher. We reject the null hypothesis, given the statistically significant difference in inequalities. Payoffs in $n1\delta_h$ are higher than $n4\delta_h$ (p=0.007), and also for $n1\delta_l$ and $n4\delta_l$ (p=0.036). The incentives to collude are highest in $n4\delta_h$, yet we find that $\bar{\pi}_{n4\delta_h}$ is lowest on average. As the difference in continuation payoff from colluding and deviating increases, the incentive to collude increases. However, we do not observe this ordering of average supergame payoffs. Furthermore, collusion is sustainable in $n4\delta_l$, but not sustainable in $n1\delta_l$. That

Table 2.8: Pairwise test on equality of coefficients.

Pairwise n	F-statistic (p-value)	Pairwise δ	F-statistic (p-value)	Pairwise $n \times \delta$	F-statistic (p-value)
$n1\delta_h = n1\delta_l$	0.62 (0.44)	$n1\delta_h = n4\delta_h$	12.46 (0.001***)	$n1\delta_h = n4\delta_l$	2.38 (0.13)
$n4\delta_h = n4\delta_l$	1.16 (0.29)	$n1\delta_l = n4\delta_l$	4.49 (0.04**)	$n1\delta_l = n4\delta_h$	17.41 (0.0001***)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.9: Hypothesis and rank-sum test on supergame payoffs between treatments.

H_0 Collusion incentives			H_{A1} Pricing frequency			H_{A2} Discount factor		
Comparison	p-value	Inequality (hypothesised)	Comparison	p-value	Inequality (hypothesised)	Comparison	p-value	Inequality (hypothesised)
$n4\delta_h = n1\delta_h$	0.007	<*** (>)	$n1\delta_h = n4\delta_h$	0.007	>*** (>)	$n1\delta_h = n1\delta_l$	0.7	~ (<)
$n1\delta_h = n4\delta_l$	0.11	~ (=)	$n1\delta_l = n4\delta_l$	0.036	>*** (>)	$n4\delta_h = n4\delta_l$	0.34	~ (<)
$n4\delta_l = n1\delta_l$	0.036	<*** (>)						

~ null of equality not rejected, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

payoffs are higher in $n1\delta_l$ supports this result. Therefore we find no evidence to support the impact of collusion incentives, as defined using repeated-game logic, on average supergame payoff $\bar{\pi}$. This leads us to our first result that collusion incentives do not necessarily lead to higher average payoffs.

Result 2 *Average supergame payoffs are higher in $n1$ treatments compared to $n4$ treatments.*

The question remains as to what is, and what is not, driving the differences between treatments. In Table 2.9 it can be observed that both inequalities for H_{A1} are statistically significant and in the correct direction. This tells us that the frequency of price changes does impact $\bar{\pi}$, with increased frequency leading to lower average payoffs.

Result 3 *Discount factor not a significant determinant of average supergame payoffs.*

Theory states that less collusion should be observed in environments where it is not sustainable in equilibrium, thus we would expect lower payoffs in $n1\delta_l$ as the induced discount factor is below $\bar{\delta}_1$. First, we test the difference between treatment

$n1\delta_h$ where collusion is theoretically sustainable, and $n1\delta_l$ where it is not, as per repeated-game logic.

We are unable to reject the null hypothesis that the average supergame payoffs for independent observations are equal ($p=0.70$). Although collusion is possible in both $n4$ treatments, we find no difference between $n4\delta_h$ and $n4\delta_l$ ($p=0.34$). Therefore, our results suggest that the critical discount factor, and discount rate more generally, is not a key determinant for collusion to arise. This finding is line with that of Feinberg et al. (2017), and Dal Bó and Fréchette (2018) who analyse meta-data on infinitely repeated prisoner's dilemmas and find that cooperation being an equilibrium action is not sufficient for cooperation. There is no difference when comparing the treatments where we control for the continuation payoff ($n1\delta_h \sim n4\delta_l$ ($p=0.11$)). This result shows that this aspect of repeated-game logic is not a key factor in determining levels of collusion and reject alternative hypothesis H_{A2} . The alternative behavioural explanation of participants aiming for a target real payoff does not appear to be the case either.

2.4.2 Dynamics

We have identified the aggregate impact of pricing frequency on collusion in the context of an infinitely repeated game. In the following sections we are interested in looking closer at decisions within markets. Figure 2.4 shows the real value of average profit over time. The difference between treatments is driven by the within-period decreases in $n4$, while period starting profit levels are similar.

Observing $n4\delta_l$ from the third supergame (period 12 in Figure 2.2) onward, the average price appears to drift upwards over time to an extent not observed in the other treatments. The upward drift of prices observed for treatment δ_l appears to be in contrast to the findings in the literature, as the higher inflation rate is potentially leading to more cooperation, whereas the opposite has been found in other studies (Feinberg and Husted, 1993; Dal Bó and Fréchette, 2018). We hypothesise that




Figure 2.4: Real value of average payoffs.

this is an experimental artefact resulting from the higher inflation rate, whereby participants start increasing their prices to compensate. Figure 2.4 supports this by showing the equalisation of average profit in real terms. The within-period competition is driving the difference between n treatments, with the additional price changes lowering average period profit, rather than the starting profit level.

Result 4 *Higher pricing frequency leads to shorter price cycles, and lower prices.*

In this section we investigate how pricing frequency impacts dynamics, noting the limitations that arise due to the short length of our experiment. Although participants played relatively competitively, we look closer at how collusive agreements, or attempts at them, were initiated. Figures showing individual group prices for all supergames, for each treatment, can be found in Appendix A.2.

In order to examine individual behaviour more clearly, we develop a ranking system from 5 to 1 to track differences in relative prices of individuals in each group over time, as follows:

- 5 Highest price ($p_i > p_j, p_k$)
- 4 Tied highest price ($p_i = p_j \ \& \ p_i > p_k \mid p_i = p_k \ \& \ p_i > p_j$)
- 3 Price in-between ($p_i < p_j \ \& \ p_i > p_k \mid p_i < p_k \ \& \ p_i > p_j$)
- 2 Tied lowest price ($p_i = p_j \ \& \ p_i < p_k \mid p_i = p_k \ \& \ p_i < p_j$)
- 1 Lowest price ($p_i < p_j, p_k$)

This ranking system is then used to determine whether a change in rank occurs between periods/sub-periods. We then generate proportion of rank changes for each group, in each supergame, by dividing the number of changes by the number of pricing opportunities. The resulting proportions can be seen in Figure 2.5. It appears that there is a marginally higher proportion of rank discrepancies for $n4$, which may indicate more attempts at signalling intentions by individual firms. Increased signalling, but lower prices, indicates failed attempts at cooperation.




Figure 2.5: Frequency of change in price rank in each supergame, divided by the number of price change opportunities.

Cycle length and adaptive pricing

Within cycle price decreases are apparent for both pricing frequencies. However, with a higher frequency of interaction a more frequently occurring focal point, or restart point, was chosen by participants. The adaptive pricing that occurred within-cycle is therefore driving down the overall average price. If we only consider the average starting price (sub-period 1), there is no difference in prices with Mann-Whitney tests between each treatment having a p-value ≥ 0.169 . There is no statistically significant difference between treatments for a given n , with respect to the discount factor ($n1\delta_h \sim n1\delta_l$, Mann-Whitney rank-sum $p=0.89$), and $n4\delta_h \sim n4\delta_l$ ($p=0.16$). If we restrict prices to sub-period 1 only, the proportion playing above Nash increases, as shown in Table 2.4.

We look closer at individual supergames. Table 2.10 shows the p-values for Mann-Whitney rank-sum tests for each treatment combination in each supergame. The significance levels are as expected based on the tests of independent observations for each treatment above. Considering sub-period 1 only, the day starting price, note the reduced significance in difference between supergame prices in each treatment. This shows that the adaptive price setting is driving the reduction in average prices.

Table 2.10: Mann-Whitney rank-sum test, p-value for difference between treatments for each supergame

Supergame	1	2	3	4	5
$n1\delta_h$ vs $n4\delta_h$	0.094*	0.0143**	0.0025***	0.16	0.011**
$n1\delta_l$ vs $n4\delta_l$	0.52	0.186	0.876	0.516	0.021**
$n1\delta_h$ vs $n1\delta_l$	0.454	0.541	0.56	0.824	0.56
$n4\delta_h$ vs $n4\delta_l$	0.271	0.171	0.21	0.47	0.879
$n1\delta_h$ vs $n4\delta_l$	0.944	0.231	0.833	0.597	0.029*
$n1\delta_l$ vs $n4\delta_h$	0.006***	0.011**	0.089*	0.101***	0.034**
Sub-period 1 only					
$n1\delta_h$ vs $n4\delta_h$	0.030**	0.191	0.653	0.87	0.072*
$n1\delta_l$ vs $n4\delta_l$	0.169	0.153	0.096*	0.406	0.336
$n1\delta_h$ vs $n1\delta_l$	0.454	0.541	0.36	0.824	0.56
$n4\delta_h$ vs $n4\delta_l$	0.47	0.382	0.159	0.239	0.675
$n1\delta_h$ vs $n4\delta_l$	0.078*	0.041**	0.525	0.459	0.231
$n1\delta_l$ vs $n4\delta_h$	0.074*	0.492	0.676	0.633	0.199

* $p \leq 0.1$, ** $p \leq 0.05$, *** $p \leq 0.01$

2.5 Conclusion

In this chapter we have demonstrated the theoretical and behavioural uncertainty in how action frequency impacts cooperation. We embedded this question within the context of a Bertrand triopoly to provide insight into the potential impact of pricing frequency regulations, inspired by the 24-hour price commitment regime in Western Australia. By incorporating the spirit of these regulations in an infinitely repeated game, it is possible to distinguish between the potential impact of repeated-game logic from the impact of other behavioural mechanisms. Does making a simultaneous decision either once, or four times, during a stage game, result in higher levels of collusion? Four explanations are tested. Our null hypothesis is that the more robust collusion is in the context of repeated games, the higher the level of tacit collusion in that market - behaviour is driven by collusion incentives. The first alternative hypothesis states that pricing frequency itself may lead to lower collusion levels if prices are set adaptively, similar to the decay observed in public goods games and in the spirit of Maskin and Tirole (1988) Edgeworth cycles. The process of undercutting can happen at a faster rate with more price changes per period. The second alternative hypothesis states that satisficing behaviour, derived from the discount factor, impacts collusion levels. A lower discount factor (higher inflation) leads to compensation in the form of higher prices to achieve a target level of real profit. The final hypothesis is that both pricing frequency and the discount factor impact collusion levels. We test these hypotheses using four treatments that capture each of these elements.

Results show that higher pricing frequency leads to lower levels of collusion. The frequency of price changes per period impacts collusion levels by promoting inter-period competition. On the aggregate level we observe Edgeworth price cycles in all treatments. Interestingly, with only one price change per period firms use the start of a new supergame as a focal point to raise prices, decreasing prices between each period. With four price changes, firms instead raise prices at the start of each

period, and decrease them within. This suggests that adaptive pricing is present in all treatments, but the frequency of price changes impacts the length of cycles by creating a different focal point to increase prices on. This could be similar to the restart effect and decay in cooperation observed in repeated prisoner's dilemmas (Chaudhuri, 2011), and by Leng et al. (2018) who find a restart effect of periods in a minimum effort coordination game.

This chapter has established that characteristics of infinitely repeated games may not provide policy relevant insights on the impact of pricing-frequency restrictions, and that behavioural mechanisms are relevant. Having bridged the gap between experiments on infinitely repeated games and behaviour in this chapter, the remaining chapters focus on testing how cooperation is impacted by pricing frequency without the constraints that infinitely repeated games impose in a laboratory setting.

Chapter 3

Continuous time

3.1 Introduction

The stage game is the cornerstone of repeated-game logic. Traditionally, repeated games in a theoretical and laboratory environment have been modelled as simultaneous or alternate move, e.g., Leufkens and Peeters (2011) follow the models of Maskin and Tirole (1988) and Fudenberg and Maskin (1986). For applications involving one-off decisions, e.g., bidding in an auction, or the decision to enter a market, these timing protocols capture the game's overall structure. However, for applications involving repeated decisions of the same type, with the same players, strict adherence to distinct decision periods, and order of moves, may obfuscate underlying behavioural impacts of policy changes.

The literature on endogenous-move games relaxes these assumptions by allowing for the order of moves in a game to be determined by players. For example, Huck et al. (2002) experimentally test the predictions of an endogenous Stackelberg leadership model, however the focus is on the order of moves, rather than the impact of timing. Differential games can model endogenously ordered decisions made in continuous time (Vives, 2001). However, a key assumption made by differential games is that players have an infinitely fine adjustment speed, which may not reflect actual

decision making by humans. Thus, we are interested in game structures that do not restrict the order of decision timing and also allow for reaction times greater than zero. Simon and Stinchcombe (1989) model continuous time using an infinitely fine grid and show that it is a useful way to model behaviour for dynamic problems. This has carried over to an emerging area of experimental literature that investigate the implementation and effect of continuous and discrete time on cooperation (Friedman and Oprea, 2012; Bigoni et al., 2015; Leng et al., 2018; Horstmann et al., 2016).¹

In Chapter 2 we investigated interaction frequency by placing stage-game payoffs on a grid that was either of size one or four. However, players were still restricted by a simultaneous-move paradigm and were unable to update prices more frequently. This is incongruous with the decision structure observed in many markets. The WA pricing frequency regulations implement a simultaneous move grid structure on an otherwise endogenously determined ordering of price change decisions which can be made as quickly as the firm can react. Understanding how the WA restrictions might impact behaviour therefore requires comparison of cooperation levels with and without restrictions in an otherwise real-time environment. We define a real-time environment as having a grid-size is smaller than the minimum reaction time of participants, and where payoffs are therefore incurred in a flow over time.

An unintended consequence of the design in Chapter 2 was the use of a different restart point by firms for each pricing frequency treatment - the start of a new day or supergame. To avoid this as an additional confound, in this chapter we use an experimental design that runs in real-time without breaks that would signify distinct periods. In addition, in Chapter 2 the game length was limited by the continuation probability and inflation rate. Thus, we now move away from the implementation of an infinitely repeated game, and observe the long-run dynamics of a single continuous game run for 60 minutes. We implement a symmetric duopoly where firms sell differentiated goods in real-time, and compare collusion levels between a treatment

¹The grid size is determined in these experiments by how frequently the experimental software is able to update.

where firms can change their prices instantaneously (and asynchronously), or are locked into a price for a set time (simultaneous updating of prices for the market). Although a triopoly may better reflect the number of competing firms in many markets such as retail petrol, we implement a duopoly to promote comparability with continuous-time experiments conducted in the literature (Friedman and Oprea, 2012; Bigoni et al., 2015).

There are three key subsets of the experimental literature relevant to our research question. The first is the literature on the more general impact of continuous time on cooperation. Friedman and Oprea (2012) run a prisoner’s dilemma in continuous time, with a flow of profits incurred over 60 seconds. The authors find that cooperation is sustainable in continuous time, but less so when choices are made on an 8-sub-period grid, and rarely for 1-shot interactions. Bigoni et al. (2015) compare cooperation levels in continuous-time social dilemmas, finding that continuous time leads to qualitative differences in behaviour compared to discrete time, and that deterministic ending rules make cooperation easier to achieve and maintain compared to games that end stochastically. Leng et al. (2018) run experiments using a coordination game and find that continuous time can help resolve strategic uncertainty that may hamper cooperation, but not in all circumstances. We contribute to the literature by exploring whether continuous time supports cooperation in more complex environments.

The second key subset of the literature is that which explores how continuous time impacts competition. This area of the literature is sparse. Horstmann et al. (2016) use an experimental design that compares cooperation levels in discrete and continuous time for differentiated Cournot and Bertrand competition, and duopolies and triopolies. The authors’ show that tacit collusion is higher in discrete time, compared to continuous. This finding is different to previous experimental studies. Thus, the impact of continuous time in an industrial organisation context warrants further investigation. We contribute to this area by providing further evidence of the impact of continuous time in oligopoly markets.

The third relevant subset of the literature is experiments that test retail petrol pricing regulations. Experimental work has attempted to provide a more controlled approach to causally identifying the impact of pricing frequency restrictions. Contrary to empirical studies that use real petrol-pricing data, Berninghaus et al. (2012) and Haucap and Müller (2012) find that the Austrian rule leads to higher prices in markets.² Contrary to our results, Haucap and Müller (2012) find the WA rule has no impact on the price level, however we note that with four firms experimental markets are generally over-competitive (Huck et al., 2004), which may result in more power being needed to identify any difference where collusion is already unlikely. Our experimental design abstracts from real world market characteristics in order to identify the causal impact of pricing frequency restrictions. These abstractions allow us to cleanly distinguish effects between real-time and the discretisation of decisions in an otherwise real-time environment.

Repeated-game logic supports the proposition of higher prices in real-time environments through easier initiation and maintenance of tacit collusion. In general, tacit collusion is more likely to be initiated and sustained in a dynamic, or repeated game context, if signalling costs and gains from cheating on an agreement are low. Initiation through price leadership is cheaper when prices can be updated more quickly. Signalling an intention to collude is less risky - other firms can more quickly respond in-kind, otherwise a failed signal can quickly be lowered. Similarly, once an agreement is reached, punishment of deviators can be implemented more swiftly, meaning that breaking an agreement is less profitable. In terms of behaviour, empirical evidence has shown that continuous time may also support higher levels of collusion (Friedman and Oprea, 2012; Bigoni et al., 2015; Leng et al., 2018).

A potential reason why higher prices might be observed in discrete settings could be that with coordination becomes easier. Byrne and de Roos (2019) show that retail petrol stations in WA were able to collude using days of the week as a coordination device. In WA, state regulations permit only one price change per day, making days

²The Austrian rule allows for a single price increase, but unlimited price decreases, per day.

of the week an obvious focal point. Without these discrete pricing periods, there is no natural coordination device, thus in other jurisdictions without restrictions it might be harder for firms to coordinate on a collusive equilibrium. Therefore *a priori* it is unclear in which direction pricing frequency restrictions might impact price levels, and whether there are any cognitive differences in how firms behave in real compared to discrete time.

In order to identify which of these factors impact collusion, we compare price levels in real-time experimental markets, with and without simultaneous discrete pricing frequency restrictions (implementing the pricing rule found in WA). We create a high information, real-time environment where two firms compete in prices over a 60 minute experiment. Using a logit demand function (Anderson and de Palma, 1992) with known demand shocks, we create an environment highly conducive to both collusion and deviation. Demand fluctuates between day (high) and night (low) every 20 seconds, requiring participants to actively set prices to maximise their payoff. In our first treatment, there are no restrictions on how often firms can change their price.³ Our second treatment implements discrete pricing periods in an otherwise real-time environment, where firms can select a different price in real time, but the change only takes effect in the market every 20 seconds.

Results show significantly higher levels of collusion in markets with unrestricted price changes. This difference quickly emerges at the start of the 60 minute experiment, and is maintained until the final seconds where collusion unravels. Looking closer at individual markets we find a significantly higher proportion of fully collusive markets in our unrestricted treatment. In markets that were not able to successfully collude, this difference in prices between treatments persists. Thus, we show that continuous time supports collusion, while the introduction of discrete pricing periods into an otherwise real-time environment leads to fewer fully collusive markets and lower prices more generally.

³Computationally there is a 0.5 second restriction, as this is how frequently the program updates. As the reaction time required for participants to view a change their strategy is greater than 0.5 seconds, from the participant's point of view, the market is perceived to run continuously.

The remainder of the paper is organised as follows. First, we outline our model, implementation and predictions. Second, we provide details on our experimental design. Third, we show the results from our experiment, before concluding.

3.2 Model and implementation

In this section we describe our model and its implementation for our experiment. We are interested in how action frequency impacts cooperation. To facilitate this we focus on price competition in long-run symmetric duopoly markets. Firms sell products that are imperfect substitutes for each other. This represents the nature of competition between petrol stations, or online sellers of branded goods, where products are similar but differentiated by location, loyalty program, or services provided.

We use a multinomial logit discrete choice demand function (Anderson and de Palma, 1992), a non-linear form of Bertrand price competition with product differentiation. The rationale for using this demand structure is twofold. First, it captures the above characteristics of consumer demand in our market, which will be explained in more detail below. Second, compared to the predominantly used linear demand function, this function allows for a wider gap between Nash and joint profit maximising (JPM) prices. This is a desirable feature for our experiments, as it allows for more pronounced treatment effects. In addition, the difference between JPM and Nash profits is significantly greater than with linear demand. The environment is therefore designed to be behaviourally conducive to collusion, as we are interested in how pricing frequency impacts the ability to cooperate. If participants play close to Nash, or competitively, treatment effects may not be as clear. Similarly, we keep in mind the competitiveness of experimental oligopolies with two and three firms (Huck et al., 2004).

3.2.1 Model

In what follows, we provide a general description of the structure of the market in our experiments. In our model two identical firms compete in price and sell products that are imperfect substitutes for each other. We use a demand function from a multinomial discrete-choice model, which was introduced to the literature by Anderson and de Palma (1992). The expressed demand for the product of firm i is expressed in terms of the probability that a randomly chosen buyer buys the good and depends on the price p_i , the competitor's price p_j , the reservation price a , the value of the outside option V_0 , and the consumers' preference for variety μ . Normalising the number of consumers to unity, we can express demand $\mathcal{P}(p_i, p_j, V_0)$ as:

$$\mathcal{P}_i(p_i, p_j, V_0) = \frac{\exp[(a - p_i)/\mu]}{\exp[(a - p_i)/\mu] + \exp[(a - p_j)/\mu] + \exp[V_0/\mu]}, \quad i \neq j \quad (3.1)$$

Note that such a demand function can be either interpreted as *ex ante* identical consumers having idiosyncratic utility shocks (McFadden, 1981), the result of consumers balancing exploration and exploitation (Weisbuch et al., 2000; Duncan, 1959), or consumers exhibiting rational inattention (Fosgerau et al., 2020).

With constant marginal cost c , firm i 's profit function is:

$$\pi_i(p_i, p_j) = (p_i - c)\mathcal{P}_i(p_i, p_j, V_0), \quad i \neq j \quad (3.2)$$

Strategies are strategic complements, such that firm i 's best response is to increase price if firm j increases their price. To see this consider the case of a duopoly market, where the best response function of firm i to the price of firm j is given by:

$$BR_i(p_j) = c + \mu + \mu W_0 \left[\frac{e^{-1 + \frac{(a-c+p_j)}{\mu}}}{e^{\frac{a}{\mu}} + e^{\frac{p_j+V_0}{\mu}}} \right], \quad (3.3)$$

where W refers to the Lambert W function. We use W_0 , which gives the principle

solution to $z = we^w$. As the Lambert W function is increasing in its argument and the argument (as can be shown) increases with p_j , we have the best-response price increase with the competitors price

$$\frac{\partial BR_i(p_j)}{\partial p_j} > 0.$$

The difference between the value of the product and the outside option ($V_0 - a$) determines the maximum price a firm would ever charge. For example, increasing the value of non-purchase relative to a results in profits decreasing as the firm's price increases, i.e., flattens the slope of the best response function. Initially the best response function increases in response to an increase in the price of the other firm. Eventually the function levels off, and increasing own price further only results in consumers taking the outside option of non purchase, rather than increasing the own profit. The taste for variety in consumer preferences, μ , impacts the intersection of the best response functions i.e., the Nash price. As $\mu \rightarrow 0$, the Nash price tends towards marginal cost (i.e., Bertrand duopoly with homogeneous goods), as the consumer now buys from whoever is cheaper. An increase in μ makes demand less elastic, which increases equilibrium prices.

3.2.2 Parameters and Static Equilibrium.

This market has a unique, symmetric Nash equilibrium (see the appendix in Anderson and de Palma, 1992, for a proof). We use two parameterisations of the above model to create high (Day) and low (Night) demand environments and restrict prices to whole numbers. The low demand scenario is characterised by consumers having less of an urge to purchase the good from either firm, which is reflected in a higher outside-option value V_0 . Furthermore, in the low-demand environment consumers have less of taste for variety μ . This can be interpreted in a fuel-market setting that at night less people require fuel but the ones that do are less willing to try out

unfamiliar petrol stations than during the day. Figure 3.1 shows the best response functions and equilibrium in the two demand conditions.

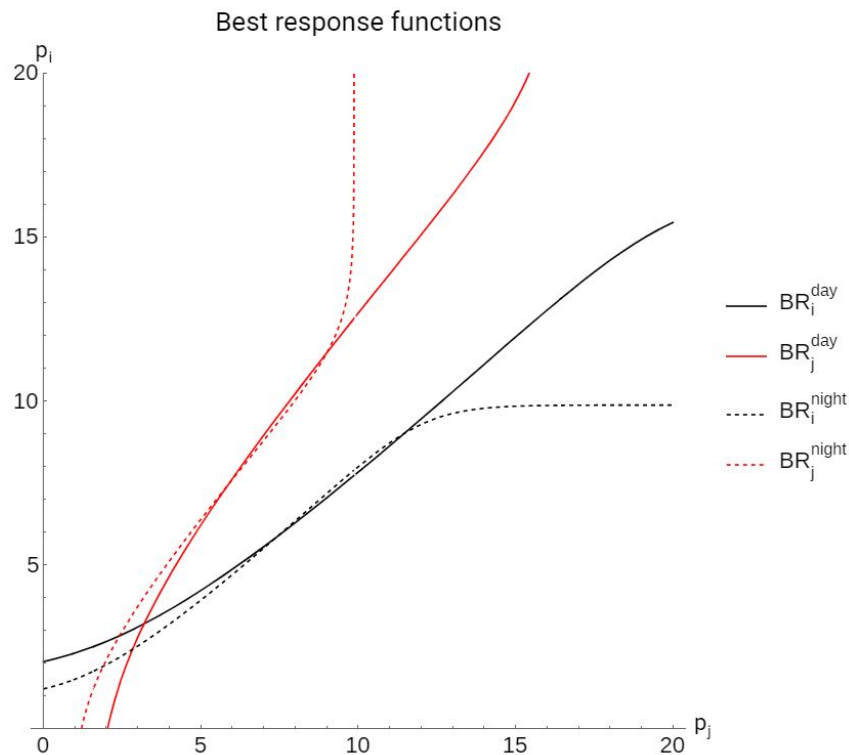


Figure 3.1: Best response functions for Day and Night demand levels, duopoly.

Table 3.1 summarises the static Nash equilibrium predictions. The JPM and one-shot deviation prices and profits are shown for the two demand conditions.⁴

The parameters were chosen to achieve a broad spread between Nash and collusion prices. For our duopoly, the Nash price p^N is 3 or 4 and the JPM price p^{JPM} is 17 during the Day. During Night, the Nash price decreases to 2 and 3, while the JPM price changes to 10. In addition, we still want a competitive outcome, below Nash, to be a possibility. With marginal cost of zero, positive profit can still be earned if both firms price below Nash.

In our experiment we implement a finitely repeated game (90 periods), where the stage game consists of one Day (with high demand) and one Night (low demand). The treatments will determine how many pricing choices participants have to make per stage game. In what follows we do not seek to theoretically characterise all

⁴The multiplicity of static NE arises from the discretisation of the strategy space.

subgame perfect Nash equilibria (SPNE) for the finitely repeated game.⁵ We are primarily interested in determining which price-change regime leads to higher levels of collusion. Once this is established, we will use the actual behaviour to investigate if the mechanisms behind our result are likely to be behavioural or strategic in the sense of repeated games.

Table 3.1: Parameters and Nash, joint-profit maximising, and best shot deviation prices and profits.

Firms	Parameters	High demand	Low demand
		(price) [profit]	(price) [profit]
		$a = 30, V = 10, \mu = 1.6, c = 0$	$a = 30, V = 18, \mu = 0.95, c = 0$
2	Nash	(3, 3), (4, 4) [149, 149], [199, 199]	(2, 2), (3, 3) [99, 99], [149, 149]
	Collusion	(17, 17) [789, 789]	(10, 10) [471, 471]
	Deviation	(14, 17) [1189, 221]	(8, 10) [703, 107]

3.3 Hypothesis and experimental design

In this section we outline our experimental design, hypothesis, and procedure. We use two treatments to identify the impact of pricing frequency restrictions on the ability for firms to tacitly collude in real-time markets. The baseline treatment is an unregulated real-time duopoly where firms are able to update prices without restriction.⁶ As price changes can be made continuously, we refer to this treatment as *continuous*. Our second treatment is still run in real-time, i.e., with flow profits, but now with prices updating in the market every 20 seconds. Prices update at the same time for all firms in the market at these set intervals. With the introduction of discrete pricing periods, we call this treatment *discrete*. The overall measure of collusiveness is the average price level throughout the experiment.

Collusion incentives and continuous time both work to support the hypothesis

⁵The multiple static equilibria and the many choices to be made give rise to a very large set of subgame-perfect Nash equilibria.

⁶The computer program updated every 0.5 seconds. While not literally continuous, the minimum amount of time required for participants to react to new information is assumed to be at least half a second, and thus can be considered continuous from their point of view (Horstmann et al., 2016).

that higher prices are observed in *continuous* compared to in *discrete*. Following repeated-game logic, generally the quicker a firm can react in the market, the more sustainable collusion will be in that environment. This is because firms can react more quickly to punish those that deviate from collusion, making deviation profits small. In the context of repeated games, the threshold level of firm patience required for a collusive outcome to be sustained as a SPNE is lower when deviation profits are low (Fudenberg and Maskin, 1986; Vives, 2001; Mailath and Samuelson, 2006; Dal Bó and Fréchette, 2011). Experiments implementing infinitely repeated games have provided evidence that cooperation being a possible equilibrium on average can increase cooperation (Dal Bó and Fréchette, 2018). In addition, experimental evidence has shown that continuous time can support cooperation (Friedman and Oprea, 2012), thus the introduction of discrete pricing periods may make collusion harder to initiate and maintain.

On the other hand, if the introduction of discrete pricing periods leads to higher prices, then the regulations may facilitate coordination through the introduction of structural breaks that are used as focal points (Schelling, 1960). Although there is evidence of firms using days of the week as focal points to support coordination between petrol stations in WA (Byrne and de Roos, 2019), it is not clear that this necessarily leads to higher price levels compared to an environment without them. A separate explanation could be that restricting pricing frequency makes it easier to keep track of competitors' actions, facilitating cooperative outcomes.

Collusion incentives and continuous time work to support the hypothesis of higher collusion levels in *continuous* compared to *discrete*. Focal points and ease of following the market would lead to lower levels of collusion in *continuous*. Thus it is unclear *ex ante* whether price levels will be significantly different between treatments. We state our null hypothesis that average prices in each treatment are equal as follows:

H1₀: *The average price level is not impacted by simultaneous pricing frequency*

restrictions. $\Rightarrow p^{\text{discrete}} = p^{\text{continuous}}$

If the null hypothesis is rejected, the alternative will depend on which direction the difference in price goes. If the price level is higher in *discrete*, the coordination aspect will be the predominant effect:

H1_a: *Simultaneous pricing frequency restrictions lead to higher price levels (more collusion) $\Rightarrow p^{\text{discrete}} > p^{\text{continuous}}$.*

If the price level is higher in *continuous*, repeated-game logic and/or the impact of continuous time are behind this result:

H1_b: *Simultaneous pricing frequency restrictions lead to lower price levels (less collusion) $\Rightarrow p^{\text{discrete}} < p^{\text{continuous}}$.*

3.3.1 Experimental implementation and procedure

At the start of the experiment participants were asked to select their starting price for the first Day period. Possible prices were the integers between 1 and 20, inclusive. Once all participants had selected their price, the game began. Participants played in the same pair for the entire experiment. Each period contained one Day (high demand), and one Night (low demand), lasting 20 seconds each. These 40 second periods were referred to as “cycles” in the experiment. This allowed a total of 90 periods to be run over 60 minutes. Our data provides us with an observation for each subject every half second, resulting in 7,200 observations per participant.

In the continuous treatment firms were able to change their price at any time. In the discrete treatment, the price selected was only implemented in the next Day or Night. There was no break between periods, and if a participant did not change their price, the previously selected price carried over. The experiment was programmed to run in quasi-continuous time (updated every 0.5 seconds) to avoid the

period break being used as a second restart point by participants.⁷ This ensured the demand shocks were the only structural point for both the discrete and continuous the treatments. Participants needed to actively adjust their price to adapt to the demand shocks - if they did not change their price, the last selected price remained.

A screenshot of the continuous experiment can be found in Figure 3.2. The graph displayed a live stream of prices with a history length of 120 seconds. After the initial 120 seconds, the graph continues to move with older prices disappearing from the graph, being replaced with new ones. The shaded areas represent Night and allowed participants to keep track of when demand changes were approaching. This feature is particularly important for the *discrete* treatment, where the flow of profits continued to be observed between pricing periods. The interface was designed to be intuitive for participants to understand the actions of their partner, and impact on their own profit. For example, the price history line changed colour to represent their current profit level. The brighter the green, the higher the profit. As the line became darker, this indicated a lower profit level.

A table on the upper right of the screen displayed the current prices and profit levels. Profit each period was calculated as the average profit level over each 40 second cycle. When a new period began, the previous period profit was calculated and shown to participants. This was the same for both treatments. In the discrete treatment the values only updated when a Day or Night started, whereas in the continuous treatment it also updated whenever a firm in the market changed their price. At the end of the experiment a computer randomly selected 10 periods for payment.

The experimental sessions were conducted at the Adelaide Laboratory for Economics Experiments in 2021 using z-Tree (Fischbacher, 2007). Participants were recruited using ORSEE (Greiner, 2015) and were mainly students (undergraduate and postgraduate) of the University of Adelaide. In total, three sessions were run

⁷Although short, the blank screen between periods in z-Tree would have disrupted the flow of time.




Figure 3.2: Example screenshot for the continuous treatment.

for each treatment with a total of 92 participants during October and November 2021, and April 2022. The session size ranged from between 10-24 participants. At the beginning of the experiment participants were given instructions to read privately.⁸ Instructions, including screen shots of the review question and the profit tables, can be found in the Appendix A.3 and A.4. Before the experiment began, participants answered review questions to help them understand the profit table, and played a practice round (with the computer). Each participant was paid, on average, AUD \$37.9.⁹ Each session, including the reading of instructions, practice round, and payment, took between one and a half and two hours.

⁸The instructions were not read aloud to participants so that they could read and process the examples/numbers in the profit table at their own pace. Understanding of profit during the Day and Night was crucial and tested using the control question. Participants could develop additional common understanding during the practice round.

⁹The first two sessions (1 continuous, 1 discrete) were run using an exchange rate of 100 Experimental Currency Units (ECU) = 1 AUD. Participants coordinated more successfully than expected, so the remaining sessions were run with an exchange rate of 130 ECU = 1 AUD. Results remain qualitatively the same when excluding these first two sessions from statistical testing.

3.4 Results

We start by presenting our experimental data on aggregate and then outline and investigate specific findings in the subsections that follow. The primary result of our experiment is that the *continuous* treatment yields higher prices on average, and therefore reduced consumer welfare, compared to *discrete*. Our result is driven by more successful initiation and maintenance of full collusion in *continuous*, as well as higher prices on average in non fully-collusive markets. Thus, we reject our null hypothesis $H1_0$. As prices are higher in *continuous*, the introduction of discrete pricing periods does not lead to more collusion between firms, as was the logic behind hypothesis $H1_a$. We instead adopt alternative hypothesis $H1_b$ that prices are lower with pricing frequency restrictions.

Figure 3.3 shows the average price for each 1 second interval during the experiment (top) and the average period price separately for Day and Night (bottom). Both figures reveal a distinct treatment effect, with the average price for Day and Night in *continuous* approaching p^{JPM} and staying well above the average price in *discrete*. Prices overlap between treatments in the first few periods, but quickly diverge. It is interesting to note the end game effect, which is particularly pronounced for the *continuous* treatment. A timer was included at the top of the screen counting down the time remaining in the experiment. However, a red line indicating the end of the experiment was only visible on the graph in the last Night, and appears to have initiated the unravelling of collusion in the final seconds.

General demographic information about participants in each treatment can be found in Table 3.2. Individual market prices over time in each treatment can be observed in Figure 3.4. A first look at the individual markets reveals a higher proportion of successful joint-profit maximising markets in *continuous*. The dynamics appear to be quite different in *discrete*, with a smaller range of prices in non-JMP markets compared to *discrete*. To play stage-game Nash or joint profit-maximising collusion would require 360 price changes per group (180 per firm) to adapt to de-

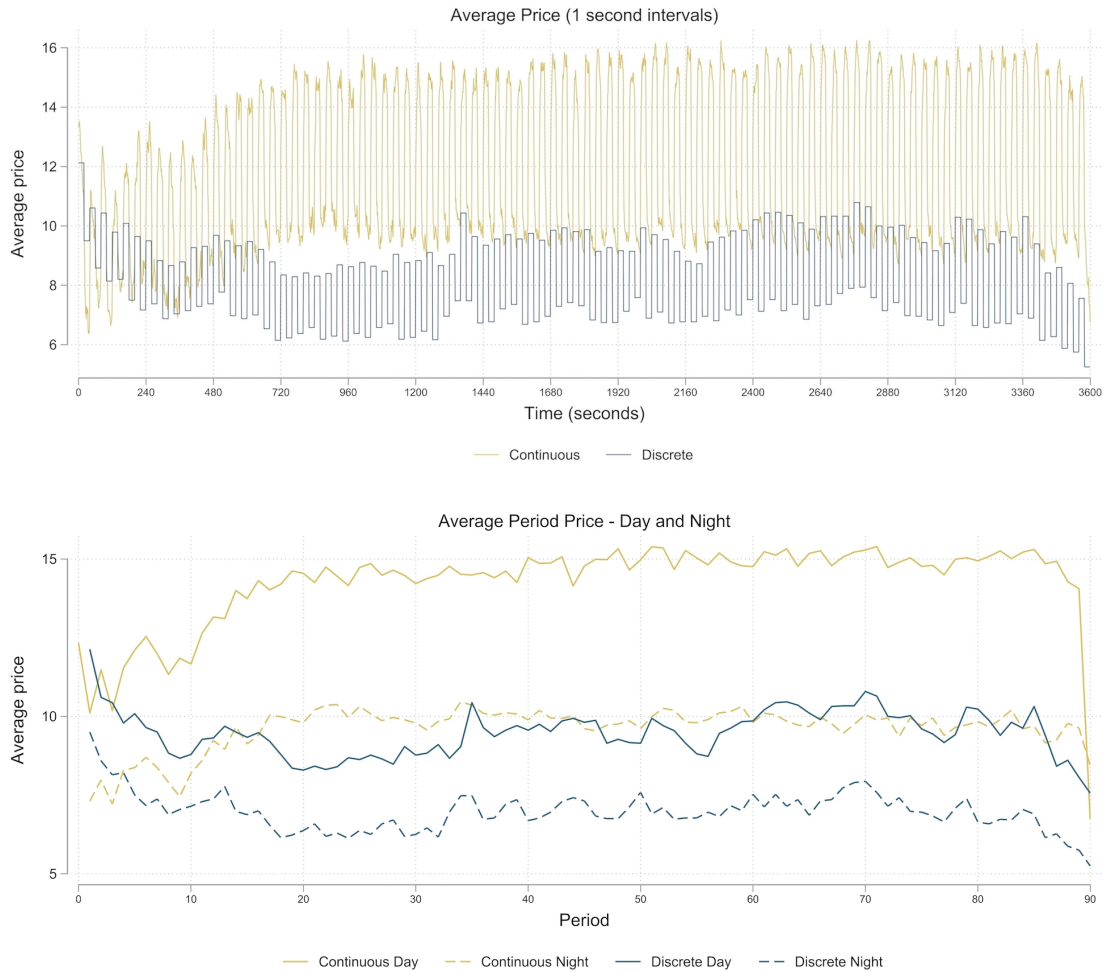


Figure 3.3: Average price over the experiment. Average price per second (top). Average period price, Day and Night (bottom).

mand shocks. For the *continuous* treatment the minimum number of price changes in a group was 3, whilst the maximum was 2,430. The mean number of price changes was 938 (median 710). We look closer at prices within individual markets in what follows.

The continuous interaction of participants means that pricing at a certain point in time is correlated with whatever happened before. This is a challenge for any regression approach. The different decision timing in the treatment also means that there is no appropriate time variable that would allow for dynamic panel analysis. Hence, the only robust and universal test is a simple comparison of prices averaged over the whole experiment and markets across treatments. A non-parametric test like the used Mann-Whitney rank-sum test is extremely robust, as it has the weak-

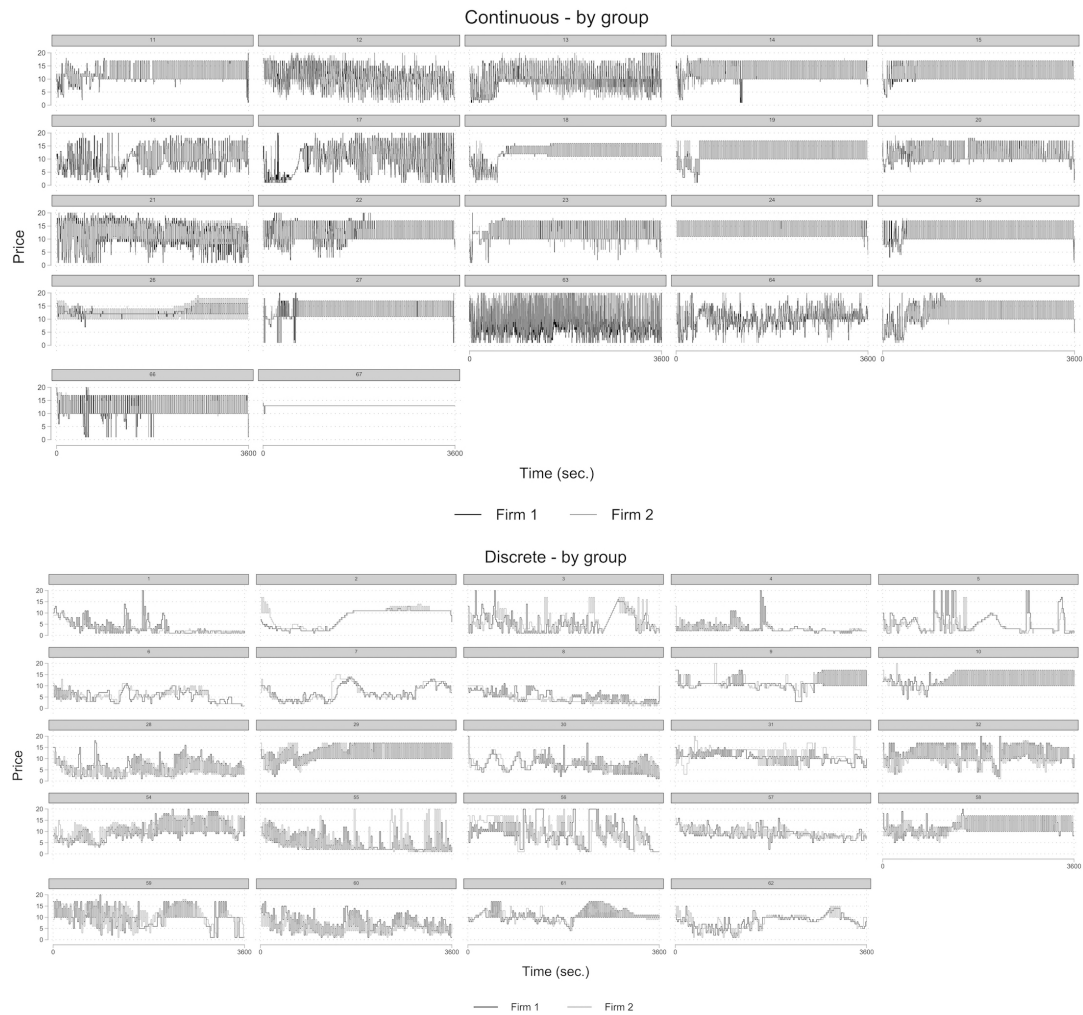


Figure 3.4: Prices in each market.

est assumptions to be valid. Averaging the prices is a perfect way of eliminating problems of correlation within markets, but implies a loss of power, which could be increased by making very restrictive and questionable assumptions allowing the use of a panel data regression. Our effect size is so large that we do not need to do this and the power of the Mann-Whitney rank-sum test is sufficient to pick it up. This is the cleanest and most robust way of establishing our result, and how the following results proceed.

Result 1 *Price levels in continuous are significantly higher than in discrete.*

On average, prices are higher in the *continuous* treatment compared to *discrete*, thus we reject H_{1_0} and adopt H_{1_b} . This also holds separately for day and night

Table 3.2: Demographics across treatments

		Continuous	Discrete
Completed high-school math %		77.27	75.00
Gender %	Female	54.55	50.00
	Male	45.45	50.00
Age group %	18-25 years	50.00	70.83
	26-30 years	27.27	8.33
	31+	22.73	20.83
Home country %	Asia	54.55	29.16
	Australia	22.73	33.33
	Other	22.73	37.50

period prices. Although this is clear from the figures above, we now calculate the average price for each independent observation. Table 3.3 provides summary statistics for the average price of independent observations in each treatment. The average price is 11.91 in the *continuous* treatment, and 8.23 for *discrete*. For reference, the fully joint-profit maximising average period price is 13.5. We conduct a Mann-Whitney rank-sum test for average prices between treatments and reject the null hypothesis that prices are equal ($p = 0.0002$). Note that the mean and median price in *discrete* still remains above stage-game Nash, thus both treatments are considered collusive, despite the stark difference between them.

Table 3.3: Average price of independent observations.

	Mean	Std. Dev.	Min.	Median	Max.	Obs.
Continuous	11.91	1.49	8.27	12.46	13.85	22
Discrete	8.23	2.99	3.09	8.08	12.46	24

To investigate the frequency of different prices in markets, we break down prices into five groups of 4. We calculate the fraction of time spent in each pricing category for Day and Night, as shown in Figures 3.5. In the *continuous* treatment, the distribution of prices is tighter and more concentrated around the p^{JPM} category for both the Day and Night. This indicates that *continuous* markets had higher prices more generally.

Higher levels of collusion in the continuous treatment is consistent with the findings of Friedman and Oprea (2012), where higher levels of cooperation in a

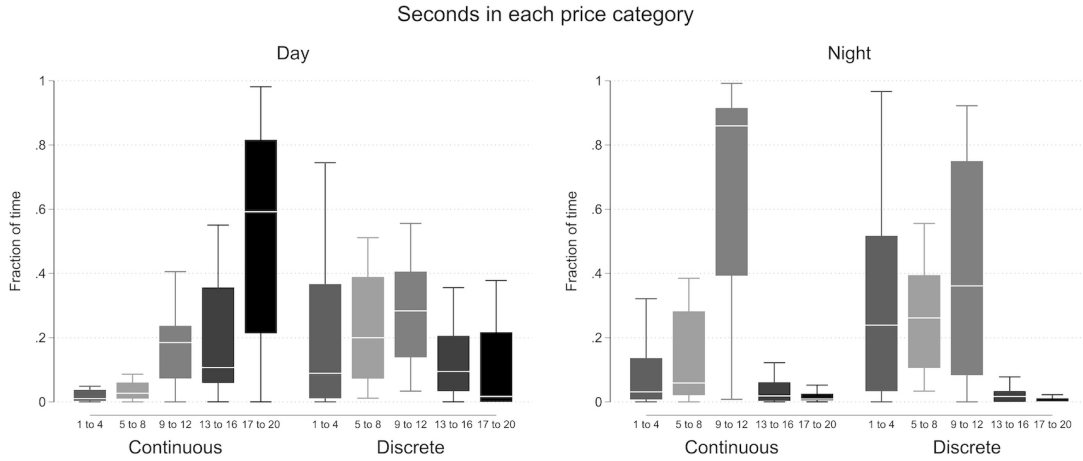


Figure 3.5: Fraction of time (seconds) played by each group in each pricing category.

repeated prisoner’s dilemma are observed in continuous time. Interestingly, our finding is the opposite to that in Horstmann et al. (2016), where more collusion was observed in discrete time oligopolies. Our action space of $[1, 20]$ is much smaller than that of $[0, 100]$ used in their experiment.¹⁰ These results suggest that comparison between continuous and discrete time experiments yield different results, thus strengthening our design of imposing pricing frequency restrictions in an otherwise continuous time market.

Result 2 *Higher fraction of fully collusive markets in continuous.*

Given the findings above, we now focus on the behaviour within individual markets that leads to higher prices in *continuous*. Observing the individual market behaviour in Figure 3.4, it appears that more firms are able to initiate and sustain collusion in *continuous*. We are interested in characterising markets as collusive or otherwise. Defining when a tacitly collusive agreement has been reached is not immediately clear. Although prices above stage-game Nash would indicate a collusive market, they do not guarantee the existence of a tacitly collusive agreement, and instead may indicate experimentation and/or attempts at signalling. Different sets of prices may also be part of a tacitly collusive agreement if a particular pattern of

¹⁰The minimum increment of 1 for the strategy space in Horstmann et al. (2016) is large, however the sliders shown in the display of the experimental software screen show labels in integers of 10.

pricing behaviour is consistently followed by both firms. Defining and identifying strategies (i.e., trigger strategies) used by firms is also not trivial, given the large strategy space.

To solve these issues, we define successful initiation of collusion through the stability of average prices over time, as well as the average price being collusive. First we characterise whether the market price over time is stable. The average price in period t for each market g is calculated. We then create two moving averages, incorporating the current and previous 5 period average prices for \bar{p} , and the current and subsequent 5 periods for \underline{p} . We define a market as stable ($\omega_t=1$) if there is minimal difference between the moving averages, as follows:

$$\omega_{t,g} = \begin{cases} 1 & \text{if } |\bar{p} - \underline{p}| \leq 0.1 \\ 0 & \text{if } |\bar{p} - \underline{p}| > 0.1 \end{cases} \quad (3.4)$$

Table 3.4 shows the summary statistics for our moving average. We believe the choice of 0.1 is a reasonably strict requirement for stability.¹¹

Table 3.4: Summary statistics for moving average $\bar{p} - \underline{p}$.

Treatment	Mean	Min	p25	p50	p75	Max
Continuous	0.59	0	0.03	0.22	0.83	6.84
Discrete	1.21	0	0.29	.83	1.75	8.17

The level of collusiveness in each market is determined using the average period price. Perfectly collusive markets are identified if the average period price is ≥ 13 , just slightly less than full JPM price of 13.5. We are also interested in categorising stable groups that are below this price, but above Nash, and allow an average period price greater than 3, but less than 13. Thus we have classify three categories of market types:

¹¹Reducing the cut-off for our moving average to be ≤ 0.05 results in the following. The fraction of observations classed as stable in each treatment for 0.1 (0.05) is 0.40, 0.13 (0.31, 0.099) for C and D, respectively. The median average period price in the low/unstable category is 11.86, 8.00 (12.13, 8.00).

- Stable Perfect Collusion ($\omega_{g,t} = 1$ & $p_{ave} \geq 13$)
- Stable Imperfect Collusion ($\omega_{g,t} = 1$ & $3 < p_{ave} < 13$)
- Unstable/Low ($\omega_{g,t} = 0$ | $p_{ave} \leq 3$)

Figure 3.6 shows distribution of the count of periods spent in each market category. The medium count of periods spent in Stable Perfect Collusion is much higher for *continuous*. There is minimal time spent in the Stable Imperfect Collusion category, indicating that the majority of markets cooperated perfectly, or not at all. The fraction of Stable Perfectly Colluding markets in *continuous* is above 40% by the half-way point of the experiment. In *discrete*, the fraction remains well below at less than 20% throughout. The end-game effects appear more pronounced using this measure, with a steep decline in Stable Perfectly Collusive markets in the last 10 periods.¹² This proportion still remains above that in *discrete*. Thus, we conclude that there are a higher fraction of markets that successfully joint profit maximised in *continuous*.




Figure 3.6: Fraction of markets within each category in each period.

Result 3 *Higher prices in continuous non-stable markets.*

¹²The length of our experiment means that these end game effects are not significant with respect to the overall results. They are interesting to note, and are consistent with end game effects observed in the literature.

One could suggest that the ease of initiating collusion in the *continuous* treatment has led to a higher fraction of fully collusive markets. We now explore average prices in Unstable/Low markets to determine whether there are differences in price levels between treatments, as well as in the proportion of fully collusive markets. In markets categorised as Unstable/Low, prices remain higher in *continuous*, as shown in Figure 3.7. The distribution of average prices in *discrete* is wider, and the median price lower, than in *continuous*. Thus, even amongst firms that were unable to fully collude, prices remain higher in *continuous*. This implication being that with discrete pricing periods, fewer markets are able to fully collude, and those who do not still have lower prices than in *continuous*.

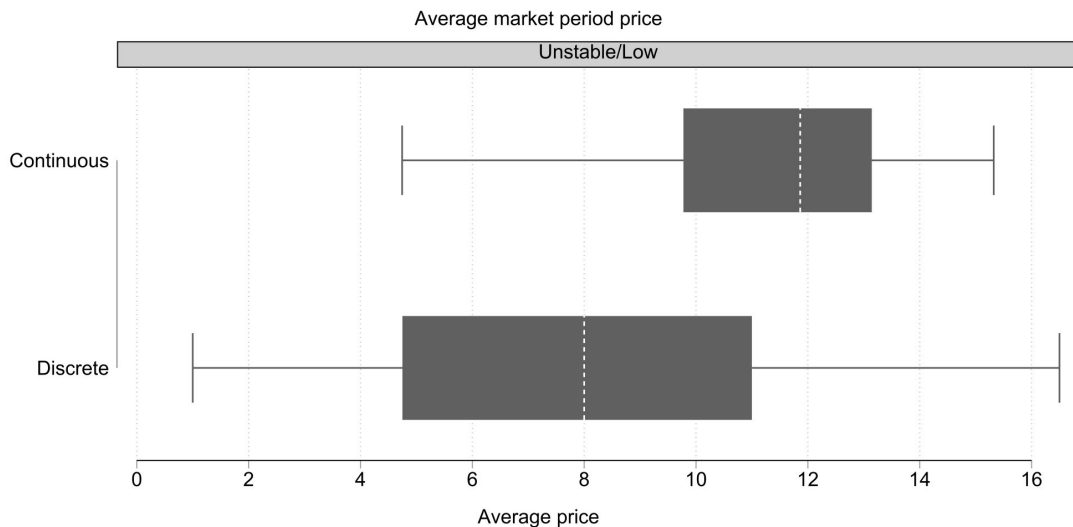


Figure 3.7: Average period price for Unstable/Low markets.

3.5 Conclusion

In this chapter we show how pricing frequency restrictions impact the level of tacit collusion in experimental real-time oligopolies. It is unclear whether implementing discrete pricing periods into an otherwise real-time market will help or hinder the ability of firms to tacitly collude. Collusion incentives and findings from repeated games in continuous time suggest that such restrictions should lower price levels.

On the other hand, introducing clear focal points may make it easier for firms to keep track of, and coordinate, their prices.

Our experimental market consists of two symmetric firms competing in prices over 60 minutes. In our first treatment, firms are able to change their price at any time (computer program updated every 0.5 seconds), asynchronously, and affecting profit immediately. The second treatment implements a new selected price at set timed intervals, otherwise termed near-continuous time by Horstmann et al. (2016). Our experiments are differentiated in two key aspects. The first, through the use of a multinomial-logit demand function (Anderson and de Palma, 1992), allows us to generate a large strategy space between Nash and JPM price within which treatment effects can be clearly identified. A second key difference between these studies and our own is the inclusion of anticipated demand shocks to replicate high and low demand periods (for example, the changes in demand for petrol before school pickup, or late at night). Schelling (1960, p. 74) recognises the importance of qualitatively identifiable signals for tacit communication, thus we include these shocks for firms to use as a coordination device. In the real world, days of the week have always existed, the key difference was the pricing restrictions making them salient in the study by Byrne and de Roos (2019). Mutually selecting symmetric prices above Nash is a Pareto improvement for both firms in our market, i.e., a coordination game with no conflict of interest, thus taking advantage of these focal points should aid tacit collusion (Sitzia and Zheng, 2019). The environment conducive to collusion will hopefully aid future investigation of treatment effects in more competitive markets.

Results show a clear treatment effect, with higher levels of collusion in the *continuous* treatment compared to *discrete*. This result is driven by a higher proportion of markets reaching the fully joint profit maximising price, as well as higher prices in markets that do not. Previous work on continuous time and cooperation has been in relatively short games, with Friedman and Oprea (2012); Leng et al. (2018); Bigoni et al. (2015) using games that last for no more than 60 seconds. Thus, we have shown that even over 3,600 seconds, and in a more complex environment,

that their finding that continuous time supports cooperation still holds.

In the next chapter we explore the underlying mechanism that has led to this result.

Chapter 4

Behavioural mechanism

4.1 Introduction

In the previous chapter we explored how implementing discrete pricing-frequency restrictions impacted collusion in real-time markets, and found a clear treatment effect with higher collusion levels in the treatment with unrestricted price changes. The question remains as to what mechanism is causing this difference. There are two potential explanations. We showed in Chapter 2 that more frequent price changes increased overall collusion incentives. With more price changes, engaging in price leadership, or signalling intention to collude, is less costly. It is quicker for other firms in the market to follow, and to reduce a failed signal. Similarly, the incentive to break a collusive agreement becomes smaller with more frequent price changes as the ‘sucker’ firm can quickly respond and lower their price. With fewer allowable price changes, other firms are forced to endure the lower payoff for the entire price period, whilst the deviator gets to enjoy deviation profits for longer. Thus, the *continuous* treatment has the lower initiation cost and benefit from breaking collusion. Firms are able to react more quickly to signals and any changes in the market. With easier initiation, and lower gain from breaking collusion, more markets successfully collude. Given the higher proportion of fully collusive markets in *continuous*, this

explanation appears feasible.

The second explanation is that something about the nature of continuous time supports cooperation, leading to higher levels of collusion in *continuous* compared to in *discrete*. Leng et al. (2018) find that with full information feedback, continuous time increases average group effort in a minimum effort coordination game. Our approach also implements a full information design where participants observe all prices and profit levels in the market, from which it follows could explain the higher levels of cooperation. Friedman and Oprea (2012) and Bigoni et al. (2015) both show that continuous time games support higher rates of cooperation in a repeated prisoner's dilemma. Friedman and Oprea (2012) find that as the frequency of decisions increases, cooperation rates increase almost linearly, which would support not only a potential behavioural explanation, but is also consistent with the first explanation above. Bigoni et al. (2015) show qualitative differences between continuous and discrete time, unlike Friedman et al. (2015), and are unable to reconcile results with equilibrium theories that could explain differences between stochastic and deterministic ending rules. The explanation behind the higher levels of cooperation observed in continuous time remains unclear, but the common finding is that it supports cooperative outcomes.

In order to discriminate between the effects of collusion incentives and the qualitative aspect of continuous time, we make collusion incentives as close to those in continuous time, while maintaining the recognisable discontinuities introduced by pricing periods. We run an additional treatment, *short-discrete*, that imposes 2.5 second pricing periods. If collusion incentives lead to higher prices in *continuous*, then this treatment will have a similar price level. If the imposition of discrete pricing periods in otherwise real-time markets leads to lower prices, then our new treatment will have price levels similar to *discrete*.

Price levels in the 2.5 second treatment are virtually identical to those in the 20 second treatment. Separating markets into the same categories as we did in Chapter

3, we find comparable proportions of fully collusive markets in *discrete* and *short-discrete*, as well as between average prices in markets that were not. We reject collusion incentives as the driving force behind overall price levels being higher in *continuous*, and conclude that a behavioural mechanism is responsible. A potential explanation is that continuous time implements a status-quo effect for participants (Samuelson and Zeckhauser, 1988), and that discrete pricing periods overcome this bias.

In what follows we outline our new treatment in context of our experimental design from Chapter 3. We explain and state our hypothesis. We then show results, before finally concluding.

4.2 Experimental Design and Hypothesis

We are interested in identifying the mechanism driving the result in Chapter 3. There are two key differences between our primary treatments. The first is the difference in collusion incentives in each treatment. Repeated-game logic tells us that the lower cost of signalling, and gain from deviating from collusion, are lower in *continuous* and therefore in general collusion is more likely than in the *discrete* treatment. The second is a potential behavioural mechanism that arises due to the implementation of discrete pricing periods in an otherwise continuous time market. If instead of collusion incentives, a behavioural mechanism is driving the difference in results, then the introduction of discrete pricing periods may lead to less collusion on average. This hypothesis is derived from the literature on continuous time supporting cooperation (Friedman and Oprea, 2012; Leng et al., 2018; Bigoni et al., 2015). Thus, there are two competing explanations as to why lower prices are observed in the *discrete* treatment. A third treatment, *short-discrete*, is implemented that captures both of these characteristics and allows us to test which factor is responsible.

We lower the price commitment period from 20 seconds to 2.5 seconds, such that in each 40 second cycle firms are able to change their price 16 times. This duration was chosen to be long enough to *feel* discrete, but not be so short such that participants were cognitively unable to identify and respond to changes in the market. We looked at the distribution of time between price changes in the *continuous* treatment and found that the majority of changes occurred within this time frame. Although 2.5 seconds may not seem long, it is sufficient to feel like a distinct pricing period for participants, who are making the same decision over the entire 60 minute game.

Table 4.1: Difference in deviation profit, and loss from signalling the JPM price instead of Nash, for each treatment.

	Continuous	Short-discrete	Discrete
Price commitment	none	2.5 sec	20 sec
Deviation profit AUD	$t_r \times \$0.14$	\$0.34	\$2.7
Signalling loss AUD	$t_r \times \$0.04$	\$0.1	\$0.77

(t_r = minimum reaction time in seconds.)

The cost of signalling and deviating is brought closer to the *continuous* treatment by implementing a 2.5 second interval, as summarised in Table 4.1. For example, assuming the minimum time required for a participant to observe a change and react is t_r seconds, the profit made by deviating in *continuous* is $t_r \times \$0.14$ AUD.¹ For *discrete*, 20 seconds of deviation profit yields a gain of \$2.7 AUD. Thus, the deviation profit in *short-discrete* results in profit of \$0.34 AUD. Signalling an intention to collude by playing the JPM price instead of Nash is decreased from \$0.77 AUD in *discrete* to \$0.10 AUD in *short-discrete*, closer to the *continuous* cost of $t_r \times \$0.04$ AUD. If collusion incentives drive the difference between our two primary treatments, then average prices in *short-discrete* should be closer to *continuous*, $\bar{p}_{SD} = \bar{p}_C$. Thus, our null hypothesis can be stated as follows:

H2₀: *Prices are higher in continuous due to repeated-game logic:*

If this is the case, then there will be no statistically significant difference in prices between *continuous* and *short-discrete* ($p^{\text{short-discrete}} = p^{\text{continuous}}$).

¹Assuming deviation occurs during night, using an exchange rate of 1 AUD per 130 ECU.

If we reject $H2_0$, and price levels are closer to *discrete*, then the introduction of discrete pricing periods is driving the primary result in Chapter 3. In the alternative, we would expect that if discrete pricing periods act as a behavioural prompt, average prices in *short-discrete* will instead be similar to *discrete*, $\bar{p}_{SD} = \bar{p}_D$:

$H2_a$: *Prices are lower in discrete due to a behavioural mechanism:*

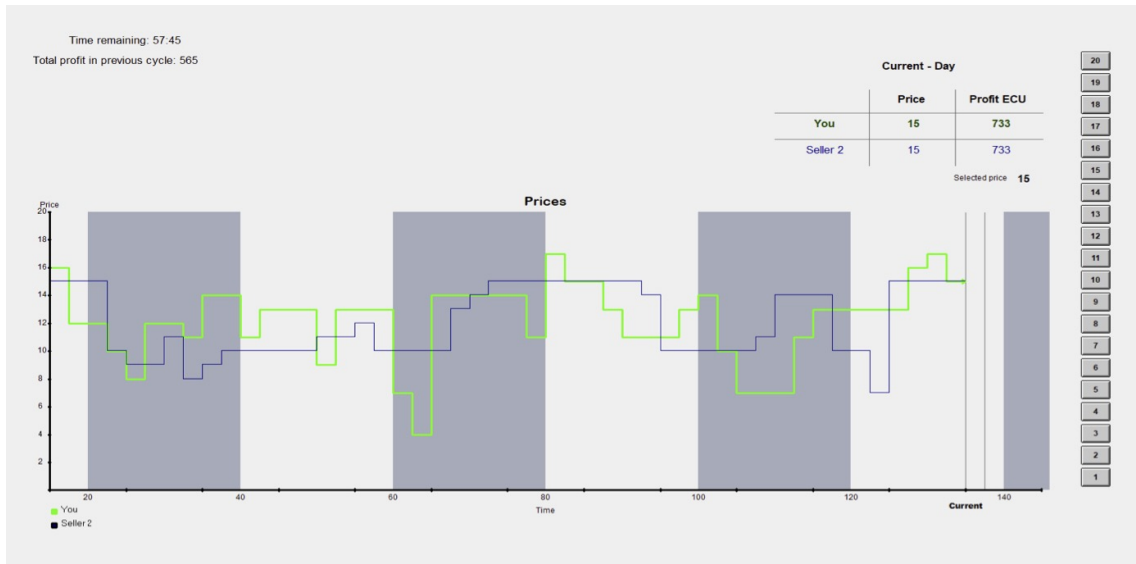


Figure 4.1: Screenshot of Short-Discrete treatment screen shown to participants.

Figure 4.1 shows a screenshot of the screen shown to participants. To further emphasise the discrete pricing periods, a vertical line appeared on the moving price graph shown to participants every 2.5 seconds to count down the time until the next price change was implemented. The moving graph remained the same as in the previous treatments, and emphasised the flow of time. The *short-discrete* treatment was run with an additional 50 participants with between 8-22 per session, in April 2022, paid each on average \$26.6 AUD, and followed the same procedure as the two primary treatments outlined in Chapter 3. The instructions used can be found in Appendix A.5, and control questions used in Appendix A.4.

4.3 Results

We present results in a similar manner to the previous chapter and include previous results for ease of comparison. Overall, *short-discrete* appears to be similar to *discrete*, in terms of the average price over time, the proportion of firms in each market category over time, and price levels for markets unable to fully collude. Thus, we reject the null hypothesis H_{20} and adopt the alternative that prices are closer to *discrete*, indicating that a behavioural factor is driving the difference in collusion levels in Chapter 3. General demographic information can be seen in Table 4.2.

Table 4.2: Demographics across treatments

		Short-discrete
Completed high-school math %		56
Gender %	Female	48
	Male	52
Age group %	18-25 years	76
	26-30 years	4
	31+	20
Home country %	Asia	44
	Australia	32
	Other	24

Result 1 *There is no statistical difference between prices in discrete and short-discrete.*

On aggregate, prices in *short-discrete* are similar to *discrete*. Figure 4.2 plots the average period, Day and Night price for our secondary treatment, with the bottom panel plotting the difference in average period price between *short-discrete*, *continuous* and *discrete*. A question that immediately comes to mind is whether there is an upward trend in the *short-discrete* average period price. It is not straight forward to test this, given the demand shocks, however in the next section we investigate prices over time more closely, and it appears that end game effect on average period price are stronger in *discrete* compared to *short-discrete*, which may result in this difference emerging.

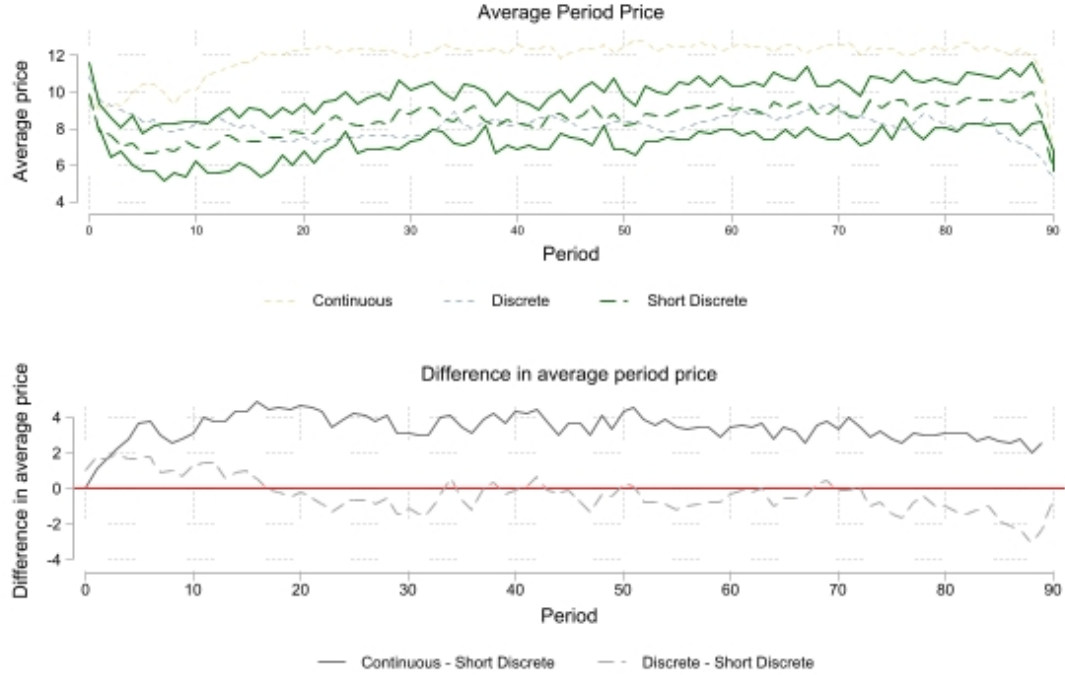


Figure 4.2: The top figure shows the average period price for each treatment (dashed lines), and separately for day and night in *short-discrete* (solid green lines). The bottom figure shows the difference in average period price between treatments.

Table 4.3 shows the summary statistics for the independent observations. The average price in *short-discrete* is significantly different to *continuous* (Mann-Whitney rank sum, $p=0.000$), but not to *discrete* ($p=0.69$). The minimum average price for independent observations in *continuous* is much higher than our two discrete treatments, indicating that the discrete treatments were far more competitive (although still collusive with respect to stage-game Nash).

Table 4.3: Average price of independent observations.

	Mean	Std. Dev.	Min	Max	Obs.
Continuous	11.91	1.49	8.27	13.85	22
Discrete	8.23	2.99	3.09	12.46	24
Short Discrete	8.54	2.88	2.74	13.30	25

We again characterise markets by their level of stability and collusion. Table 4.4 shows the summary statistics for our moving average.²

²Reducing the cut-off for our moving average to be ≤ 0.05 results in the following. The fraction of observations classed as stable in each treatment for 0.1 (0.05) is 0.40, 0.13, 0.16 (0.31, 0.099, 0.099) for C, D and SD, respectively. The median average period price in the low/unstable category is 11.86, 8.00, and 7.94 (12.13, 8, 9.09).

Table 4.4: Summary statistics for moving average $\bar{p} - p$

Treatment	Mean	Min	p25	p50	p75	Max
Continuous	0.59	0	0.03	0.22	0.83	6.84
Discrete	1.21	0	0.29	.83	1.75	8.17
Short Discrete	1.06	0	0.18	.65	1.52	11.52

Figure 4.3 shows distribution of the count of periods spent in each market category. The medium count of periods spent in Stable Perfect Collusion is much higher for *continuous*. For *discrete* and *short-discrete* the median count of periods being Unstable/Low are similar and high. There is minimal time spent in the Stable Imperfect Collusion category, indicating that the majority of markets that did manage to cooperate did so fully, and the remainder were unable to do so at all.

Figure 4.5 plots the fraction of groups in each category over time in each treatment. There is a stark difference between the *continuous* and two discrete treatments. By half-way through the experiment around half of *continuous* markets were perfectly colluding. The majority of time in the discrete treatments is spent in the Unstable/Low categories, confirming that the count distribution in each category reflects what happened over the entire experiment. Overall, there are a higher fraction of Stable Fully Collusive markets in *continuous* for almost the entire 90 periods of the experiment.

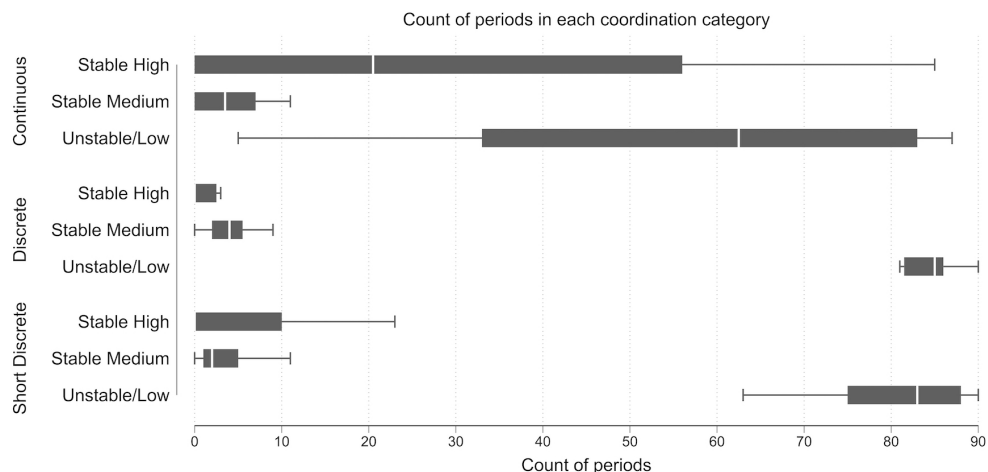


Figure 4.3: Average market period price in each category.

Individual markets are shown in Figure 4.4. Despite the similarities in aggregate

prices between *discrete* and *short-discrete*, the overall dynamics observed in markets not perfectly colluding appears to be more similar to markets in *continuous*. The number of available price changes is therefore not impacting collusion levels. In Chapter 3 it could be claimed that with more price changes, eventually more markets would successfully collude. The results of *short-discrete* appear to go against this.

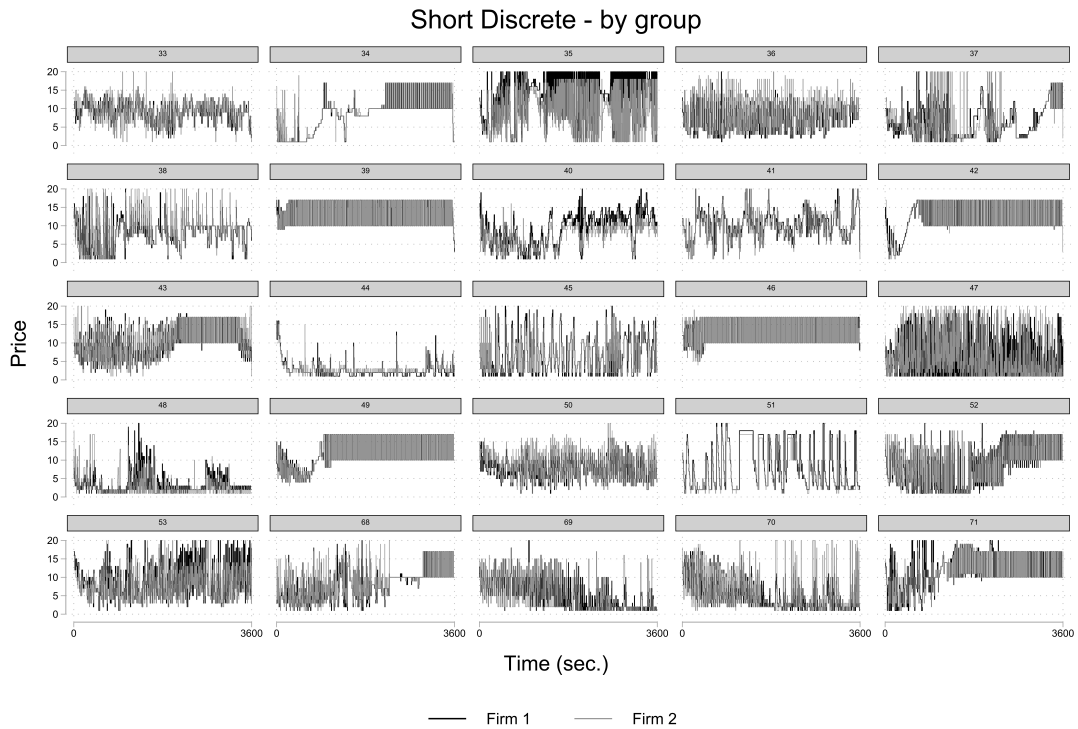


Figure 4.4: Prices in each market.

Result 2 *Discrete treatments have lower prices in Unstable/Low markets.*

We now investigate the price level in markets that were unable to perfectly collude. Following the argument that it is easier to initiate and maintain tacit collusion in *continuous*, is there a treatment effect with respect to price levels in those that were unable to do so? Figure 4.6 shows the distribution of prices in each treatment for the Unstable/Low category. The median average period price is almost identical in the two discrete treatments, while in *continuous* it is sitting slightly below the joint profit maximising average period price.

To see how this might change over the experiment, we plot the average period

price for each treatment over time in Figure 4.7. The starting points in each treatment are similar, with prices diverging upwards in *continuous* and downwards in *discrete*. There does appear to be a slight upward trend in average price in *short-discrete*, however it remains well below that in *continuous*.

Rank-sum tests on the fraction of time each market spends in each category between treatments confirms this result. Comparing *continuous* with *discrete* and again with *short-discrete* confirms a significantly higher time spent in Stable Perfect Collusion by *continuous* markets ($p < 0.05$). There is no statistically significant difference between the two discrete treatments with respect to the fraction of time spent in Stable Perfect Collusion.




Figure 4.5: Fraction of markets in each category over time.

Figure 4.6 shows the distribution of prices in each treatment for the Unstable/Low category, and in Figure 4.7 the average period price in the Unstable/Low category over the experiment. The median price in both discrete treatments is similar and significantly lower than markets in *continuous*. This shows that even in markets that were not perfectly colluding, there is less collusion with pricing-frequency restrictions.

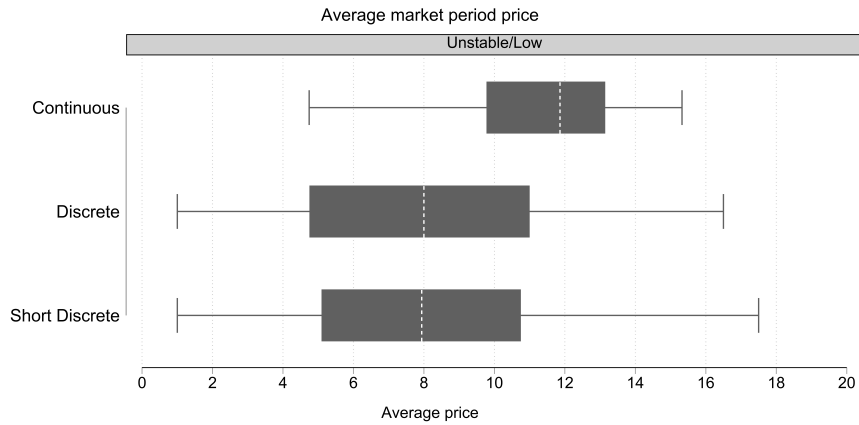


Figure 4.6: Distribution of average period prices in markets that are Unstable/Low.




Figure 4.7: Average period price for markets that are Unstable/Low, over time.

4.4 Conclusion

We have shown that restrictions on pricing frequency can reduce tacit collusion in otherwise real-time environments. Discrete pricing periods act as a behavioural prompt for participants to re-strategise for the next pricing period, resulting in a lower proportion of fully-collusive markets, and a lower price level more generally.

Which behavioural mechanism is responsible for this effect requires further investigation. Our findings are opposite to those by Horstmann et al. (2016), who are the only other study that investigates oligopoly markets in continuous time. The reason for this is unclear. Similarly, our finding that cooperation levels do not linearly increase with reductions in grid size, as found by Friedman and Oprea (2012),

is also notable. Our model is more complex than a simple prisoner's dilemma, and we also run our game for 3,600 seconds, rather than 60 seconds. The results in this thesis do not support repeated-game logic or collusion incentives driving cooperation levels.

One potential behavioural explanation for higher cooperation levels in *continuous* could be a status quo effect (Samuelson and Zeckhauser, 1988). The introduction of pricing periods appears to have prompted the decay in cooperation levels on aggregate. Although it is difficult to observe in the individual markets, it is interesting to note how prices jump up at the start of each Day, before drifting down. This is similar to the within-period competitive effect we observed in Chapter 2. Direct comparison of results from Chapter 2 and those in Chapters 3 and 4 should be made with caution. The markets in each experiment are characteristically different, including the number of firms competing, the length of interaction between firms, and the overall collusion incentives derived from the underlying model used. Nonetheless, the qualitative differences between these experiments highlights the need for experimentalists to be cautious in choosing their timing protocol, which may impact behaviour in an unexpected manner.

Chapter 5

Conclusion

Tacit collusion in markets is a key issue for policy makers in designing regulations to support consumer welfare. Retail petrol markets are a common target for pricing regulations in many jurisdictions around the world. A common regulation supports pricing transparency for consumers, ensuring up-to-date prices are easily found online, thus reducing consumer search costs and supposedly promoting competition. A second, unintended, consequence of these regulations is that communication between firms becomes cheaper - checking your competitors' prices is now a click away. The regulations explored in this thesis are based on those in the WA retail petrol market, where firms are restricted to simultaneously announcing their prices that are then fixed for 24 hours. The limitations of theory on this topic, and complexities of real-world markets, prompted us to experimentally investigate this question. We have contributed to the experimental literature on repeated games, cooperation in continuous time, and competition policy.

We designed experiments to test how interaction frequency within a stage game impacts cooperation levels. In Chapter 2 we embedded behavioural hypotheses within the context of a repeated game, thus bridging the gap between the two areas of the literature. Results support that the frequency of price changes impacts collusion. Collusion levels were higher with only one price change per period, compared to the

level in the treatment that allowed for four changes. Cooperation levels appeared to decay over time in all treatments. It is interesting to note the choice of focal point and restart effect, which was the start of a supergame for one price change, or the start of each period in the treatments with 4. The within-period competition led to overall lower average collusion levels.

In Chapter 3 we tested whether interaction frequency impacts cooperation in real-time markets. Results from a 60 minute duopoly experiment confirm that continuous time can support cooperation, even in more complex environments and over a long period of time. With unrestricted price changes, more markets successfully initiated and maintained full collusion. In addition, prices were higher in markets that were not able to do so. At first glance this result appears to contradict the results from Chapter 2. More available price changes in our unrestricted treatment lead to higher levels of collusion, while higher collusion levels were observed in Chapter 2 with only one price change per period. In Chapter 4 we test for the mechanism behind our finding in Chapter 3.

In Chapter 4 we run an additional treatment, reducing the pricing period from 20 seconds down to 2.5 seconds. This allows us to determine whether collusion incentives, i.e., the cost of signalling and gain from deviating, lead to lower prices in our 20 second treatment, or whether the discrete pricing periods themselves are behaviourally relevant. Price levels in the 2.5 second treatment are virtually identical to those in the 20 second treatment. Thus, we reject that collusion incentives impact collusion levels, and instead conclude that discrete pricing periods are prompting firms to compete. This explains the seemingly contrary results between Chapters 2 and 3, as more pricing periods in Chapter 2 led to more opportunities to undercut.

Further research is still required in this area. Our results support the previous findings of Friedman and Oprea (2012); Bigoni et al. (2015); Leng et al. (2018) that continuous time supports cooperation. However, Horstmann et al. (2016) run experiments to test the impact of discrete and continuous time in an oligopoly context

and find results opposite to ours. The behavioural mechanism identified in Chapter 4 requires further investigation. We hypothesise that the status quo effect of continuous time is broken with the introduction of discrete decision periods. It is interesting that a prompt leads to undercutting, as opposed to attempts to raise prices. This could be explained as players simply best responding to the previous price, which in context of the stage game is to lower their price.

Another area that warrants future research is on the impact of Artificial Intelligence (AI) and pricing algorithms. It has been shown that algorithms are able to reach collusive outcomes in both simultaneous and sequential environments (Calvano et al., 2020; Klein, 2021). Calvano et al. (2020) run experiments with AI subjects using a similar model to ours in Chapters 3 and 4, and find that a collusive outcome is indeed possible and robust to various market characteristics. Interestingly, the introduction of additional AI firms does not appear to reduce profit levels as quickly as has been observed with human participants (Horstmann et al., 2018; Potters and Suetens, 2013). Thus questions arise as to how pricing frequency regulations impact cooperation in the presence of at least one firm setting prices using AI. Our finding that discrete pricing periods, through a behavioural mechanism, drives down the price level may not apply to AI.

Appendices

A.1 Chapter 2 - Instructions

Chapter 2 instructions for treatment $n4\delta_h$:

General information

Welcome to the experiment. Please read these instructions carefully. If you have a question please raise your hand. Please switch off your mobile phone and do not communicate with other participants.

During this experiment you can earn Experimental Currency Units (ECU) which will be exchanged into Australian dollars. You will receive payment at the end of the experiment, privately, in cash. The exchange rate will be given later in the instructions.

As you read through these instructions you will be prompted to answer questions displayed on your screen. These questions are designed to help you understand the game. If your answer is incorrect you may try again. If you are unsure at any time, please raise your hand. By answering these questions you will earn 10 AUD in addition to the profit you earn in the game.

There will be a practice game (competing with a computer) before the experiment begins.

Your task

In today's session all participants are sellers of a fictitious product. Your objective is to set and adjust your price such that you make as much profit as possible. In each game you will compete with two other sellers, who also set prices. A game consists of multiple days (which last for 60 seconds each).

Before the first day begins you will be asked to set your price, which you can change during the day. Your profit for the day depends on the prices you and your competitors charge throughout the day.

At the end of each day you will see your final profit. It is randomly determined if a new day starts, or if the game ends.

How to play

When a new game starts you will see a screen asking you to select your starting price. Possible prices are whole numbers between 1 and 7. A day is divided into four periods of 15 seconds each.

The day begins with period 1, and you will observe your price and profit, as well as that of the other two sellers you are competing with. If you change your price, it will be effective in the next period. For example, if you select a new price during period 1, it will become effective in period 2. If you select a new price during period 2 it will be effective in period 3, and so on.

Please complete Question 1 on your screen.

A screenshot of the day screen is shown in Figure 1. At the top of the screen there is information about the time remaining in the current period (countdown from 15 seconds to 0 seconds).

The top left-hand side of the screen shows a table with the current period's prices and profits for you and your competitors. On the right-hand side of the screen are two graphs showing the day's profit (top) and price (bottom) levels. The table and graphs will update automatically when each new period begins.

At the bottom left of your screen your selected price is displayed. There you can also change your price. Once the timer reaches zero, the price you have currently selected will be entered as your price for the next period.

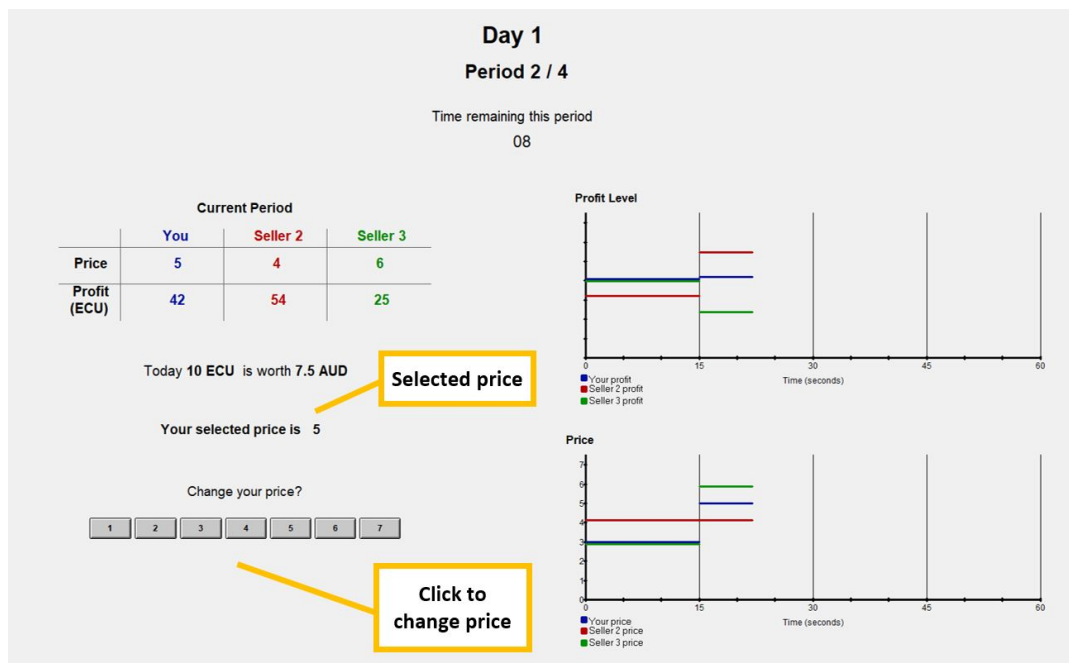


Figure 1: Screenshot of day screen

At the end of the day you will see an overnight screen, as shown in Figure 2. The screen summarises the prices and day profit for you and the other sellers you are competing with. You also have the option to change your starting price for the following day, which will take effect once the new day begins. If you don't change your price, the last price selected will be your starting price for period 1 tomorrow.

Number of days in each game

The number of days in each game is randomly determined and unknown until the game has ended. As seen in Figure 2, the spinning of a computerised wheel determines if the game continues. After each day there is a 20% probability the game will end. This means there is an 80% probability that the game will continue to the next day.

If the game continues, you will compete with the same participants.

If the game ends, and there is enough time for another game, you will be rematched with two other participants and a new game begins. In case that there is not enough time in the session for another game, the experiment ends.

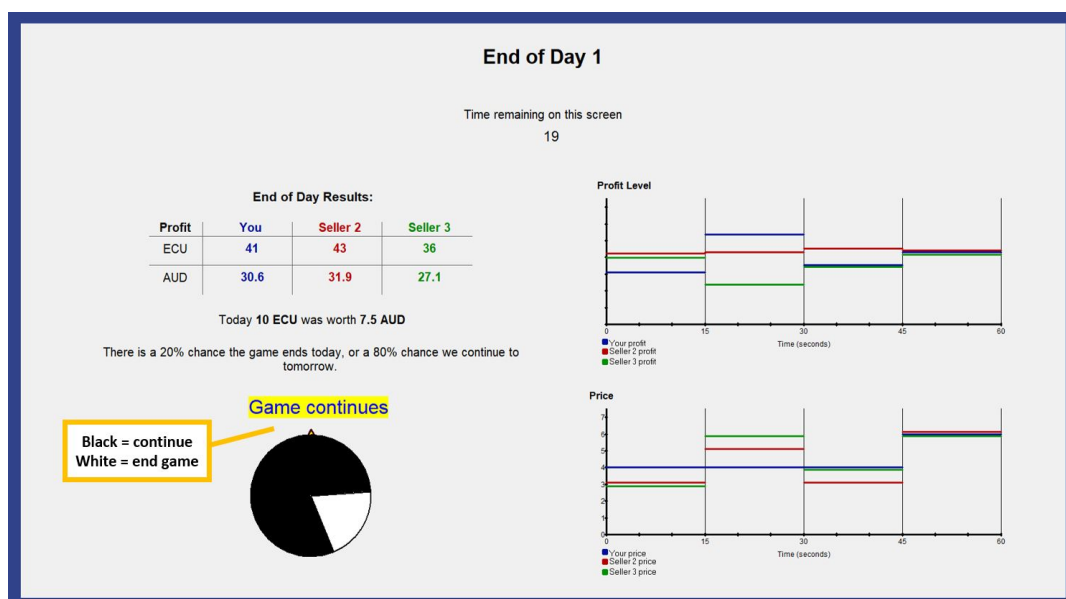


Figure 2: Screenshot of overnight screen

Profit

Below we explain how period profit is calculated, and then how day profit is determined.

Period profit - the profit table

How much you are able to sell depends on your price, and the average price of your competitors (Seller 2 and Seller 3). Profit, for each combination of prices, is shown in the attached Profit Table.

Example of how to look up profit

Let's look at an example of how to use the Profit Table. The table is replicated below in Table 1. Your price is 3. Seller 2's price is 4. Seller 3's price is 6:

You -

- the average of your competitors prices is $(4+6)/2 = 5$. If you set your price to 3, and the average of the other firms' price is 5, your profit for the period is 52 ECU. This is coloured blue in Table 1.

Seller 2 -

- the average price of you and Seller 3 is $(3+6)/2=4.5$. If Seller 2's price is 4 and the average price is 4.5, Seller 2's profit is 45 ECU. This is coloured red in Table 1.

Seller 3 -

- the average of your price and Seller 2's price is $(3+4)/2=3.5$. If Seller 3's price is 6 and the average price of others is 3.5, Seller 3's profit is 13 ECU. This is coloured green in Table 1.

Now complete Question 2 on your screen.

Table 1: Profit table example (ECU)

		Average of other sellers' prices												
		1	1.5	2	2.5	3	3.5	4	4.5	5	5.5	6	6.5	7
Price	1	25	29	32	35	37	40	42	45	47	47	47	47	47
	2	26	30	33	36	39	42	45	48	51	54	56	56	56
	3	22	26	29	33	37	41	44	48	52	56	60	62	63
	4	13	18	22	27	31	36	40	45	49	54	58	63	67
	5	0	6	11	16	21	27	32	37	42	48	53	58	63
	6	0	0	0	4	7	13	19	25	31	37	43	49	55
	7	0	0	0	0	0	2	3	10	16	23	29	36	42

How does changing your price impact your profit?

Increasing your price impacts your profit in two ways. A higher price means you receive more money per unit sold, positively impacting your profit. However, there is also a negative effect, as a higher price means you will sell a lower quantity. Therefore, raising your price may increase or decrease your profit:

- Increasing a low price will increase your profit - For example, if your price is 1 and the average price of your competitors is 5, your profit is 47 ECU. If you were to raise your price to 3, your profit increases to 52 ECU.
- Increasing a high price will decrease your profit - For example, if your price is 6 and the average price of your competitors is 5, your profit is 31 ECU. If you were to raise your price to 7, your profit decreases to 16 ECU.

This is true for any given average price set by your competitors. The same logic can be applied to the impact of decreasing your price. Reducing a low price will decrease your profit, but reducing a high price will increase your profit.

You can see how changing your price (given the average price of your competitors) impacts your profit by looking down each column in the Profit Table.

How do your competitor's prices impact your profit?

If your competitors increase their prices (and yours is held constant), then you sell more and your profit increases. If your competitors decrease their prices, you sell less and your profit decreases.

For example, if you set your price to 5 and your competitors average price is also 5, your profit is 42 ECU. If the average of your competitors price decreases to 4.5, your profit is 37 ECU. If their average price increases to 5.5, your profit is 48.

You can see how changes in your competitors average price impacts your profit (for your given price) by looking across each row in the Profit Table.

Day profit - average of period profits

On the day screen, period profit is displayed. Your final profit for the day is shown on the overnight screen. Overall, your profit for the day is the average of your period profits.

For example, if your profit in each period during the day is (30, 30, 20, 40 ECUs), adding these values up gives us 120 ECU. Then we take the average by dividing 120 ECU by 4. Your final profit for the day is therefore 30 ECU.

Exchange rate & final payment

The value of profit will decrease by approximately 9% with every new day. The starting exchange rate is 1 ECU = 0.75 AUD. The exchange rate can be found in Table 2 for the first 10 days (also displayed on the day screen).

For example, on the first day 10 ECU is worth 7.5 AUD. On the second day, 10 ECU is worth 6.8 AUD, and so on. Let's say your day profit is 40 ECU. If the game ends on day 1, your payment for this game would be 30 AUD. However, if the game ends on day 6, your payment would be 18.6 AUD. Each new game will start again from day 1.

Please complete Question 3 on your screen.

Table 2: Exchange rate - The amount of AUD received for 10 ECU

Day	1	2	3	4	5	6	7	8	9	10	...
AUD per 10 ECU	7.50	6.82	6.20	5.63	5.12	4.66	4.23	3.85	3.50	3.18	...

At the end of the experiment a random game will be selected. You will be paid only for profit (in AUD) you earn on the last day before this game ended.

Please complete Question 4 on your screen.

Closing

After the last game of the experiment a questionnaire will appear on your screen for you to fill out before you can receive your payment. Please remain seated and silent.

Please fill out and sign your receipt form. The experimenter will then provide instructions for you to receive your payment.

Chapter 2 instructions for treatment n4 δ _i:

General information

Welcome to the experiment. Please read these instructions carefully. If you have a question please raise your hand. Please switch off your mobile phone and do not communicate with other participants.

During this experiment you can earn Experimental Currency Units (ECU) which will be exchanged into Australian dollars. You will receive payment at the end of the experiment, privately, in cash. The exchange rate will be given later in the instructions.

As you read through these instructions you will be prompted to answer questions displayed on your screen. These questions are designed to help you understand the game. If your answer is incorrect you may try again. If you are unsure at any time, please raise your hand. By answering these questions you will earn 10 AUD in addition to the profit you earn in the game.

There will be a practice game (competing with a computer) before the experiment begins.

Your task

In today's session all participants are sellers of a fictitious product. Your objective is to set and adjust your price such that you make as much profit as possible. In each game you will compete with two other sellers, who also set prices. A game consists of multiple days (which last for 60 seconds each).

Before the first day begins you will be asked to set your price, which you can change during the day. Your profit for the day depends on the prices you and your competitors charge throughout the day.

At the end of each day you will see your final profit. It is randomly determined if a new day starts, or if the game ends.

How to play

When a new game starts you will see a screen asking you to select your starting price. Possible prices are whole numbers between 1 and 7. A day is divided into four periods of 15 seconds each.

The day begins with period 1, and you will observe your price and profit, as well as that of the other two sellers you are competing with. If you change your price, it will be effective in the next period. For example, if you select a new price during period 1, it will become effective in period 2. If you select a new price during period 2 it will be effective in period 3, and so on.

Please complete Question 1 on your screen.

A screenshot of the day screen is shown in Figure 1. At the top of the screen there is information about the time remaining in the current period (countdown from 15 seconds to 0 seconds).

The top left-hand side of the screen shows a table with the current period's prices and profits for you and your competitors. On the right-hand side of the screen are two graphs showing the day's profit (top) and prices (bottom) levels. The table and graphs will update automatically when each new period begins.

At the bottom left of your screen your selected price is displayed. There you can also change your price. Once the timer reaches zero, the price you have currently selected will be entered as your price for the next period.

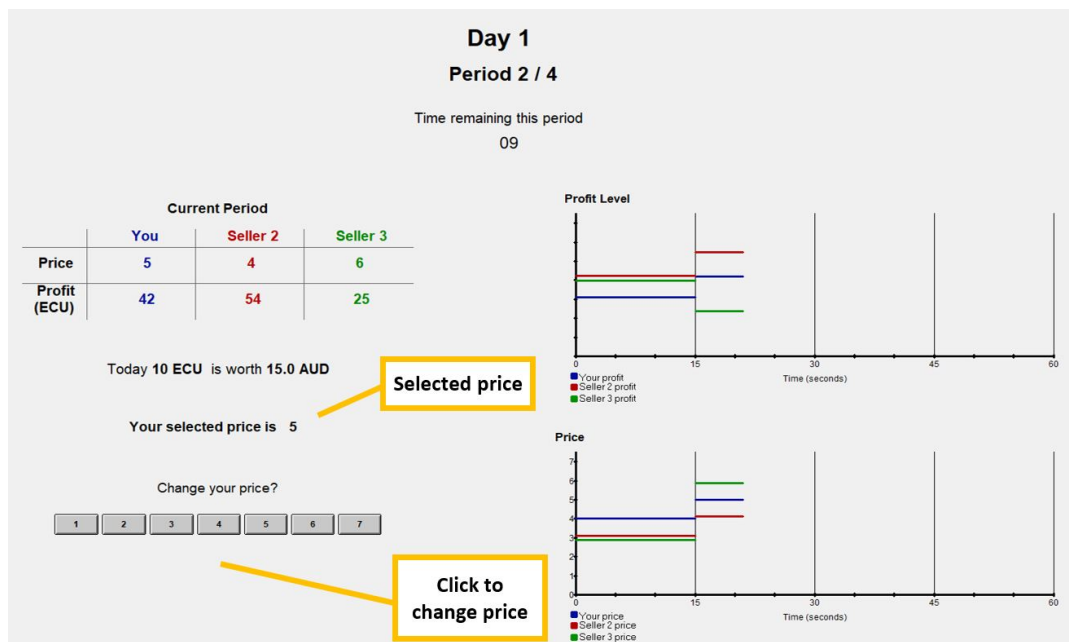


Figure 1: Screenshot of day screen

At the end of the day you will see an overnight screen, as shown in Figure 2. The screen summarises the prices and day profit for you and the other sellers you are competing with. You also have the option to change your starting price for the following day, which will take effect once the new day begins. If you don't change your price, the last price selected will be your starting price for period 1 tomorrow.

Number of days in each game

The number of days in each game is randomly determined and unknown until the game has ended. As seen in Figure 2, the spinning of a computerised wheel determines if the game continues. After each day there is a 20% probability the game will end. This means there is an 80% probability that the game will continue to the next day.

If the game continues, you will compete with the same participants.

If the game ends, and there is enough time for another game, you will be rematched with two other participants and a new game begins. In case that there is not enough time in the session for another game, the experiment ends.

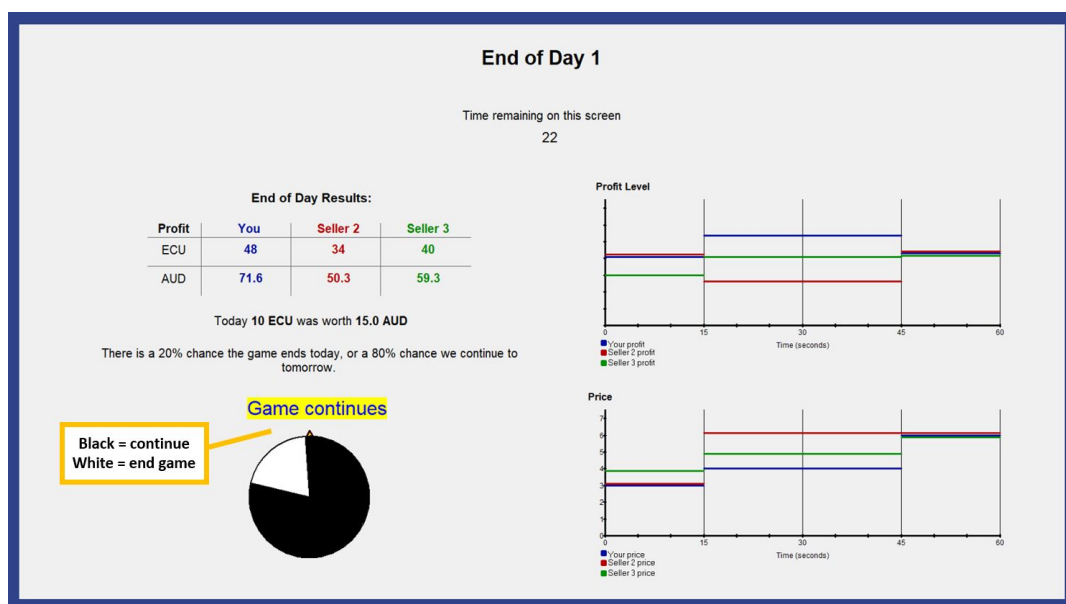


Figure 2: Screenshot of overnight screen

Profit

Below we explain how period profit is calculated, and then how day profit is determined.

Period profit - the profit table

How much you are able to sell depends on your price, and the average price of your competitors (Seller 2 and Seller 3). Profit, for each combination of prices, is shown in the attached Profit Table.

Example of how to look up profit

Let's look at an example of how to use the Profit Table. The table is replicated below in Table 1. Your price is 3. Seller 2's price is 4. Seller 3's price is 6:

You -

- the average of your competitors prices is $(4+6)/2 = 5$. If you set your price to 3, and the average of the other firms' price is 5, your profit for the period is 52 ECU. This is coloured blue in Table 1.

Seller 2 -

- the average price of you and Seller 3 is $(3+6)/2=4.5$. If Seller 2's price is 4 and the average price is 4.5, Seller 2's profit is 45 ECU. This is coloured red in Table 1.

Seller 3 -

- the average of your price and Seller 2's price is $(3+4)/2=3.5$. If Seller 3's price is 6 and the average price of others is 3.5, Seller 3's profit is 13 ECU. This is coloured green in Table 1.

Now complete Question 2 on your screen.

Table 1: Profit table example (ECU)

		Average of other sellers' prices												
		1	1.5	2	2.5	3	3.5	4	4.5	5	5.5	6	6.5	7
Price	1	25	29	32	35	37	40	42	45	47	47	47	47	47
	2	26	30	33	36	39	42	45	48	51	54	56	56	56
	3	22	26	29	33	37	41	44	48	52	56	60	62	63
	4	13	18	22	27	31	36	40	45	49	54	58	63	67
	5	0	6	11	16	21	27	32	37	42	48	53	58	63
	6	0	0	0	4	7	13	19	25	31	37	43	49	55
	7	0	0	0	0	0	2	3	10	16	23	29	36	42

How does changing your price impact your profit?

Increasing your price impacts your profit in two ways. A higher price means you receive more money per unit sold, positively impacting your profit. However, there is also a negative effect, as a higher price means you will sell a lower quantity. Therefore, raising your price may increase or decrease your profit:

- Increasing a low price will increase your profit - For example, if your price is 1 and the average price of your competitors is 5, your profit is 47 ECU. If you were to raise your price to 3, your profit increases to 52 ECU.
- Increasing a high price will decrease your profit - For example, if your price is 6 and the average price of your competitors is 5, your profit is 31 ECU. If you were to raise your price to 7, your profit decreases to 16 ECU.

This is true for any given average price set by your competitors. The same logic can be applied to the impact of decreasing your price. Reducing a low price will decrease your profit, but reducing a high price will increase your profit.

You can see how changing your price (given the average price of your competitors) impacts your profit by looking down each column in the Profit Table.

How do your competitor's prices impact your profit?

If your competitors increase their prices (and yours is held constant), then you sell more and your profit increases. If your competitors decrease their prices, you sell less and your profit decreases.

For example: if you set your price to 5 and your competitors average price is also 5, your profit is 42 ECU. If the average of your competitors price decreases to 4.5, your profit is 37 ECU. If their average price increases to 5.5, your profit is 48.

You can see how changes in your competitors average price impacts your profit (for your given price) by looking across each row in the Profit Table.

Day profit - average of period profits

On the day screen, period profit is displayed. Your final profit for the day is shown on the overnight screen. Overall, your profit for the day is the average of your period profits.

For example:

If your profit in each period during the day is (30, 30, 20, 40 ECUs), adding these values up gives us 120 ECU. Then we take the average by dividing 120 ECU by 4. Your final profit for the day is therefore 30 ECU.

Exchange rate & final payment

The value of profit will decrease by approximately 22% with every new day. The starting exchange rate is 1 ECU = 1.5 AUD. The exchange rate can be found in Table 2 for the first 10 days (also displayed on the day screen).

For example, on the first day 10 ECU is worth 15 AUD. On the second day, 10 ECU is worth 11.6 AUD, and so on. Let's say your day profit is 40 ECU. If the game ends on day 1, your payment for this game would be 60 AUD. However, if the game ends on day 6, your payment would be 16.8 AUD. Each new game will start again from day 1.

Please complete Question 3 on your screen.

Table 2: Exchange rate - The amount of AUD received for 10 ECU

Day	1	2	3	4	5	6	7	8	9	10
AUD per 10 ECU	15.00	11.63	9.01	6.99	5.42	4.20	3.26	2.52	1.96	1.52

At the end of the experiment a random game will be selected. You will be paid only for profit (in AUD) you earn on the last day before this game ended.

Please complete Question 4 on your screen.

Closing

After the last game of the experiment a questionnaire will appear on your screen for you to fill out before you can receive your payment. Please remain seated and silent.

Please fill out and sign your receipt form. The experimenter will then provide instructions for you to receive your payment.

Chapter 2 instructions for treatment n1 δ_h :

General information

Welcome to the experiment. Please read these instructions carefully. If you have a question please raise your hand. Please switch off your mobile phone and do not communicate with other participants.

During this experiment you can earn Experimental Currency Units (ECU) which will be exchanged into Australian dollars. You will receive payment at the end of the experiment, privately, in cash. The exchange rate will be given later in the instructions.

As you read through these instructions you will be prompted to answer questions displayed on your screen. These questions are designed to help you understand the game. If your answer is incorrect you may try again. If you are unsure at any time, please raise your hand. By answering these questions you will earn 10 AUD in addition to the profit you earn in the game.

There will be a practice game (competing with a computer) before the experiment begins.

Your task

In today's session all participants are sellers of a fictitious product. Your objective is to set and adjust your price such that you make as much profit as possible. In each game you will compete with two other sellers, who also set prices. A game consists of multiple days (which last for 60 seconds each).

Before the first day begins you will be asked to set your price. Your profit for the day depends on the price charged by you and your competitors during the day.

At the end of each day you will see your final profit. It is randomly determined if a new day starts, or if the game ends.

How to play

When a new game starts you will see a screen asking you to select your starting price. Possible prices are whole numbers between 1 and 7.

When the day begins you will observe your price and profit, as well as that of the other two sellers you are competing with.

Please complete Question 1 on your screen.

A screenshot of the day screen is shown in Figure 1. At the top of the screen there is information about the time remaining in the current day (countdown from 60 seconds to 0 seconds).

The top left-hand side of the screen shows a table with the current period's prices and profits for you and your competitors. On the right-hand side of the screen are two graphs showing the day's profit (top) and price (bottom) levels.

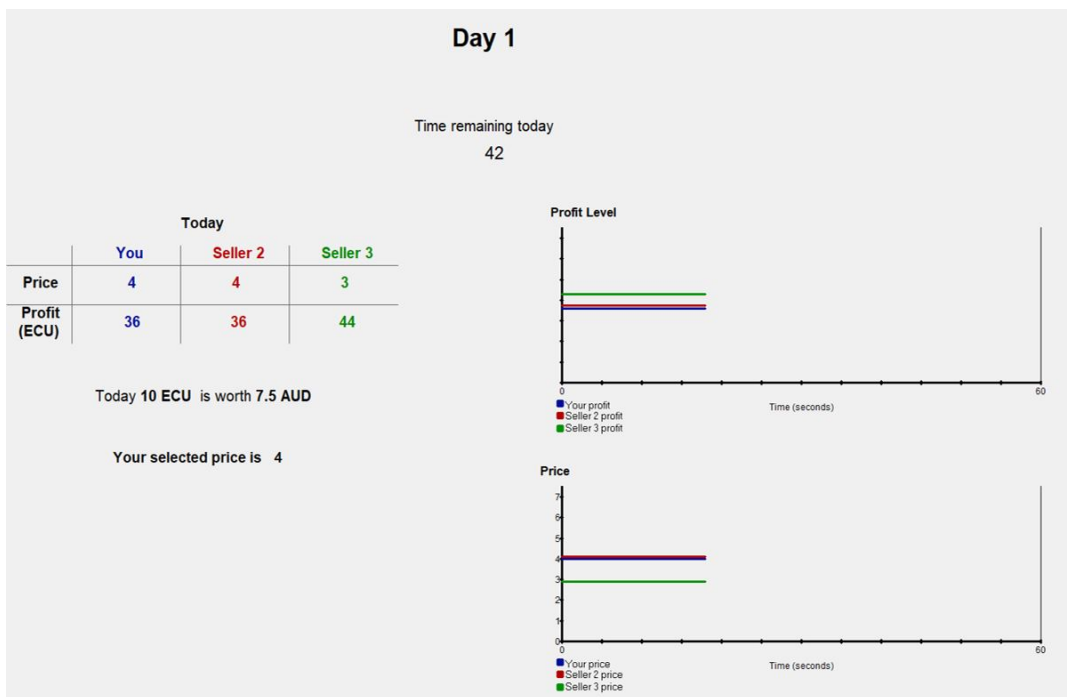


Figure 1: Screenshot of day screen

At the end of the day you will see an overnight screen, as shown in Figure 2. The screen summarises the price and profit for you and the other sellers you are competing with. You also have the option to change your starting price for the following day, which will take effect once the new day begins. If you don't change your price, the last price selected will be your starting price tomorrow.

Number of days in each game

The number of days in each game is randomly determined and unknown until the game has ended. As seen in Figure 2, the spinning of a computerised wheel determines if the game continues. After each day there is a 20% probability the game will end. This means there is an 80% probability that the game will continue to the next day.

If the game continues, you will compete with the same participants.

If the game ends, and there is enough time for another game, you will be rematched with two other participants and a new game begins. In case that there is not enough time in the session for another game, the experiment ends.

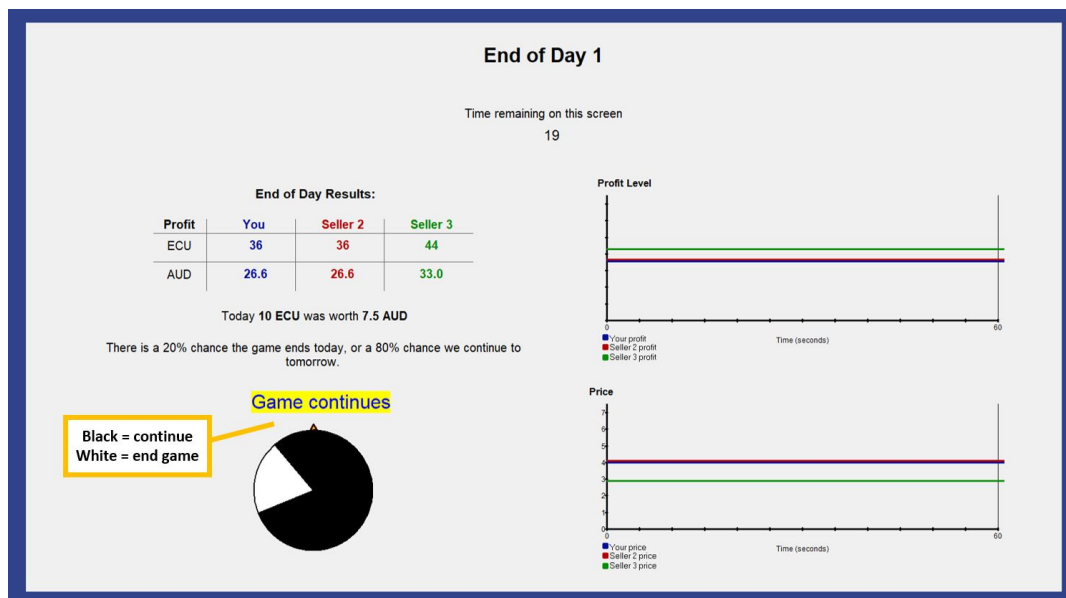


Figure 2: Screenshot of overnight screen

Profit

Below we explain how profit is determined.

Day profit - the profit table

How much you are able to sell depends on your price, and the average price of your competitors (Seller 2 and Seller 3). Profit, for each combination of prices, is shown in the attached Profit Table.

Example of how to look up profit

Let's look at an example of how to use the Profit Table. The table is replicated below in Table 1. Your price is 3. Seller 2's price is 4. Seller 3's price is 6:

You -

- the average of your competitors prices is $(4+6)/2 = 5$. If you set your price to 3, and the average of the other firms' price is 5, your profit for the period is 52 ECU. This is coloured blue in Table 1.

Seller 2 -

- the average price of you and Seller 3 is $(3+6)/2=4.5$. If Seller 2's price is 4 and the average price is 4.5, Seller 2's profit is 45 ECU. This is coloured red in Table 1.

Seller 3 -

- the average of your price and Seller 2's price is $(3+4)/2=3.5$. If Seller 3's price is 6 and the average price of others is 3.5, Seller 3's profit is 13 ECU. This is coloured green in Table 1.

Now complete [Question 2](#) on your screen.

Table 1: Profit table example (ECU)

		Average of other sellers' prices												
		1	1.5	2	2.5	3	3.5	4	4.5	5	5.5	6	6.5	7
Price	1	25	29	32	35	37	40	42	45	47	47	47	47	47
	2	26	30	33	36	39	42	45	48	51	54	56	56	56
	3	22	26	29	33	37	41	44	48	52	56	60	62	63
	4	13	18	22	27	31	36	40	45	49	54	58	63	67
	5	0	6	11	16	21	27	32	37	42	48	53	58	63
	6	0	0	0	4	7	13	19	25	31	37	43	49	55
	7	0	0	0	0	0	2	3	10	16	23	29	36	42

How does changing your price impact your profit?

Increasing your price impacts your profit in two ways. A higher price means you receive more money per unit sold, positively impacting your profit. However, there is also a negative effect, as a higher price means you will sell a lower quantity. Therefore, raising your price may increase or decrease your profit:

- Increasing a low price will increase your profit - For example, if your price is 1 and the average price of your competitors is 5, your profit is 47 ECU. If you were to raise your price to 3, your profit increases to 52 ECU.
- Increasing a high price will decrease your profit - For example, if your price is 6 and the average price of your competitors is 5, your profit is 31 ECU. If you were to raise your price to 7, your profit decreases to 16 ECU.

This is true for any given average price set by your competitors. The same logic can be applied to the impact of decreasing your price. Reducing a low price will decrease your profit, but reducing a high price will increase your profit.

You can see how changing your price (given the average price of your competitors) impacts your profit by looking down each column in the Profit Table.

How do your competitor's prices impact your profit?

If your competitors increase their prices (and yours is held constant), then you sell more and your profit increases. If your competitors decrease their prices, you sell less and your profit decreases.

For example, if you set your price to 5 and your competitors average price is also 5, your profit is 42 ECU. If the average of your competitors price decreases to 4.5, your profit is 37 ECU. If their average price increases to 5.5, your profit is 48.

You can see how changes in your competitors average price impacts your profit (for your given price) by looking across each row in the Profit Table.

Exchange rate & final payment

The value of profit will decrease by approximately 9% with every new day. The starting exchange rate is 1 ECU = 0.75 AUD. The exchange rate can be found in Table 2 for the first 10 days (also displayed on the day screen).

For example, on the first day 10 ECU is worth 7.5 AUD. On the second day, 10 ECU is worth 6.8 AUD, and so on. Let's say your day profit is 40 ECU. If the game ends on day 1, your payment for this game would be 30 AUD. However, if the game ends on day 6, your payment would be 18.6 AUD. Each new game will start again from day 1.

Please complete Question 3 on your screen.

Table 2: Exchange rate - The amount of AUD received for 10 ECU

Day	1	2	3	4	5	6	7	8	9	10	...
AUD per 10 ECU	7.50	6.82	6.20	5.63	5.12	4.66	4.23	3.85	3.50	3.18	...

At the end of the experiment a random game will be selected. You will be paid only for profit (in AUD) you earn on the last day before this game ended.

Please complete Question 4 on your screen.

Closing

After the last game of the experiment a questionnaire will appear on your screen for you to fill out before you can receive your payment. Please remain seated and silent.

Please fill out and sign your receipt form. The experimenter will then provide instructions for you to receive your payment.

Chapter 2 instructions for treatment n1 δ_i :

General information

Welcome to the experiment. Please read these instructions carefully. If you have a question please raise your hand. Please switch off your mobile phone and do not communicate with other participants.

During this experiment you can earn Experimental Currency Units (ECU) which will be exchanged into Australian dollars. You will receive payment at the end of the experiment, privately, in cash. The exchange rate will be given later in the instructions.

As you read through these instructions you will be prompted to answer questions displayed on your screen. These questions are designed to help you understand the game. If your answer is incorrect you may try again. If you are unsure at any time, please raise your hand. By answering these questions you will earn 10 AUD in addition to the profit you earn in the game.

There will be a practice game (competing with a computer) before the experiment begins.

Your task

In today's session all participants are sellers of a fictitious product. Your objective is to set and adjust your price such that you make as much profit as possible. In each game you will compete with two other sellers, who also set prices. A game consists of multiple days (which last for 60 seconds each).

Before the first day begins you will be asked to set your price. Your profit for the day depends on the price charged by you and your competitors during the day.

At the end of each day you will see your final profit. It is randomly determined if a new day starts, or if the game ends.

How to play

When a new game starts you will see a screen asking you to select your starting price. Possible prices are whole numbers between 1 and 7.

When the day begins you will observe your price and profit, as well as that of the other two sellers you are competing with.

Please complete Question 1 on your screen.

A screenshot of the day screen is shown in Figure 1. At the top of the screen there is information about the time remaining in the current day (countdown from 60 seconds to 0 seconds).

The top left-hand side of the screen shows a table with the current period's prices and profits for you and your competitors. On the right-hand side of the screen are two graphs showing the day's profit (top) and price (bottom) levels.

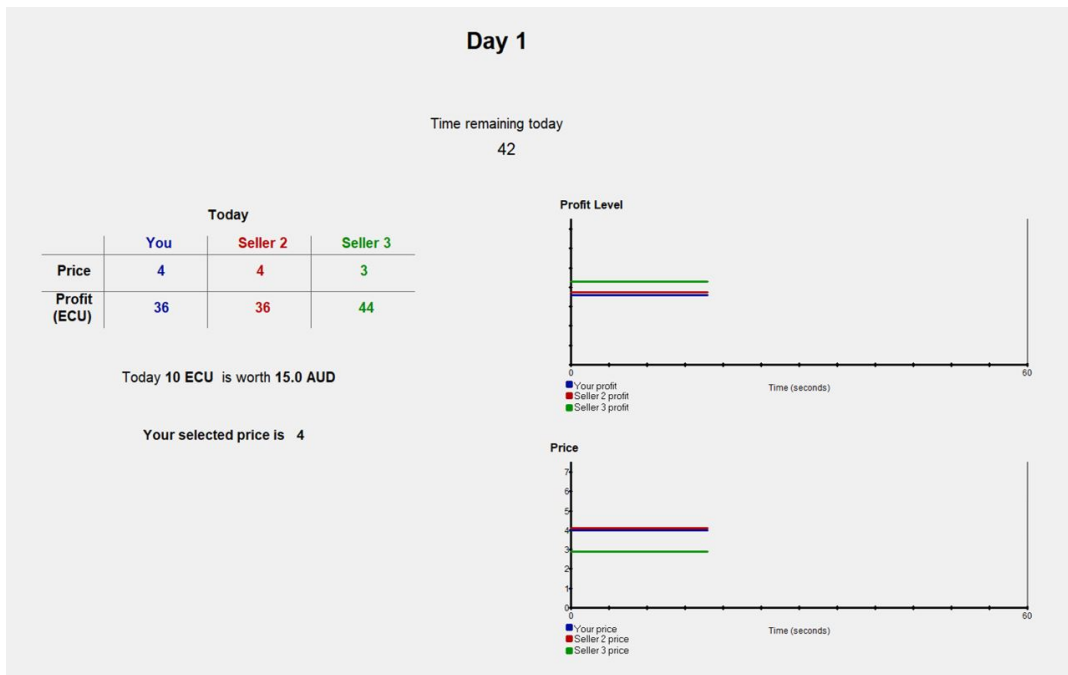


Figure 1: Screenshot of day screen

At the end of the day you will see an overnight screen, as shown in Figure 2. The screen summarises the price and profit for you and the other sellers you are competing with. You also have the option to change your starting price for the following day, which will take effect once the new day begins. If you don't change your price, the last price selected will be your starting price tomorrow.

Number of days in each game

The number of days in each game is randomly determined and unknown until the game has ended. As seen in Figure 2, the spinning of a computerised wheel determines if the game continues. After each day there is a 20% probability the game will end. This means there is an 80% probability that the game will continue to the next day.

If the game continues, you will compete with the same participants.

If the game ends, and there is enough time for another game, you will be rematched with two other participants and a new game begins. In case that there is not enough time in the session for another game, the experiment ends.

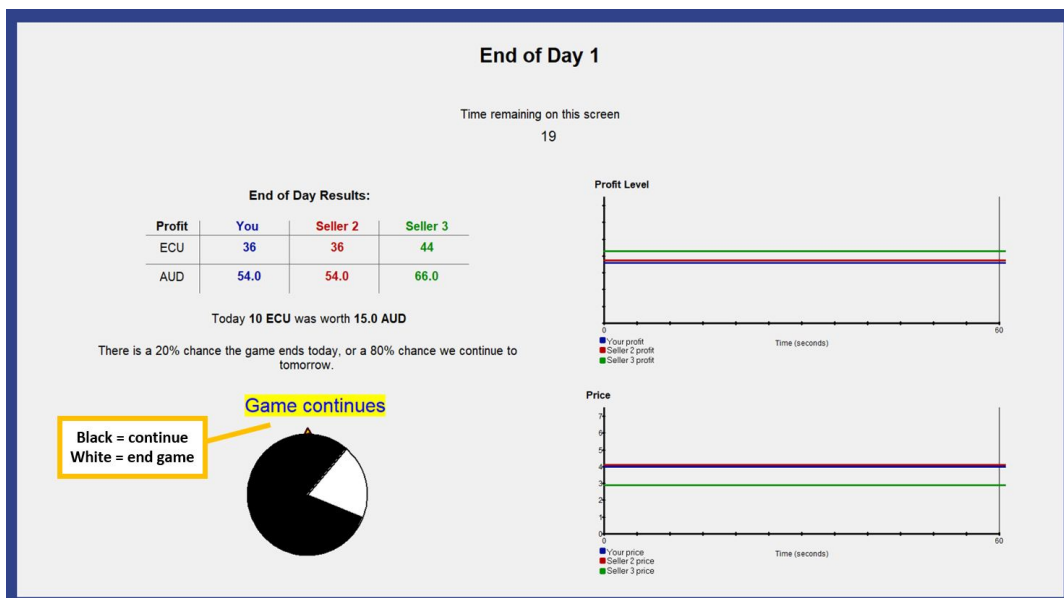


Figure 2: Screenshot of overnight screen

Profit

Below we explain how profit is determined.

Day profit - the profit table

How much you are able to sell depends on your price, and the average price of your competitors (Seller 2 and Seller 3). Profit, for each combination of prices, is shown in the attached Profit Table.

Example of how to look up profit

Let's look at an example of how to use the Profit Table. The table is replicated below in Table 1. Your price is 3. Seller 2's price is 4. Seller 3's price is 6:

You -

- the average of your competitors prices is $(4+6)/2 = 5$. If you set your price to 3, and the average of the other firms' price is 5, your profit for the period is 52 ECU. This is coloured blue in Table 1.

Seller 2 -

- the average price of you and Seller 3 is $(3+6)/2=4.5$. If Seller 2's price is 4 and the average price is 4.5, Seller 2's profit is 45 ECU. This is coloured red in Table 1.

Seller 3 -

- the average of your price and Seller 2's price is $(3+4)/2=3.5$. If Seller 3's price is 6 and the average price of others is 3.5, Seller 3's profit is 13 ECU. This is coloured green in Table 1.

Now complete [Question 2](#) on your screen.

Table 1: Profit table example (ECU)

		Average of other sellers' prices												
		1	1.5	2	2.5	3	3.5	4	4.5	5	5.5	6	6.5	7
Price	1	25	29	32	35	37	40	42	45	47	47	47	47	47
	2	26	30	33	36	39	42	45	48	51	54	56	56	56
	3	22	26	29	33	37	41	44	48	52	56	60	62	63
	4	13	18	22	27	31	36	40	45	49	54	58	63	67
	5	0	6	11	16	21	27	32	37	42	48	53	58	63
	6	0	0	0	4	7	13	19	25	31	37	43	49	55
	7	0	0	0	0	0	2	3	10	16	23	29	36	42

How does changing your price impact your profit?

Increasing your price impacts your profit in two ways. A higher price means you receive more money per unit sold, positively impacting your profit. However, there is also a negative effect, as a higher price means you will sell a lower quantity. Therefore, raising your price may increase or decrease your profit:

- Increasing a low price will increase your profit - For example, if your price is 1 and the average price of your competitors is 5, your profit is 47 ECU. If you were to raise your price to 3, your profit increases to 52 ECU.
- Increasing a high price will decrease your profit - For example, if your price is 6 and the average price of your competitors is 5, your profit is 31 ECU. If you were to raise your price to 7, your profit decreases to 16 ECU.

This is true for any given average price set by your competitors. The same logic can be applied to the impact of decreasing your price. Reducing a low price will decrease your profit, but reducing a high price will increase your profit.

You can see how changing your price (given the average price of your competitors) impacts your profit by looking down each column in the Profit Table.

How do your competitor's prices impact your profit?

If your competitors increase their prices (and yours is held constant), then you sell more and your profit increases. If your competitors decrease their prices, you sell less and your profit decreases.

For example, if you set your price to 5 and your competitors average price is also 5, your profit is 42 ECU. If the average of your competitors price decreases to 4.5, your profit is 37 ECU. If their average price increases to 5.5, your profit is 48.

You can see how changes in your competitors average price impacts your profit (for your given price) by looking across each row in the Profit Table.

Exchange rate & final payment

The value of profit will decrease by approximately 22% with every new day. The starting exchange rate is 1 ECU = 1.5 AUD. The exchange rate can be found in Table 2 for the first 10 days (also displayed on the day screen).

For example, on the first day 10 ECU is worth 15 AUD. On the second day, 10 ECU is worth 11.6 AUD, and so on. Let's say your day profit is 40 ECU. If the game ends on day 1, your payment for this game would be 60 AUD. However, if the game ends on day 6, your payment would be 16.8 AUD. Each new game will start again from day 1.

Please complete Question 3 on your screen.

Table 2: Exchange rate - The amount of AUD received for 10 ECU

Day	1	2	3	4	5	6	7	8	9	10
AUD per 10 ECU	15.00	11.63	9.01	6.99	5.42	4.20	3.26	2.52	1.96	1.52

At the end of the experiment a random game will be selected. You will be paid only for profit (in AUD) you earn on the last day before this game ended.

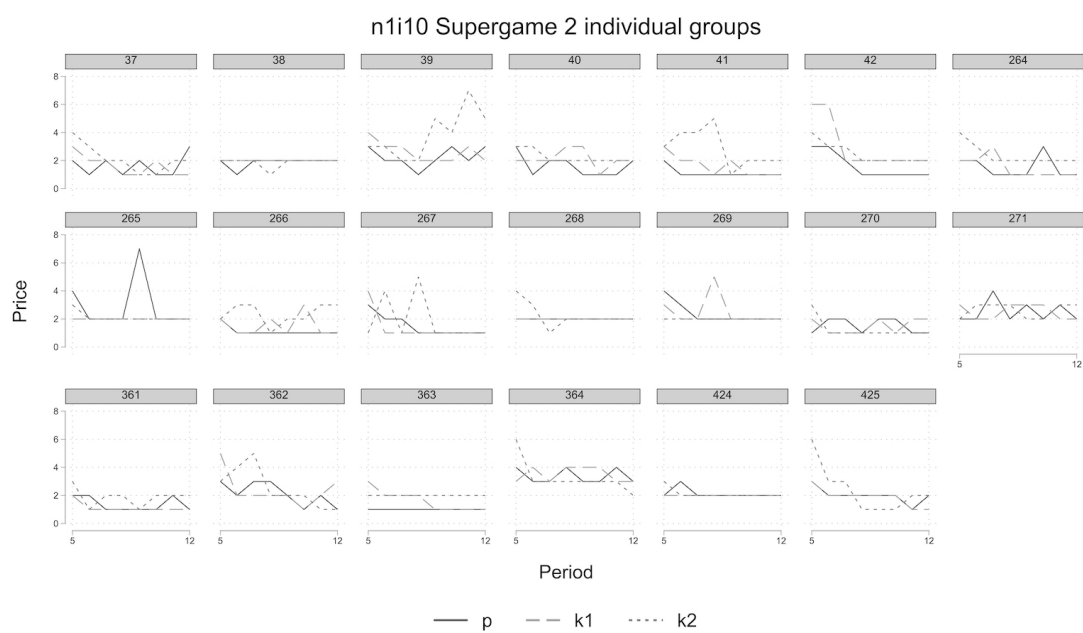
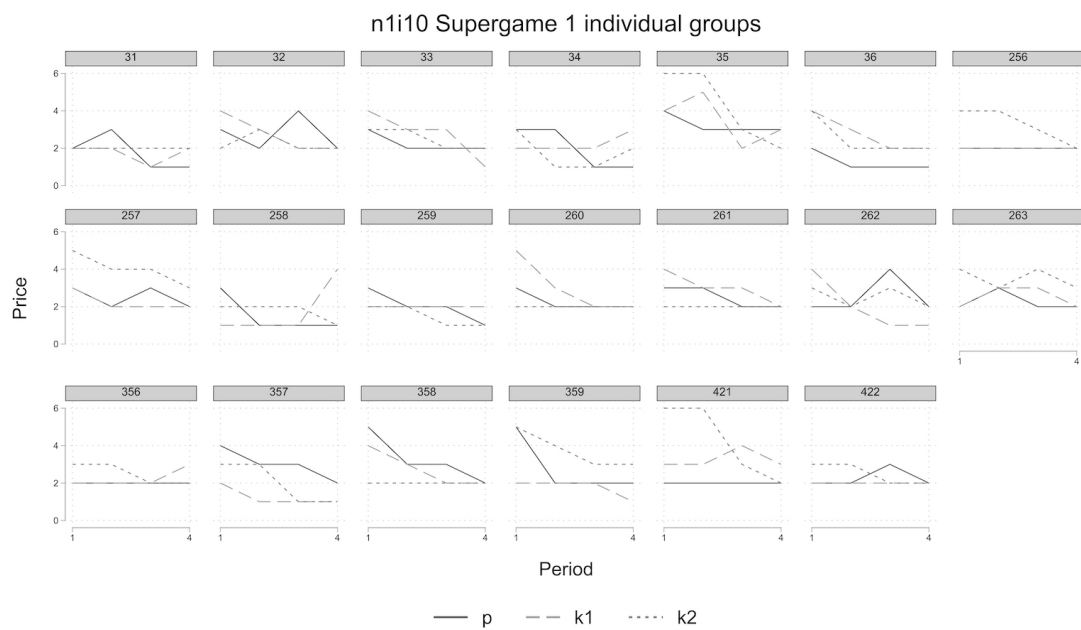
Please complete Question 4 on your screen.

Closing

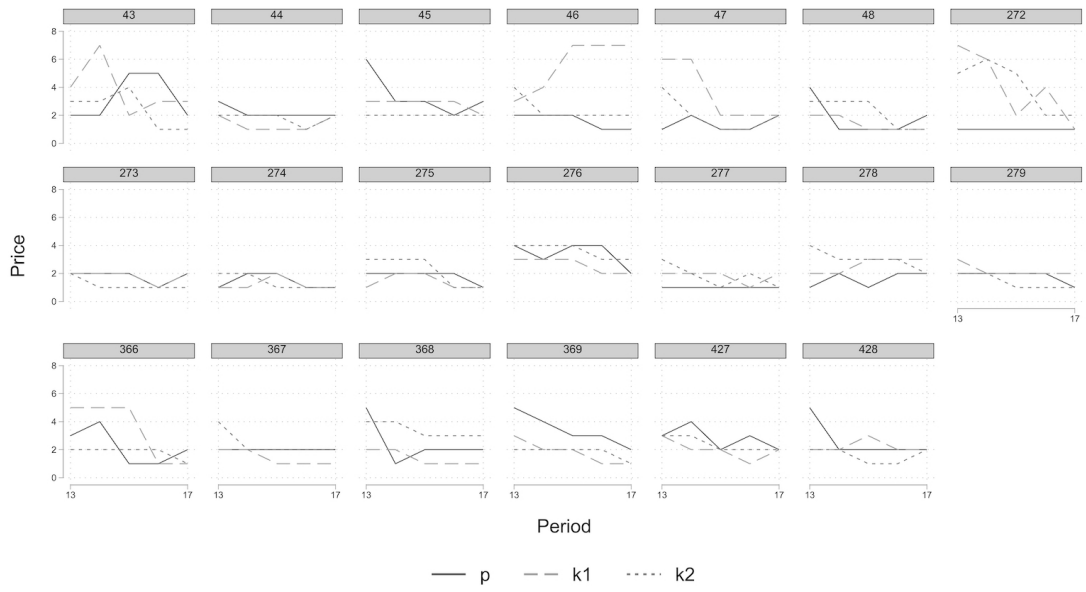
After the last game of the experiment a questionnaire will appear on your screen for you to fill out before you can receive your payment. Please remain seated and silent.

Please fill out and sign your receipt form. The experimenter will then provide instructions for you to receive your payment.

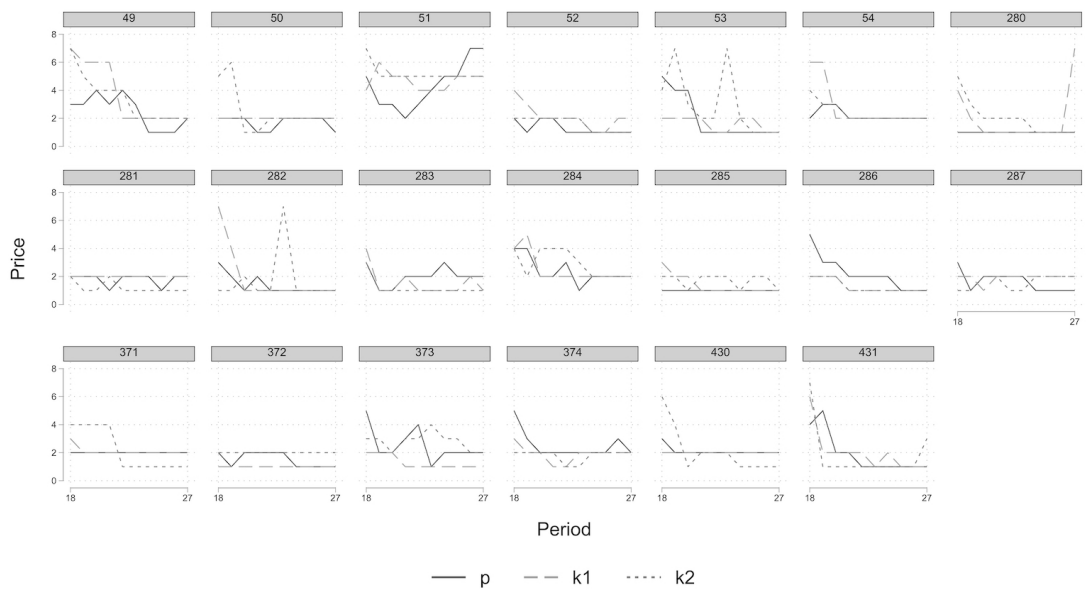
A.2 Chapter 2 - Individual markets



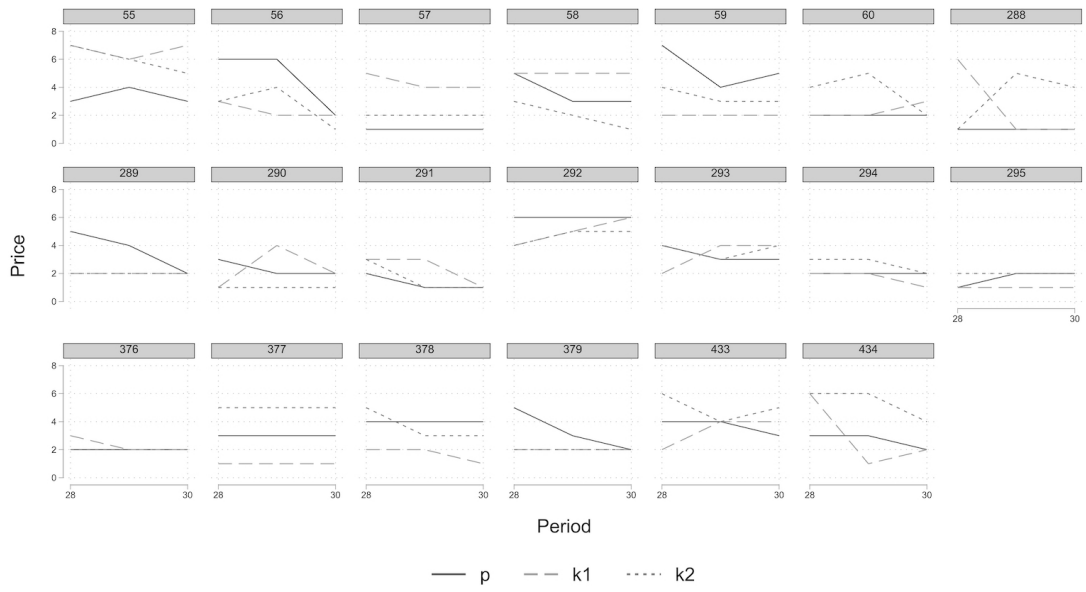
n1i10 Supergame 3 individual groups



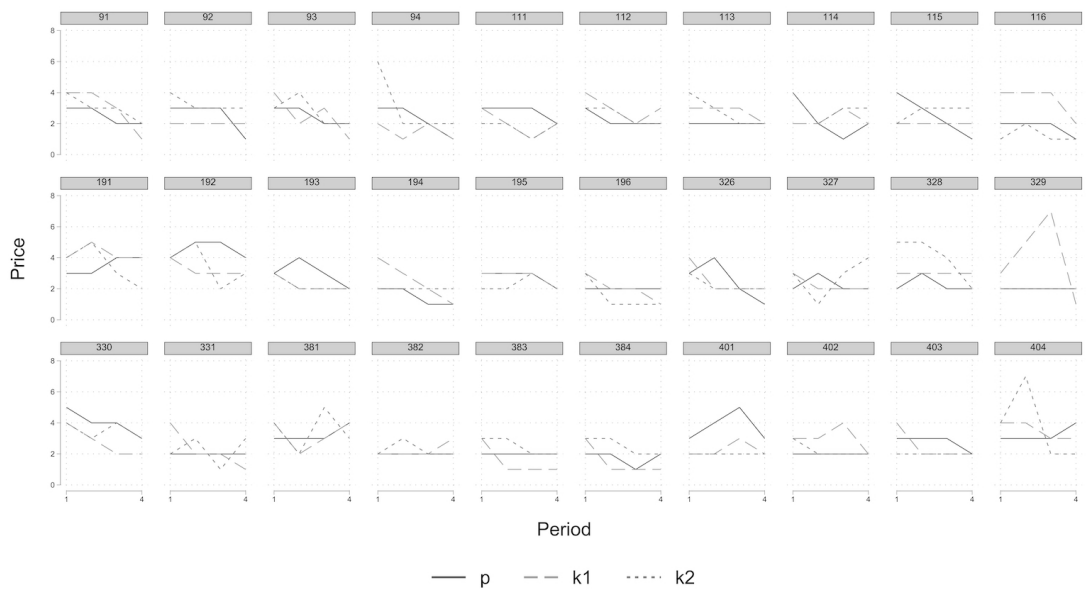
n1i10 Supergame 4 individual groups



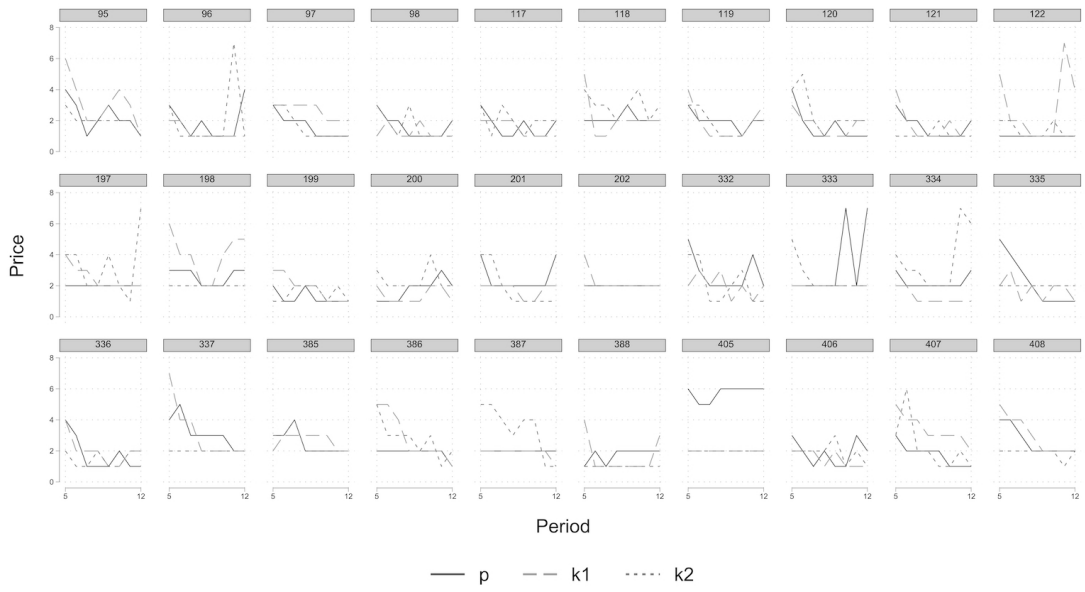
n1i10 Supergame 5 individual groups



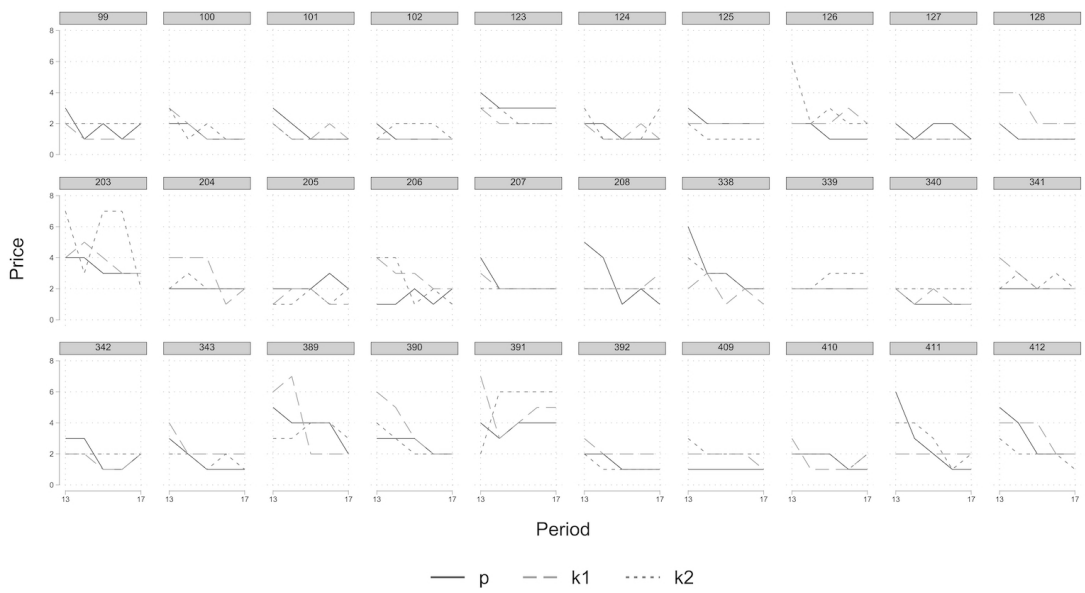
n1i29 Supergame 1 individual groups




n1i29 Supergame 2 individual groups



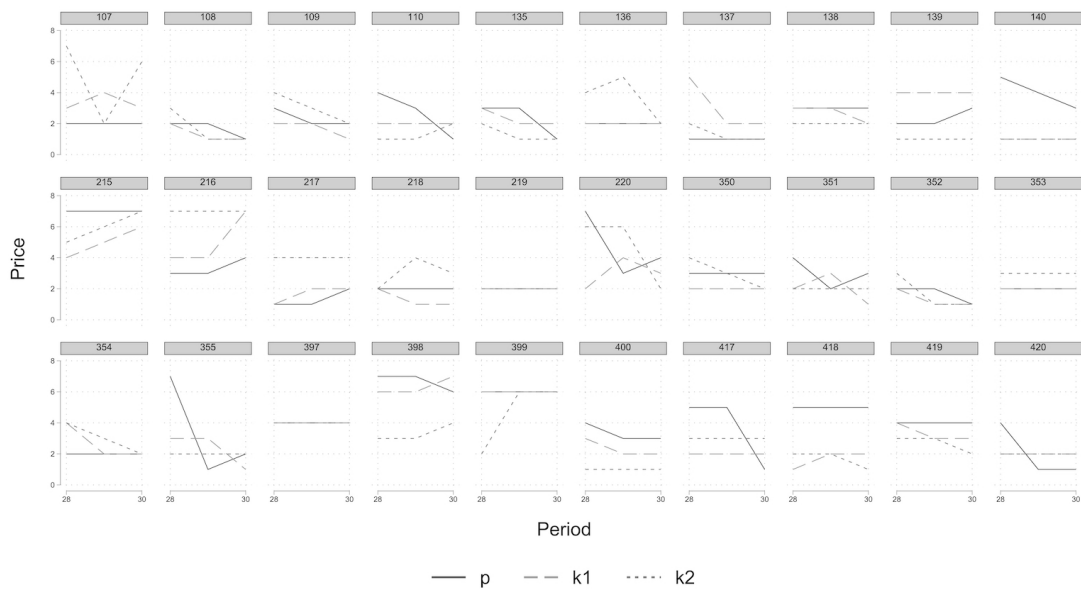
n1i29 Supergame 3 individual groups




n1i29 Supergame 4 individual groups




n1i29 Supergame 5 individual groups




n4i10 Supergame 1 individual groups



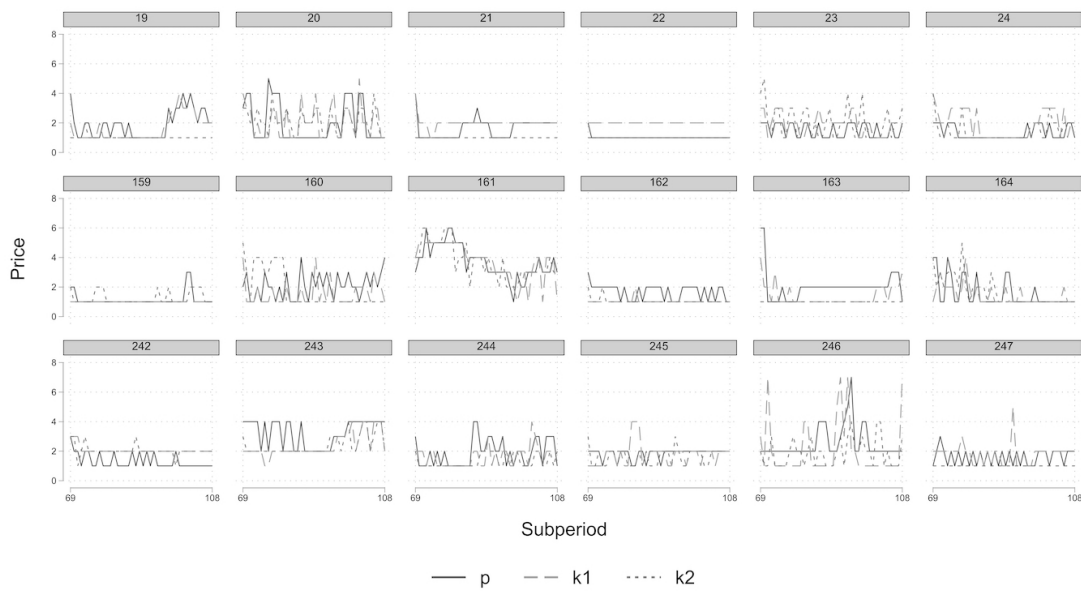
n4i10 Supergame 2 individual groups



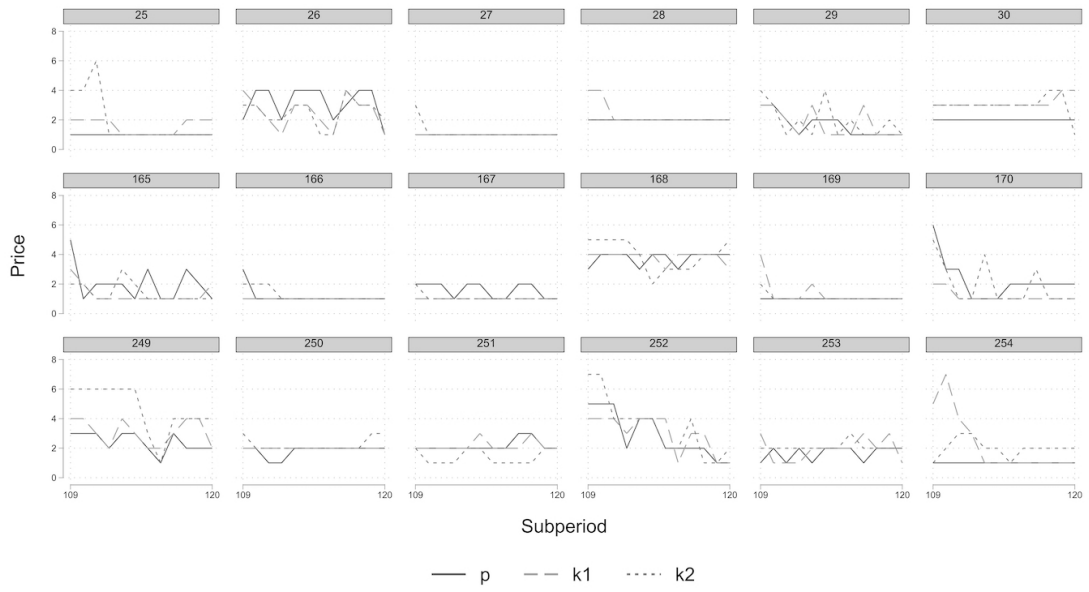
n4i10 Supergame 3 individual groups




n4i10 Supergame 4 individual groups



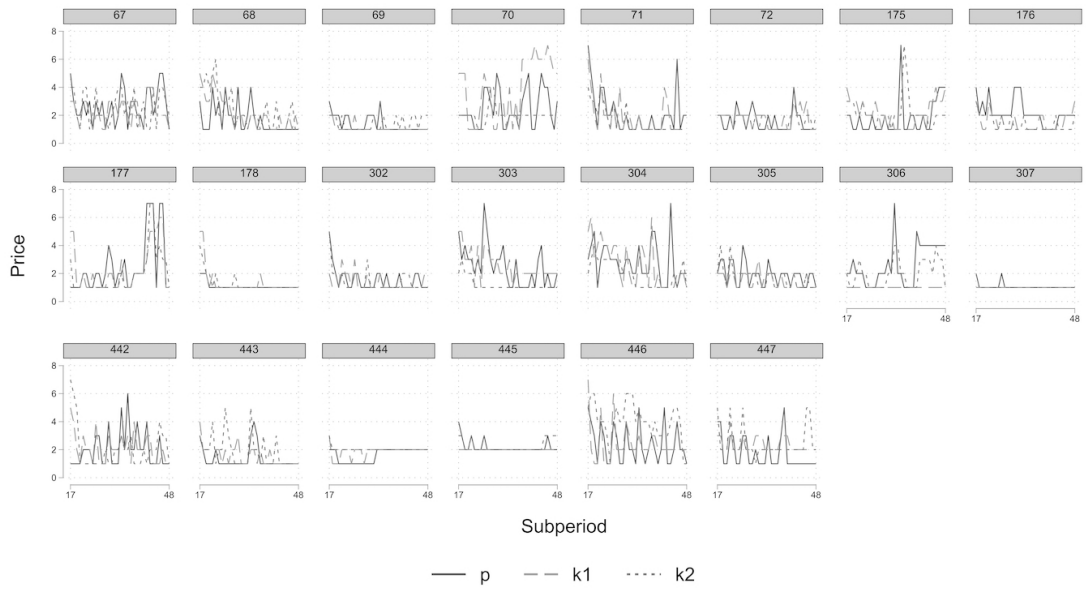
n4i10 Supergame 5 individual groups



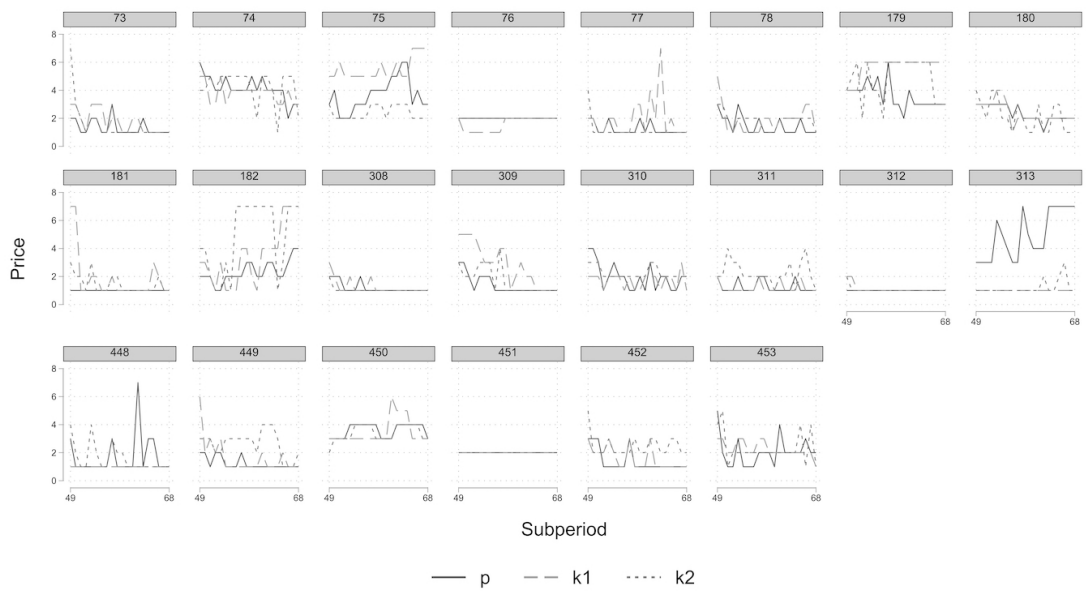
n4i29 Supergame 1 individual groups



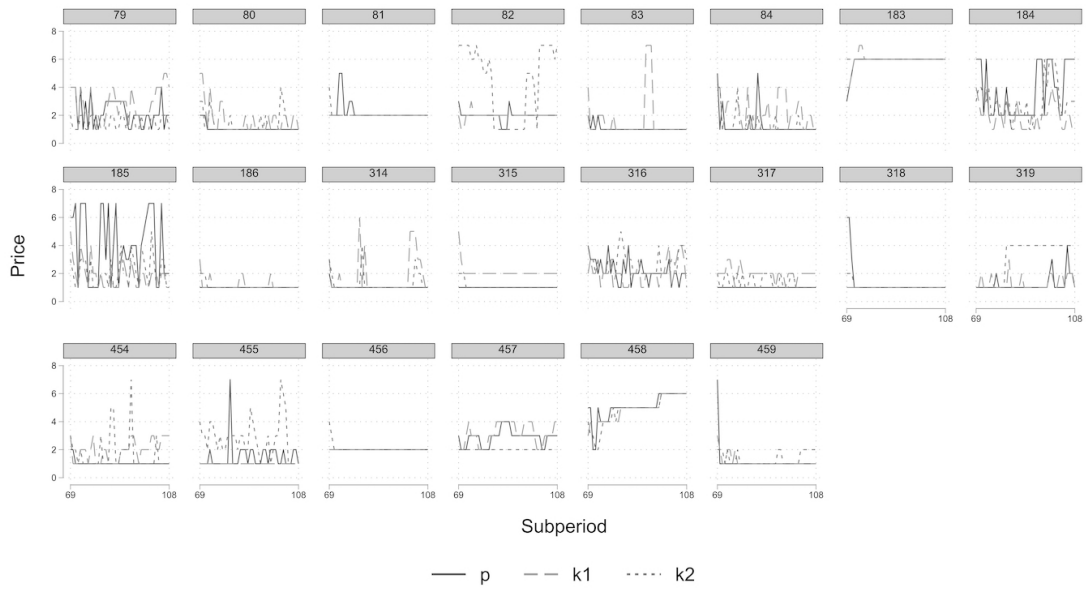
n4i29 Supergame 2 individual groups



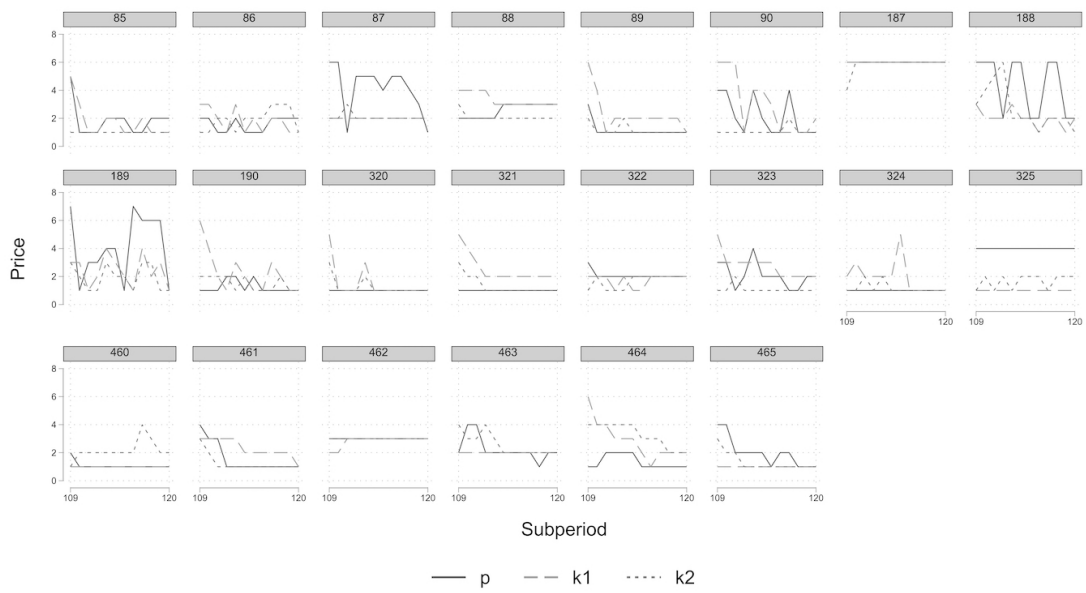
n4i29 Supergame 3 individual groups



n4i29 Supergame 4 individual groups



n4i29 Supergame 5 individual groups



A.3 Chapter 3 - Instructions

Chapter 3 Instructions for Treatment *continuous*:

General information

Welcome to the experiment. Please switch off your mobile phone and do not communicate with other participants. Raise your hand if you have a question, an experimenter will assist you.

During this experiment you can earn Experimental Currency Units (ECU) which will be exchanged into Australian dollars. You will receive payment at the end of the experiment, privately, in cash.

The exchange rate is $130 \text{ ECU} = 1 \text{ AUD}$.

There will be review questions and a practice round (playing with a computer) to help you understand the game before the experiment begins.

Your task

In today's session all participants are sellers of a fictitious product. Your objective is to set and adjust your price to make as much profit as possible. You will be paired with another participant (Seller 2), who also sets prices. The two of you are the only sellers in the market. The product you sell is similar to the one sold by Seller 2, and so the amount you are able to sell depends on your price, as well as Seller 2's. This will be explained in more detail later in the instructions.

You will be partnered with the same Seller for the entire experiment, which will run for 60 minutes.

Throughout the experiment there are two different types of intervals: Day and Night. This will impact the profit you can earn at a given price - you can think of this as there being more buyers in the market during the Day, but fewer around at Night.

The experiment will start during the Day, and after 20 seconds will switch to Night. Then, 20 seconds later, it will switch back to Day time. This cycle of Day and Night will continue for the entire 60 minutes of the experiment. Each cycle (1 Day + 1 Night) your total profit is calculated as the average profit over that 40 second interval.

You are able to change your price at any time. Your price change will be implemented immediately.

The next section will explain how to play and the screen you will see. We will then explain how profit is calculated. Finally, we explain how your final payment for today's session will be determined.

How to play

At the start of the experiment you will be asked to choose your starting price. Figure 1 shows the screen you will see once the experiment begins.

At the top right of the screen there is a table that shows current prices and profits.

In the middle there is a graph that plots the prices of you and Seller 2. Your price line is thicker, while Seller 2's is represented using a thin blue line.

To help you keep track of your current profit, the colour of your entire price line will change between black to light green. The darker the line, the lower your profit. The brighter green the line is, the higher your profit. Seller 2's line does not change colour on your graph. You can see their current profit in the table above the price graph.

The shaded blocks on the graph represent Night time, while the un-shaded areas represent Day time. You can use this to keep track of when Day/Night is approaching.

On the right of the screen there are buttons to change your price. Possible prices are between 1 and 20. You can change your price at any time, which will take effect immediately. Seller 2 can also change their price at any time. When this happens, the prices and profits in the table and graph will update immediately.

At the top left of the screen is a timer counting down from 60 minutes. Below that you can see your total profit for the previous cycle, the average of your profit for that 40 second interval.

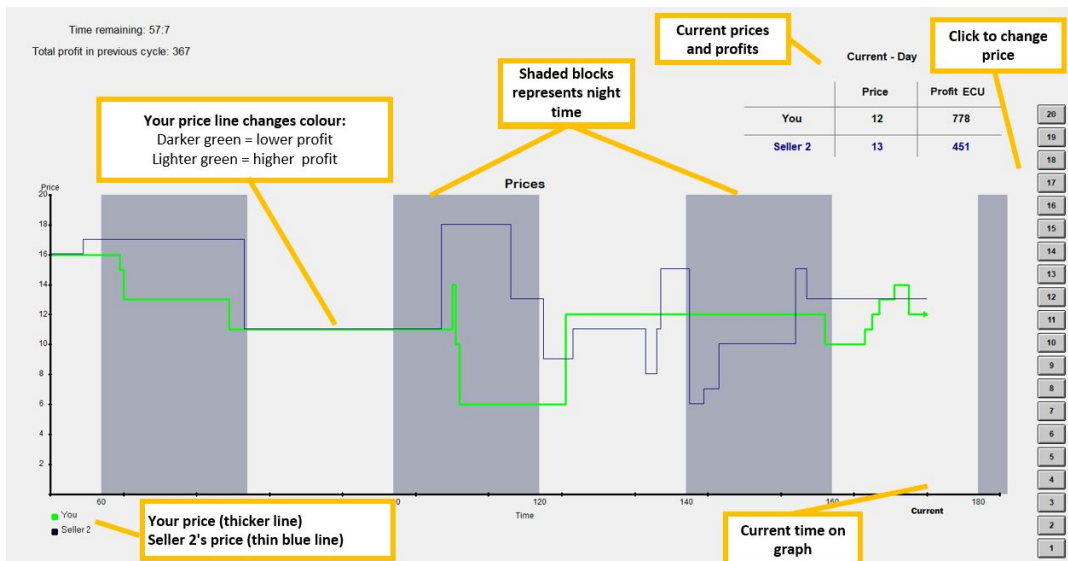


Figure 1: Screenshot of the user interface

Profit

In this section we explain how to calculate profit, and how it is impacted by Day/Night and relative prices. Please note the two profit tables attached.

During the experiment, your profit will depend on your price, and the price set by Seller 2. Your total profit each cycle is the average profit during each 40 second interval. Profit Table 1 shows profits during the Day, while Profit Table 2 shows profits during the Night. You can think of this as there being fewer buyers around at Night, while there are more buyers during the Day.

How to use the profit table

You can use each profit table to determine your profit, depending on your price, and the price set by Seller 2. The column on the left is your price. As you look across each row, you can see how your profit changes with the price selected by Seller 2.

For example, if your price is 13 and Seller 2's price is 16 your profit is 1,114 ECU during the Day, and 332 ECU at Night.

You can also use the same process to figure out what Seller 2's profit would be, as you both are using the same table. Instead, now use the 'Your price' column to represent Seller 2.

If Seller 2's price is 16, and your price is 13, then Seller 2's profit is 210 ECU during the Day, and 17 at Night.

If your prices are the same, your profits are also the same.

For example, if you and Seller 2 both set your prices to 17, during the Day each of you has a profit of 789 ECU, while at Night your profits would both be 8 ECU.

Understanding prices and profit

You and Seller 2 make similar products. This means that buyers are more likely to purchase from the seller with the lower price. Your profit depends on the amount you can sell (your price, *relative* to Seller 2), and how much you earn from selling each unit (your price):

- If you set a low price, you will sell more units, but your profit won't be high as you receive a low price per unit sold.
- As you start increasing your price, you may sell fewer units, but overall you earn more profit due to the higher price per unit.
- However, after a certain point the higher price is outweighed by a decrease in units sold - so your profit will decrease.
- To see this, look down the profit tables to see how your profit changes as you increase price from 1 to 20, for any given price set by Seller 2.

Although more buyers will purchase from the seller with the lower price, if you both set your prices higher you may both earn higher profit levels:

- For example, if you both set a lower price (i.e. 3), you both make 149 ECU. If your both increase your price to 10, you each have a profit of 499 ECU during the Day, and 472 ECU at Night.

- You can see how Seller 2's price impacts your profit by looking across the profit tables from left to right, for any price you choose. Your profit increases if Seller 2's price increases. This works the same for Seller 2 - if your price increases, Seller 2's profit increases.

Therefore, when choosing your price to earn as much profit as possible, consider your price relative to that of Seller 2, and the profit that each of you earn at those prices.

Payment

After the final day, a computer will randomly select 10 intervals of 40 seconds from the experiment. Each interval is equally likely to be chosen. You will only be paid for the average profit you made during these 10 intervals. The selected intervals and profits in ECU will be shown on your screen once the experiment ends.

Chapter 3 Instructions for Treatment *discrete*:

General information

Welcome to the experiment. Please switch off your mobile phone and do not communicate with other participants. Raise your hand if you have a question, an experimenter will assist you.

During this experiment you can earn Experimental Currency Units (ECU) which will be exchanged into Australian dollars. You will receive payment at the end of the experiment, privately, in cash.

The exchange rate is $130 \text{ ECU} = 1 \text{ AUD}$.

There will be review questions and a practice round (playing with a computer) to help you understand the game before the experiment begins.

Your task

In today's session all participants are sellers of a fictitious product. Your objective is to set and adjust your price to make as much profit as possible. You will be paired with another participant (Seller 2), who also sets prices. The two of you are the only sellers in the market. The product you sell is similar to the one sold by Seller 2, and so the amount you are able to sell depends on your price, as well as Seller 2's. This will be explained in more detail later in the instructions.

You will be partnered with the same Seller for the entire experiment, which will run for 60 minutes.

Throughout the experiment there are two different types of intervals: Day and Night. This will impact the profit you can earn at a given price - you can think of this as there being more buyers in the market during the Day, but fewer around at Night.

The experiment will start during the Day, and after 20 seconds will switch to Night. Then, 20 seconds later, it will switch back to Day time. This cycle of Day and Night will continue for the entire 60 minutes of the experiment. Each cycle (1 Day + 1 Night) your total profit is calculated as the average profit over that 40 second interval.

You are able to change your price at any time - but your price will only take effect at the start of the next Day or Night.

The next section will explain how to play and the screen you will see. We will then explain how profit is calculated. Finally, we explain how your final payment for today's session will be determined.

How to play

At the start of the experiment you will be asked to choose your starting price. Figure 1 shows the screen you will see once the experiment begins.

At the top right of the screen there is a table that shows current prices and profits.

In the middle there is a graph that plots the prices of you and Seller 2. Your price line is thicker, while Seller 2's is represented using a thin blue line.

To help you keep track of your current profit, the colour of your entire price line will change between black to light green. The darker the line, the lower your profit. The brighter green the line is, the higher your profit. Seller 2's line does not change colour on your graph. You can see their current profit in the table above the price graph.

The shaded blocks on the graph represent Night time, while the un-shaded areas represent Day time. You can use this to keep track of when Day/Night is approaching.

On the right of the screen there are buttons to change your price. Possible prices are between 1 and 20. You can change your price at any time, however it won't take effect until the next Day or Night. For example, if you change your price during the Day, this new price will only take effect once Night starts. You can see your selected price below the table on the right. This is also the same for Seller 2. When there is a switch from Day to Night (or Night to Day), the prices and profits in the table and graph will update automatically.

At the top left of the screen is a timer counting down from 60 minutes. Below that you can see your total profit for the previous cycle, the average of your profit for that 40 second interval.

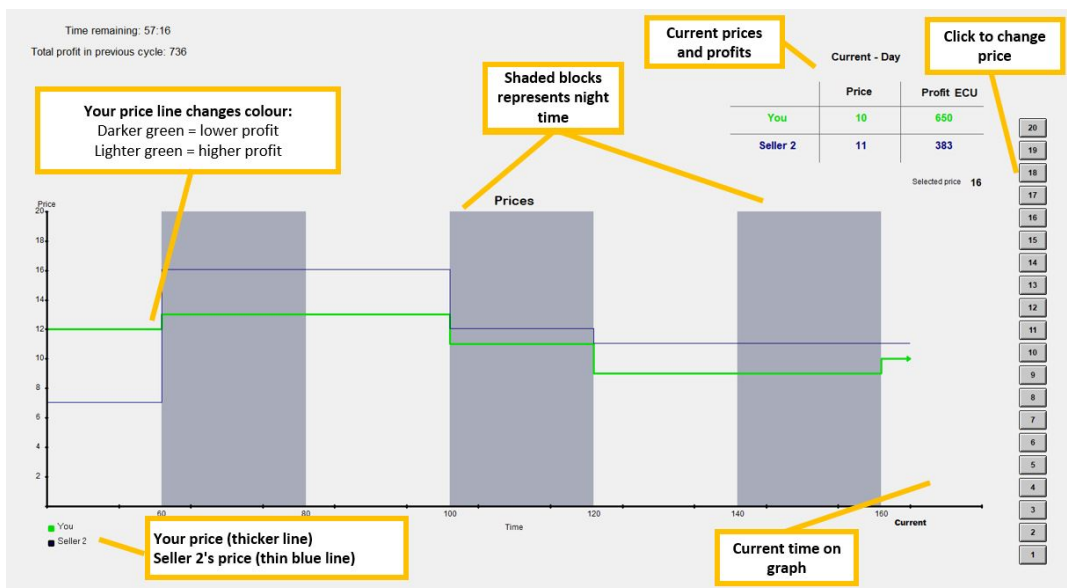


Figure 1: Screenshot of the user interface

Profit

In this section we explain how to calculate profit, and how it is impacted by Day/Night and relative prices. Please note the two profit tables attached.

During the experiment, your profit will depend on your price, and the price set by Seller 2. Your total profit each cycle is the average profit during each 40 second interval. Profit Table 1 shows profits during the Day, while Profit Table 2 shows profits during the Night. You can think of this as there being fewer buyers around at Night, while there are more buyers during the Day.

How to use the profit table

You can use each profit table to determine your profit, depending on your price, and the price set by Seller 2. The column on the left is your price. As you look across each row, you can see how your profit changes with the price selected by Seller 2.

For example, if your price is 13 and Seller 2's price is 16 your profit is 1,114 ECU during the Day, and 332 ECU at Night.

You can also use the same process to figure out what Seller 2's profit would be, as you both are using the same table. Instead, now use the 'Your price' column to represent Seller 2.

If Seller 2's price is 16, and your price is 13, then Seller 2's profit is 210 ECU during the Day, and 17 at Night.

If your prices are the same, your profits are also the same.

For example, if you and Seller 2 both set your prices to 17, during the Day each of you has a profit of 789 ECU, while at Night your profits would both be 8 ECU.

Understanding prices and profit

You and Seller 2 make similar products. This means that buyers are more likely to purchase from the seller with the lower price. Your profit depends on the amount you can sell (your price, *relative* to Seller 2), and how much you earn from selling each unit (your price):

- If you set a low price, you will sell more units, but your profit won't be high as you receive a low price per unit sold.
- As you start increasing your price, you may sell fewer units, but overall you earn more profit due to the higher price per unit.
- However, after a certain point the higher price is outweighed by a decrease in units sold - so your profit will decrease.
- To see this, look down the profit tables to see how your profit changes as you increase price from 1 to 20, for any given price set by Seller 2.

Although more buyers will purchase from the seller with the lower price, if you both set your prices higher you may both earn higher profit levels:

- For example, if you both set a lower price (i.e. 3), you both make 149 ECU. If your both increase your price to 10, you each have a profit of 499 ECU during the Day, and 472 ECU at Night.

- You can see how Seller 2's price impacts your profit by looking across the profit tables from left to right, for any price you choose. Your profit increases if Seller 2's price increases. This works the same for Seller 2 - if your price increases, Seller 2's profit increases.

Therefore, when choosing your price to earn as much profit as possible, consider your price relative to that of Seller 2, and the profit that each of you earn at those prices.

Payment

After the final day, a computer will randomly select 10 intervals of 40 seconds from the experiment. Each interval is equally likely to be chosen. You will only be paid for the average profit you made during these 10 intervals. The selected intervals and profits in ECU will be shown on your screen once the experiment ends.

Profit table provided to participants:

Profit Table 1 - Day (ECU)

		Seller 2's price																			
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	49	65	77	86	92	95	97	98	99	99	99	99	99	99	99	99	99	99	99	99	99
2	69	99	130	155	173	184	191	195	197	198	199	199	199	199	199	199	199	199	199	199	199
3	66	104	149	195	233	260	277	287	293	296	297	298	299	299	299	299	299	299	299	299	299
4	53	89	139	199	260	310	346	369	383	390	395	397	398	399	399	399	399	399	399	399	399
5	37	66	111	174	249	325	388	433	462	478	488	493	496	498	498	498	499	499	499	499	499
6	25	45	79	133	209	299	390	466	520	554	574	586	592	595	597	598	599	599	599	599	599
7	16	29	53	93	155	244	349	455	543	606	646	670	683	691	695	697	698	699	699	699	699
8	9	18	33	60	106	178	278	399	520	621	693	738	765	781	789	794	796	798	798	799	799
9	6	11	20	37	68	119	200	313	449	585	699	779	830	861	878	887	893	895	897	898	898
10	3	6	12	22	42	75	132	222	348	499	650	776	865	922	956	975	985	991	994	996	996
11	2	3	7	13	25	46	83	146	244	383	549	714	852	950	1,013	1,050	1,070	1,082	1,088	1,092	1,092
12	1	2	4	8	14	27	50	90	159	266	417	597	778	927	1,034	1,102	1,142	1,164	1,177	1,184	1,184
13	0	1	2	4	8	16	29	54	98	172	288	451	645	839	1,000	1,114	1,187	1,230	1,254	1,268	1,268
14	0	0	1	2	5	9	17	32	58	106	185	310	484	691	898	1,068	1,189	1,266	1,311	1,337	1,337
15	0	0	0	1	2	5	10	18	34	63	113	198	330	515	733	949	1,127	1,252	1,332	1,378	1,378
16	0	0	0	0	1	3	5	10	19	36	67	120	210	349	542	768	989	1,169	1,295	1,374	1,374
17	0	0	0	0	0	1	3	6	11	21	38	71	127	221	366	562	789	1,006	1,180	1,300	1,300
18	0	0	0	0	0	0	1	3	6	12	22	41	74	133	230	376	570	787	988	1,144	1,144
19	0	0	0	0	0	0	1	1	3	6	12	23	43	78	138	235	378	558	749	917	917
20	0	0	0	0	0	0	0	1	2	3	7	13	24	44	80	141	234	364	517	666	666

Your price

Profit Table 2 - Night (ECU)

		Seller 2's price																			
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	49	74	89	95	98	99	99	99	99	99	99	99	99	99	99	99	99	99	99	99	99
2	51	99	148	178	191	197	198	199	199	199	199	199	199	199	199	199	199	199	199	199	199
3	32	77	149	222	267	287	295	298	299	299	299	299	299	299	299	299	299	299	299	299	299
4	16	43	103	199	296	356	383	394	397	399	399	399	399	399	399	399	399	399	399	399	399
5	7	20	54	129	249	370	445	479	492	497	498	499	499	499	499	499	499	499	499	499	499
6	3	8	24	65	155	299	444	533	574	590	595	597	598	598	598	598	598	598	598	598	598
7	1	3	10	28	75	180	349	516	621	668	686	692	695	695	696	696	696	696	696	696	696
8	0	1	4	11	32	86	206	397	586	703	756	776	784	786	787	788	788	788	788	788	788
9	0	0	1	4	13	36	97	230	440	646	772	829	851	859	861	862	863	863	863	863	863
10	0	0	0	1	5	14	40	107	250	471	679	804	858	879	887	889	890	891	891	891	891
11	0	0	0	0	1	5	16	44	115	261	468	647	747	790	806	812	814	815	815	815	815
12	0	0	0	0	0	2	6	17	47	246	400	510	565	565	567	595	598	599	599	599	599
13	0	0	0	0	0	2	6	18	47	107	193	267	308	326	332	335	335	336	336	336	336
14	0	0	0	0	0	0	0	2	6	18	42	80	115	137	146	150	151	151	151	151	151
15	0	0	0	0	0	0	0	0	2	6	16	31	45	54	58	60	61	61	61	61	61
16	0	0	0	0	0	0	0	0	0	2	6	11	17	20	22	23	23	23	23	23	23
17	0	0	0	0	0	0	0	0	0	0	2	4	6	7	8	8	8	8	8	8	8
18	0	0	0	0	0	0	0	0	0	0	0	0	2	2	3	3	3	3	3	3	3
19	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
20	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Your price

A.4 Chapters 3 and 4 - Control questions

Review Question

	Price
You	14
Seller 2	15

This table identifies the prices set by you and Seller 2.

How much profit does each seller make (in ECU)?

What is your **Day** profit **Correct answer!**

What is Seller 2's **Day** profit? **Correct answer!**

What is your **Night** profit? **Correct answer!**

What is Seller 2's **Night** profit? **Correct answer!**

A.5 Chapter 4 - Instructions

Instructions for Treatment *short-discrete*:

General information

Welcome to the experiment. Please switch off your mobile phone and do not communicate with other participants. Raise your hand if you have a question, an experimenter will assist you.

During this experiment you can earn Experimental Currency Units (ECU) which will be exchanged into Australian dollars. You will receive payment at the end of the experiment, privately, in cash.

The exchange rate is $130 \text{ ECU} = 1 \text{ AUD}$.

There will be review questions and a practice round (playing with a computer) to help you understand the game before the experiment begins.

Your task

In today's session all participants are sellers of a fictitious product. Your objective is to set and adjust your price to make as much profit as possible. You will be paired with another participant (Seller 2), who also sets prices. The two of you are the only sellers in the market. The product you sell is similar to the one sold by Seller 2, and so the amount you are able to sell depends on your price, as well as Seller 2's. This will be explained in more detail later in the instructions.

You will be partnered with the same Seller for the entire experiment, which will run for 60 minutes.

Throughout the experiment there are two different types of intervals: Day and Night. This will impact the profit you can earn at a given price - you can think of this as there being more buyers in the market during the Day, but fewer around at Night.

The experiment will start during the Day, and after 20 seconds will switch to Night. Then, 20 seconds later, it will switch back to Day time. This cycle of Day and Night will continue for the entire 60 minutes of the experiment. Each cycle (1 Day + 1 Night) your total profit is calculated as the average profit over that 40 second interval.

You are able to change your price at any time - but prices only take effect every 2.5 seconds. This equates to 8 price changes per Day, and 8 per Night. Thus, you have the option to change your price 16 times per cycle.

The next section will explain how to play and the screen you will see. We will then explain how profit is calculated. Finally, we explain how your final payment for today's session will be determined.

How to play

At the start of the experiment you will be asked to choose your starting price. Figure 1 shows the screen you will see once the experiment begins.

At the top right of the screen there is a table that shows current prices and profits.

In the middle there is a graph that plots the prices of you and Seller 2. Your price line is thicker, while Seller 2's is represented using a thin blue line.

To help you keep track of your current profit, the colour of your entire price line will change between black to light green. The darker the line, the lower your profit. The brighter green the line is, the higher your profit. Seller 2's line does not change colour on your graph. You can see their current profit in the table above the price graph.

The shaded blocks on the graph represent Night time, while the un-shaded areas represent Day time. You can use this to keep track of when Day/Night is approaching.

On the right of the screen there are buttons to change your price. Possible prices are between 1 and 20. You can change your price at any time, however it will only take effect once the next 2.5 second block begins. You can see your selected price below the table on the right. This is also the same for Seller 2. When there is a switch from Day to Night (or Night to Day), the prices and profits in the table and graph will also update automatically.

At the top left of the screen is a timer counting down from 60 minutes. Below that you can see your total profit for the previous cycle, the average of your profit for that 40 second interval.




Figure 1: Screenshot of the user interface

Profit

In this section we explain how to calculate profit, and how it is impacted by Day/Night and relative prices. Please note the two profit tables attached.

During the experiment, your profit will depend on your price, and the price set by Seller 2. Your total profit each cycle is the average profit during each 40 second interval. Profit Table 1 shows profits during the Day, while Profit Table 2 shows profits during the Night. You can think of this as there being fewer buyers around at Night, while there are more buyers during the Day.

How to use the profit table

You can use each profit table to determine your profit, depending on your price, and the price set by Seller 2. The column on the left is your price. As you look across each row, you can see how your profit changes with the price selected by Seller 2.

For example, if your price is 13 and Seller 2's price is 16 your profit is 1,114 ECU during the Day, and 332 ECU at Night.

You can also use the same process to figure out what Seller 2's profit would be, as you both are using the same table. Instead, now use the 'Your price' column to represent Seller 2.

If Seller 2's price is 16, and your price is 13, then Seller 2's profit is 210 ECU during the Day, and 17 at Night.

If your prices are the same, your profits are also the same.

For example, if you and Seller 2 both set your prices to 17, during the Day each of you has a profit of 789 ECU, while at Night your profits would both be 8 ECU.

Understanding prices and profit

You and Seller 2 make similar products. This means that buyers are more likely to purchase from the seller with the lower price. Your profit depends on the amount you can sell (your price, *relative* to Seller 2), and how much you earn from selling each unit (your price):

- If you set a low price, you will sell more units, but your profit won't be high as you receive a low price per unit sold.
- As you start increasing your price, you may sell fewer units, but overall you earn more profit due to the higher price per unit.
- However, after a certain point the higher price is outweighed by a decrease in units sold - so your profit will decrease.
- To see this, look down the profit tables to see how your profit changes as you increase price from 1 to 20, for any given price set by Seller 2.

Although more buyers will purchase from the seller with the lower price, if you both set your prices higher you may both earn higher profit levels:

- For example, if you both set a lower price (i.e. 3), you both make 149 ECU. If your both increase your price to 10, you each have a profit of 499 ECU during the Day, and 472 ECU at Night.

- You can see how Seller 2's price impacts your profit by looking across the profit tables from left to right, for any price you choose. Your profit increases if Seller 2's price increases. This works the same for Seller 2 - if your price increases, Seller 2's profit increases.

Therefore, when choosing your price to earn as much profit as possible, consider your price relative to that of Seller 2, and the profit that each of you earn at those prices.

Payment

After the final day, a computer will randomly select 10 intervals of 40 seconds from the experiment. Each interval is equally likely to be chosen. You will only be paid for the average profit you made during these 10 intervals. The selected intervals and profits in ECU will be shown on your screen once the experiment ends.

Profit table provided to participants:

Profit Table 1 - Day (ECU)

		Seller 2's price																			
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	49	65	77	86	92	95	97	98	99	99	99	99	99	99	99	99	99	99	99	99	99
2	69	99	130	155	173	184	191	195	197	198	199	199	199	199	199	199	199	199	199	199	199
3	66	104	149	195	233	260	277	287	293	296	297	298	299	299	299	299	299	299	299	299	299
4	53	89	139	199	260	310	346	369	383	390	395	397	398	399	399	399	399	399	399	399	399
5	37	66	111	174	249	325	388	433	462	478	488	493	496	498	498	498	499	499	499	499	499
6	25	45	79	133	209	299	390	466	520	554	574	586	592	595	597	598	599	599	599	599	599
7	16	29	53	93	155	244	349	455	543	606	646	670	683	691	695	697	698	699	699	699	699
8	9	18	33	60	106	178	278	399	520	621	693	738	765	781	789	794	796	798	798	799	799
9	6	11	20	37	68	119	200	313	449	585	699	779	830	861	878	887	893	895	897	898	898
10	3	6	12	22	42	75	132	222	348	499	650	776	865	922	956	975	985	991	994	994	996
11	2	3	7	13	25	46	83	146	244	383	549	714	852	950	1,013	1,050	1,070	1,082	1,088	1,088	1,092
12	1	2	4	8	14	27	50	90	159	266	417	597	778	927	1,034	1,102	1,142	1,164	1,177	1,184	1,184
13	0	1	2	4	8	16	29	54	98	172	288	451	645	839	1,000	1,114	1,187	1,230	1,254	1,268	1,268
14	0	0	1	2	5	9	17	32	58	106	185	310	484	691	898	1,068	1,189	1,266	1,311	1,337	1,337
15	0	0	0	1	2	5	10	18	34	63	113	198	330	515	733	949	1,127	1,252	1,332	1,378	1,378
16	0	0	0	0	1	3	5	10	19	36	67	120	210	349	542	768	989	1,169	1,295	1,374	1,374
17	0	0	0	0	0	1	3	6	11	21	38	71	127	221	366	562	789	1,006	1,180	1,300	1,300
18	0	0	0	0	0	0	1	3	6	12	22	41	74	133	230	376	570	787	988	1,144	1,144
19	0	0	0	0	0	0	1	1	3	6	12	23	43	78	138	235	378	558	749	917	917
20	0	0	0	0	0	0	0	1	2	3	7	13	24	44	80	141	234	364	517	666	666

Your price

Profit Table 2 - Night (ECU)

		Seller 2's price																			
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	49	74	89	95	98	99	99	99	99	99	99	99	99	99	99	99	99	99	99	99	99
2	51	99	148	178	191	197	198	199	199	199	199	199	199	199	199	199	199	199	199	199	199
3	32	77	149	222	267	287	295	298	299	299	299	299	299	299	299	299	299	299	299	299	299
4	16	43	103	199	296	356	383	394	397	399	399	399	399	399	399	399	399	399	399	399	399
5	7	20	54	129	249	370	445	479	492	497	498	499	499	499	499	499	499	499	499	499	499
6	3	8	24	65	155	299	444	533	574	590	595	597	598	598	598	598	598	598	598	598	598
7	1	3	10	28	75	180	349	516	621	668	686	692	695	695	696	696	696	696	696	696	696
8	0	1	4	11	32	86	206	397	586	703	756	776	784	786	787	788	788	788	788	788	788
9	0	0	1	4	13	36	97	230	440	646	772	829	851	859	861	862	863	863	863	863	863
10	0	0	0	1	5	14	40	107	250	471	679	804	858	879	887	889	890	891	891	891	891
11	0	0	0	0	1	5	16	44	115	261	468	647	747	790	806	812	814	815	815	815	815
12	0	0	0	0	0	2	6	17	47	117	246	400	510	565	587	595	598	599	599	599	599
13	0	0	0	0	0	2	6	18	47	107	193	267	308	326	332	335	335	336	336	336	336
14	0	0	0	0	0	0	0	2	6	18	42	80	115	137	146	150	151	151	151	151	151
15	0	0	0	0	0	0	0	0	2	6	16	31	45	54	58	60	61	61	61	61	61
16	0	0	0	0	0	0	0	0	0	2	6	11	17	20	22	23	23	23	23	23	23
17	0	0	0	0	0	0	0	0	0	0	2	4	6	7	8	8	8	8	8	8	8
18	0	0	0	0	0	0	0	0	0	0	0	0	2	2	3	3	3	3	3	3	3
19	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
20	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Your price

Bibliography

- Abbink, K. and Brandts, J. (2008). Pricing in bertrand competition with increasing marginal costs. *Games and Economic Behavior*, 63(1):1–31.
- Alós-Ferrer, C. and Ania, A. (2005). The evolutionary stability of perfectly competitive behavior. *Economic Theory*, 26(3):497–516.
- Anderson, L., Freeborn, B., and Holt, C. (2010). Tacit collusion in price-setting duopoly markets: Experimental evidence with complements and substitutes. *Southern Economic Journal*, 76(3):577–591.
- Anderson, S. P. and de Palma, A. (1992). The logit as a model of product differentiation. *Oxford Economic Papers*, 44(1):51–67.
- Anderson, S. P., Erkal, N., and Piccinin, D. (2020). Aggregative games and oligopoly theory: short-run and long-run analysis. *The RAND Journal of Economics*, 51(2):470–495.
- Becker, M., Pfeifer, G., and Schweikert, K. (2021). Price effects of the austrian fuel price fixing act: A synthetic control study. *Energy Economics*, 97:105207.
- Berninghaus, S., Hesch, M., and Hildenbrand, A. (2012). Zur wirkung regulatorischer preiseingriffe auf dem tankstellenmarkt. *Wirtschaftsdienst*, 92(1):46–50.
- Bertrand, J. (1883). Review of ‘theorie mathematique de la richesse sociale and of recherches sur les principes mathematiques de la theorie des richesses’. *Journal des Savants*, pages 499–508.

- Bigoni, M., Casari, M., Skrzypacz, A., and Spagnolo, G. (2015). Time horizon and cooperation in continuous time. *Econometrica*, 83(2):587–616.
- Byrne, D. P. and de Roos, N. (2019). Learning to coordinate: A study in retail gasoline. *American Economic Review*, 109(2):591–619.
- Calvano, E., Calzolari, G., Denicolo, V., and Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10):3267–97.
- Caplin, A. and Nalebuff, B. (1991). Aggregation and imperfect competition: On the existence of equilibrium. *Econometrica*, 59(1):25–59.
- Chandrasekhar, A. G. and Xandri, J. P. (2014). A note on payments in the lab for infinite horizon dynamic games with discounting. *Working paper*.
- Chaudhuri, A. (2011). Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature. *Experimental Economics*, 14(1):47–83.
- Corchón, L. C. (1994). Comparative statics for aggregative games the strong concavity case. *Mathematical Social Sciences*, 28(3):151–165.
- Cournot, A.-A. (1838). *Recherches sur les principes mathématiques de la théorie des richesses par Augustin Cournot*. chez L. Hachette.
- 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.
- 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.

- Dewenter, R., Heimeshoff, U., and Lüth, H. (2017). Less Pain at the Pump? The Effects of Regulatory Interventions in Retail Gasoline Markets. *Applied Economics Quarterly (formerly: Konjunkturpolitik)*, 63(3):259–274.
- Dixon, H. D. (2000). Keeping up with the joneses: competition and the evolution of collusion. *Journal of Economic Behavior & Organization*, 43(2):223–238.
- Dufwenberg, M. and Gneezy, U. (2000). Price competition and market concentration: an experimental study. *International Journal of Industrial Organization*, 18(1):7–22.
- Duncan, L. R. (1959). *Individual Choice Behavior; a theoretical analysis*. New York: Wiley.
- Eckert, A. (2013). Empirical studies of gasoline retailing: A guide to the literature. *Journal of Economic Surveys*, 27(1):140–166.
- Feinberg, R. and Husted, T. (1993). An experimental test of discount-rate effects on collusive behaviour in duopoly markets. *Journal of Industrial Economics*, 41(2):153–60.
- Feinberg, R., Husted, T., and Waskiewicz, M. (2017). An experimental test of tradeoffs between discount rates and number of firms in supporting collusion. *Journal of Applied Business and Economics*, 19(6).
- Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.
- Fonseca, M. A. and Normann, H.-T. (2012). Explicit vs. tacit collusion—the impact of communication in oligopoly experiments. *European Economic Review*, 56(8):1759 – 1772.
- Fosgerau, M., Melo, E., de Palma, A., and Shum, M. (2020). Discrete choice and rational inattention: A general equivalence result. *International Economic Review*, 61(4):1569–1589.

- Friedman, D., Huck, S., Oprea, R., and Weidenholzer, S. (2015). From imitation to collusion: Long-run learning in a low-information environment. *Journal of Economic Theory*, 155:185–205.
- Friedman, D. and Oprea, R. (2012). A continuous dilemma. *American Economic Review*, 102(1):337–63.
- Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *The Review of Economic Studies*, 38(1):1–12.
- Fudenberg, D. and Maskin, E. (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3):533–554.
- Green, E. J., Marshall, R. C., and Marx, L. M. (2014). Tacit Collusion in Oligopoly. In *The Oxford Handbook of International Antitrust Economics, Volume 2*. Oxford University Press.
- Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125.
- Harsanyi, J. C. (1995). A new theory of equilibrium selection for games with complete information. *Games and Economic Behavior*, 8(1):91–122.
- Harsanyi, J. C. and Selton, R. (1988). *A general theory of equilibrium selection in games*. MIT Press, Cambridge, Mass.
- Haucap, J. and Müller, H. C. (2012). The effects of gasoline price regulations: Experimental evidence. DICE Discussion Papers 47, University of Düsseldorf, Düsseldorf Institute for Competition Economics (DICE).
- Heinemann, F., Nagel, R., and Ockenfels, P. (2009). Measuring strategic uncertainty in coordination games. *The Review of Economic Studies*, 76(1):181–221.
- Horstmann, N., Krämer, J., and Schnurr, D. (2016). Oligopoly competition in continuous time. *SSRN Electronic Journal*.

- Horstmann, N., Krämer, J., and Schnurr, D. (2018). Number Effects and Tacit Collusion in Experimental Oligopolies. *Journal of Industrial Economics*, 66(3):650–700.
- Hotelling, H. (1929). Stability in competition. *The Economic Journal*, 39(153):41–57.
- Huck, S., Leutgeb, J., and Oprea, R. (2017). Payoff information hampers the evolution of cooperation. *Nature Communications*, 8(15147).
- Huck, S., Müller, W., and Normann, H.-T. (2002). To commit or not to commit: Endogenous timing in experimental duopoly markets. *Games and Economic Behavior*, 38(2):240–264.
- Huck, S., Normann, H.-T., and Oechssler, J. (2004). Two are few and four are many: number effects in experimental oligopolies. *Journal of Economic Behavior & Organization*, 53(4):435–446.
- Jensen, M. (2010). Aggregative games and best-reply potentials. *Economic Theory*, 43(1):45–66.
- Klein, T. (2021). Autonomous algorithmic collusion: Q-learning under sequential pricing. *The RAND Journal of Economics*, 52(3):538–558.
- Leng, A., Friesen, L., Kalayci, K., and Man, P. (2018). A minimum effort coordination game experiment in continuous time. *Experimental Economics*, 21(3):549–572.
- Leufkens, K. and Peeters, R. (2011). Price dynamics and collusion under short-run price commitments. *International Journal of Industrial Organization*, 29(1):134–153. Special Issue: Experiments in Industrial Organization.
- Mailath, G. and Samuelson, L. (2006). *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press.

- Maskin, E. and Tirole, J. (1988). A theory of dynamic oligopoly, ii: Price competition, kinked demand curves, and edgeworth cycles. *Econometrica*, 56(3):571–599.
- McFadden, D. (1981). Econometric models of probabilistic choice. *Structural analysis of discrete data with econometric applications*, 198272.
- Potters, J. and Suetens, S. (2009). Cooperation in experimental games of strategic complements and substitutes. *Review of Economic Studies*, 76(3):1125–1147.
- Potters, J. and Suetens, S. (2013). Oligopoly experiments in the current millennium. *Journal of Economic Surveys*, 27(3):439–460.
- Samuelson, W. and Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of risk and uncertainty*, 1:7–59.
- Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard University Press.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological review*, 63(2):129.
- Simon, L. K. and Stinchcombe, M. B. (1989). Extensive form games in continuous time: Pure strategies. *Econometrica*, 57(5):1171–1214.
- Sitzia, S. and Zheng, J. (2019). Group behaviour in tacit coordination games with focal points – an experimental investigation. *Games and Economic Behavior*, 117:461–478.
- Stigler, G. J. (1964). A theory of oligopoly. *Journal of Political Economy*, 72(1):44–61.
- Vives, X. (2001). *Oligopoly Pricing: Old Ideas and New Tools*, volume 1 of *MIT Press Books*. The MIT Press.
- Wang, Z. (2009). (mixed) strategy in oligopoly pricing: Evidence from gasoline price cycles before and under a timing regulation. *Journal of Political Economy*, 117(6):987–1030.

Weisbuch, G., Kirman, A., and Herreiner, D. (2000). Market organisation and trading relationships. *The Economic Journal*, 110(463):411–436.